

A five-year milestone: reflections on advances and limitations in GeoAI research

Yingjie Hu, Michael Goodchild, A-Xing Zhu, May Yuan, Orhun Aydin, Budhendra Bhaduri, Song Gao, Wenwen Li, Dalton Lunga & Shawn Newsam

To cite this article: Yingjie Hu, Michael Goodchild, A-Xing Zhu, May Yuan, Orhun Aydin, Budhendra Bhaduri, Song Gao, Wenwen Li, Dalton Lunga & Shawn Newsam (29 Jan 2024): A five-year milestone: reflections on advances and limitations in GeoAI research, Annals of GIS, DOI: [10.1080/19475683.2024.2309866](https://doi.org/10.1080/19475683.2024.2309866)

To link to this article: <https://doi.org/10.1080/19475683.2024.2309866>



© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group, on behalf of Nanjing Normal University.



Published online: 29 Jan 2024.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)

A five-year milestone: reflections on advances and limitations in GeoAI research

Yingjie Hu^a, Michael Goodchild^b, A-Xing Zhu^c, May Yuan^d, Orhun Aydin^{e,f}, Budhendra Bhaduri^g, Song Gao^c, Wenwen Li^h, Dalton Lunga^g and Shawn Newsamⁱ

^aDepartment of Geography, University at Buffalo, Buffalo, NY, USA; ^bDepartment of Geography, University of California, Santa Barbara, CA, USA; ^cDepartment of Geography, University of Wisconsin-Madison, Madison, WI, USA; ^dSchool of Economic, Political and Policy Sciences, University of Texas Dallas, Dallas, TX, USA; ^eDepartment of Earth and Atmospheric Sciences, Saint Louis University, Saint Louis, MO, USA; ^fTaylor Geospatial Institute, Saint Louis, MO, USA; ^gOak Ridge National Laboratory, Oak Ridge, TN, USA; ^hSchool of Geographical Sciences and Urban Planning, Arizona State University, Tempe, AZ, USA; ⁱDepartment of Computer Science and Engineering, University of California, Merced, CA, USA

ABSTRACT

The Annual Meeting of the American Association of Geographers (AAG) in 2023 marked a five-year milestone since the first Geospatial Artificial Intelligence (GeoAI) Symposium was held at AAG in 2018. In the past five years, progress has been made while open questions remain. In this context, we organized an AAG panel and invited five panellists to discuss the advances and limitations in GeoAI research. The panellists commended the successes, such as the development of spatially explicit models, the production of large-scale geographic datasets, and the use of GeoAI to address real-world problems. The panellists also shared their thoughts on limitations in current GeoAI research, which were considered as opportunities to engage theories in geography, enhance model explainability, quantify uncertainty, and improve model generalizability. This article summarizes the presentations from the panellists and also provides after-panel thoughts from the organizers. We hope that this article can make these thoughts more accessible to interested readers and help stimulate new ideas for future breakthroughs.

ARTICLE HISTORY

Received 14 June 2023
Accepted 7 January 2024

KEYWORDS

GeoAI; theory; perspectives;
GIScience; geography

1. Introduction

Geospatial artificial intelligence (GeoAI) is an interdisciplinary field related to a wide range of disciplines, such as geography, GIScience, computer science, data science, remote sensing, Earth system science, urban planning, civil engineering, and public health. With the potential to advance solutions to societal challenges, GeoAI research has received tremendous attention from both academia and industry (Hu, Li, et al. 2019; W. Li 2020; Gao 2021; Chiappinelli 2022). In the Annual Meeting of the American Association of Geographers (AAG) in April 2018, we organized the first AAG GeoAI Symposium to facilitate the exchange of ideas and promote research on this topic. A series of GeoAI Symposia were organized in AAG in the following years. Another thread of activities that were held in parallel is the GeoAI Workshop series that we organized at the Association for Computing Machinery (ACM) Special Interest Group on Spatial Information (SIGSPATIAL) conferences, with the first ACM SIGSPATIAL GeoAI Workshop held in November 2017 (Hu, Gao, et al. 2019; Lunga et al. 2022; Mao et al. 2018). The AAG GeoAI Symposium series and the SIGSPATIAL GeoAI Workshop series complement each other in that the former attracts

more participants from the geography community while the latter draws more participants from the computer science community. These symposia and workshops, along with many other efforts such as special issues in the *International Journal of Geographical Information Science* (Janowicz et al. 2020), *Transactions in GIS* (Mai, Hu, et al. 2022), and *Geoinformatica* (Gao et al. 2023), have contributed to the remarkable growth of GeoAI research in recent years.

The Annual Meeting of AAG in 2023 marked a five-year milestone since the first AAG GeoAI Symposium in 2018. In the past five years, we have witnessed many advances in using GeoAI to address societal challenges, such as those related to disaster resilience (Peng et al. 2020; B. Zhou et al. 2022), public health (M. N. K. Boulos, Peng, and VoPham 2019; R. Z.; B. Zhou et al. 2022), ecosystem conservation (Ma et al. 2022; Nguyen, Kellenberger, and Tuia 2022), terrain analysis (W. Li and Hsu 2020; W. Li, Hsu, and Hu 2021), humanitarian mapping (Gaikwad et al. 2022; Lunga et al. 2018), geoprivacy protection (K. Boulos et al. 2022; Rao et al. 2021), transportation management (M. Li et al. 2019; Y. Zhang et al. 2020), and smart cities and communities (R. Wang and Biljecki 2022; R. Zhang et al. 2021). In addition,

CONTACT Yingjie Hu  yhu42@buffalo.edu  Department of Geography, University at Buffalo, Buffalo, NY, USA

© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group, on behalf of Nanjing Normal University.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

methodological advances were also made to increase model explainability (Cheng et al. 2021; Hsu and Li 2023), improve spatial interpolation (D. Zhu et al. 2020), and develop spatially explicit models (Gupta et al. 2021; Mai et al. 2022). Meanwhile, various limitations remain, such as how to better incorporate human feedback (Zheng and Sieber 2022) and better engage and advance theories in geography (Janowicz, Sieber, and Crampton 2022).

In this context, we organized a panel in the 2023 AAG GeoAI Symposium, and invited five panellists to discuss the progress made in GeoAI research so far and potential future directions. Before the panel session, we sent four questions to the panellists, which served to guide the conversation and provide a general theme for the panel. The four questions were:

- Which areas do you think GeoAI research has excelled in over the past five years?
- Which areas do you think GeoAI research has fallen short or could benefit from further improvement?
- What advancements in GeoAI research would you like to see in the next five years?
- Any other thoughts you would like to share?

In the following, we summarize the thoughts shared by the five panellists and also those from the organizers after the panel. We note that this paper is intended to be a *position paper* or *comment paper*, in which we summarize the thoughts shared by scholars from different perspectives and present these thoughts to more researchers, especially those who cannot attend the 2023 AAG GeoAI panel session in person. For a comprehensive survey of the GeoAI literature, interested readers may refer to review articles, such as Gao (2021), Zhu et al. (2017), Liu and Biljecki (2022), and Yuan et al. (2020).

2. Thoughts from the panellists

2.1. Thoughts from Michael Goodchild

AI has exploded into the public consciousness over the past five years, most recently with the public furore over ChatGPT, and GeoAI has seen similar growth. But while the term 'GeoAI' may date from only about five years ago, interest in the applications of AI to geographic research has a much longer history. Dobson was writing as early as 1983 about what he termed 'automated geography' (Dobson 1983), and in 1986 the National Science Foundation (NSF) identified 'artificial intelligence and expert systems' as one of the five research areas of its proposed National Center for Geographic

Information and Analysis (Abler 1987). Openshaw and Openshaw's book *Artificial Intelligence in Geography* (Openshaw and Openshaw 1997), collecting much of the extensive work on AI by the authors and others, appeared in 1997 and according to the publisher's blurb marked 'the beginning of the AI revolution in geography'. To be sure, the term 'GeoAI' was coined and came into widespread use only about five years ago, but it captured and helped to integrate what had already accumulated as a large and extensive body of work.

So how should we understand this recent surge of interest in GeoAI? Is the coining of the term perhaps having a stimulating effect similar to that of the coining of 'GIS' in the late 1960s? Or have we reached some threshold in the availability of massive quantities of geospatial data, or high-performance computing power, or ready-to-use software? Or is something much more fundamental at work?

Dobson, the Openshaws, and other early enthusiasts were clearly seeing AI as a means of advancing spatial analysis, by enabling the analysis of larger data sets at a time when computing resources were severely constrained, by permitting researchers to examine large numbers of competing hypotheses, or by automating repetitive human activities. The fundamental objectives of spatial analysis remained as they had been since the early days of the quantitative revolution: a search for understanding and explanation, and the discovery of the principles by which the human and environmental worlds operated. The basic tenets of science, including reproducibility and replicability, openness, and the nomothetic search for simple processes that apply everywhere and at all times, remained firmly in place.

But GeoAI as it is understood today seems to me to be causing a major disruption, by questioning and augmenting this simple sense of purpose that had endured for more than half a century. I would like to give three examples. First, much of the excitement over AI from 10 or so years ago concerned its apparent success in predicting the stock market, the weather, and outbreaks of flu. Yet, prediction has never been a central goal of science, and critics labelled such applications as 'curve-fitting' (Pearl and Mackenzie 2018), a general lack of transparency and generalizability, and of little interest to science. Second, many applications of GeoAI focus on an operation that often appears in GIS functionality as 'find similar', the ability to search across space and sometimes time for instances that match a feature of interest. This is no more than an extended version of classification, powered often by high-performance computing, and very useful in classifying massive numbers of remotely sensed pixels. But conceptually it is a very simple

operation that is not likely to offer what scientists have traditionally thought of as explanation or understanding. Finally, the use of generative AI to create ‘deep fakes’, misinformation and disinformation, and various kinds of imaginary landscapes is certainly a powerful and interesting area of GeoAI, but hardly one that advances the traditional aims of geographic science.

These developments disrupt what was for more than half a century a community consensus on the purpose of spatial analysis and quantitative geography. They add new purposes, some that open commercial opportunities, others that present intellectual challenges for any geographic scientist, and some that pose important ethical questions. As such, they help us to understand what has happened in GeoAI in the past five years. It is important that we rethink the nature of science in the era of GeoAI, and the new purposes that underlie the culture of geographic science and will drive further developments in GeoAI.

2.2. Thoughts from A-Xing Zhu

My thoughts and comments will focus on the first and third questions, i.e. advances and challenges in GeoAI. Around 2015, I was asked to write a piece on artificial neural networks (ANN) in geography for *The International Encyclopedia of Geography*. In the summary section of that paper (Zhu 2017), I wrote a paragraph about three limitations in the applications of ANNs in geography at that time. First, it is expensive to obtain a large set of representative samples to train an ANN model. Second, the heterogeneity of geographic phenomena makes it difficult to generalize from one application to another. Third, it is difficult to understand the relationships between the input and output variables using an ANN model. ANN is a very small part of AI and GeoAI. One more point that was not included in that paper, which I would like to add here, is incorporating spatial structure of the input data as part of the training process. My comments are organized around these four points with the first two points focusing more on advances and the next two points focusing more on challenges.

First, we have made advances in collecting sample data for training AI models. As geographers, we know that it is time-consuming and labour-intensive to collect a large set of representative samples. One advance that we have made to address this difficulty is the utilization of volunteered geographic information (VGI) (Goodchild 2007), in which citizens contribute data on phenomena that they observe in the field, especially those phenomena that may vanish very quickly. The result of VGI is a vast amount of data that can be used to train models

through supervised, weakly-supervised, or unsupervised learning. Meanwhile, VGI can have the issues of spatial bias that may affect the trained models. We have also made advances in this direction using similarity models to mitigate spatial bias in VGI and improve the trained model (G. Zhang and Zhu 2020). Spatial bias may be less of a concern for image-based machine learning problems, as sampling from remote sensing images would not suffer from spatial bias as much.

The second point related to advances is the incorporation of spatial structure. In this thread, it is great to see the works from Yan et al. (2017) and Li et al. (2021) who incorporated spatial structures into GeoAI models. Those spatial structures include not only spatial distance but also spatial arrangement that is highly important in the genesis of geographic phenomena. Spatial arrangement also largely affects the development and evolution of geographic phenomena. The incorporation of spatial structures, and the use of VGI to supplement training data are the two advances that I think we have made significant progress in, although there is certainly room for further improvements.

The third point, which is more related to challenges, is complexity tied to spatial heterogeneity. In their paper, Hu et al. (2019) discussed the challenge of model generalizability in GeoAI research, i.e. the difficulty for a model trained based on the data from one geographic area to be generalized to other geographic areas. The lack of model generalizability is an important issue for geographic studies, because many geographic phenomena are highly complex, and their spatial relationships vary across different geographic areas. Currently, our models are more like ‘average’ models that learn based on the frequency of data; meanwhile, we often have more training data from some geographic areas and less or no training data from some other geographic areas. Consequently, the model learns more about the geographic phenomena in areas where training data are more abundant. In addition, some phenomena happen less frequently (thus less training data), but they are still quite important. Thus, a challenge is how to train models so that they can still work in the geographic areas with less training data or for the geographic phenomena that happen less frequently. A possible solution is using the so-called similarity model (A.-X. Zhu and Turner 2022): instead of using training samples simply based on their frequencies, we can use them based on their uniqueness. To do so, we may need to first screen the training data to determine their uniqueness and contribution to the understanding of the particular problem.

Last but not least, model interpretability is still a major challenge in GeoAI research. By default, most machine learning models learn the relationships

between the input and output variables in the form of model weights. Those weights are certainly important for the models to make predictions, but they do not provide much information for us to interpret those relationships. What can we do to make the results from machine learning more interpretable? This can be a challenge for GeoAI research to address in the coming years.

2.3. Thoughts from May Yuan

Artificial intelligence has been advancing very fast since the early 1990s, with progress made in models, such as artificial neural networks and convolutional neural networks, and algorithms, such as backpropagation developed to solve weights in multiple layers of neurons. Based on my role serving as a programme officer in NSF, the funding agency has been very supportive for AI-based research. If we search the keyword 'AI' in the NSF award database, we can find over 3,000 funded projects spanning through all directorates in NSF. Geography is under the NSF directorate of Social, Behavioral and Economic Sciences (SBE), and there are also many AI-based projects funded by NSF under SBE. How about funded projects related to GeoAI? A keyword search on 'GeoAI' in the NSF award database returned 10 awards. This is a relatively small number, and we would hope to have more GeoAI proposals to SBE and the NSF programmes that I serve, i.e. the Human-Environment and Geographical Sciences (HEGS) programme, the Human Networks and Data Science (HNDS) programme and the CyberTraining programme, all of which are very supportive for AI-based research. As of March 2023, HEGS programme funded three GeoAI-related projects, but we declined many more GeoAI proposals. So, why were certain proposals funded while some others were not? The key difference is whether the proposal is merely a data-driven exercise, i.e. simply applying some analyses to some data, or it involves innovative GeoAI approaches to answer important research questions. Proposals in the second group are more likely to be funded.

As the Editor-in-Chief of the *International Journal of Geographical Information Science* (IJGIS), I have been reading GeoAI-related paper submissions almost every day, and probably 95% of the submissions to IJGIS have some references to machine learning. The biggest challenge for me is to learn the growing acronyms that authors fashionably make in many manuscripts. However, if we really unpack those fancy acronyms, most of them appear as repackaging algorithms already established elsewhere with geographic coordinates or geographic variables as additional input.

These submissions are neither innovative nor contributing to our understanding of geography: nothing wrong but nothing exciting. Figure 1 shows a word cloud visualization based on the titles of the papers accepted or published in IJGIS between January 2022 and March 2023. As we can see, there are a lot of papers related to 'learning', 'method', and 'model'; meanwhile, the biggest word in this word cloud is 'using' which probably suggests many of our works are mostly using others' algorithms. In sum, I think the GeoAI community has been trying to get up to speed in the past five years with trendy computational memes. It is time to attend to learning geography, incorporating geography into learning algorithms, and promoting curiosity-driven GeoAI research.

For the areas that GeoAI research has excelled in, I think we have successfully applied an expansive suite of machine learning algorithms to a wide range of structured and unstructured data in diverse domains across macro and micro spatial scales. However, most research that I was exposed to is limited to extracting facts, e.g. extracting spatial features and distribution patterns. While these facts are important, this line of research is like using data to create more data. Although we may think that we are discovering knowledge, we are in fact formulating existing data into new data from another perspective, or extracting higher level data from lower level data. For example, we have a lot of micro-location data (e.g. cell phone data and GPS data), and we may extract spatiotemporal trajectories and human movements from such data; we may further contextualize the extracted trajectories and movements to look at the geographic environments, transportation modes, and human activities. This is in essence a data assembling process from lower level data (i.e. micro-location data) to higher level data (i.e. trajectories, movement flows, community networks, and community zones), or aggregating location data across different spatial scales. To some extent, this type of research can be considered as acquiring additional data computationally, and only a small number of papers contributed new insights to improving our understanding of geography. Another area that GeoAI research has excelled in is treating spatial heterogeneity more effectively and more explicitly compared with research five years ago, and I think that is a major contribution of GeoAI.

For the areas that GeoAI research has fallen short, one area is model explainability. Our current research mostly focuses on the importance of input features for explainability, i.e. how important an input feature is for a model to make predictions, rather than the function or mechanism through which an input variable affects the model prediction. Researchers have already used



Figure 1. A word cloud visualization based on the titles of papers published or accepted by IJGIS between January 2022 and March 2023.

frameworks, such as SHAP (Shapley Additive Explanations), to improve our ability to explain the effects of input variables on model predictions, but the mechanistic processes of input variables and their interactions to generate model output remain difficult, if not impossible, to discern. Also, many GeoAI papers lack theories or hypotheses and are focused on finding patterns from data or making predictions, not leading to any new understanding. It is time to look into novel conceptual frameworks and general principles to make our results generalizable beyond the study area and case datasets. In this way, we create new knowledge rather than summarizing what we already know in a different aggregated level. In addition, I think that the AI learning paradigm has improved significantly over the years, from supervised learning, unsupervised learning, and transferred learning to reinforcement learning, and recently, we see much progress in self-supervised learning and transformers for natural language processing (NLP). While both supervised learning and reinforcement learning are constrained by the training data that are expensive to obtain, self-supervised learning can learn continuously and is less constrained by training data. Self-supervised learning may have important meaning to geography: due to spatial heterogeneity, an algorithm should not be fixed to a particular geographic location;

when it is applied to other locations, the algorithm should be able to continue to learn in order to become effective in those new locations. Similarly, the idea of transformers may be adaptable to modelling geographic predictions over space and time.

Finally, there are three statements that I would like to share to help stimulate our thinking. First, *the past, the present, and the future no longer help understand each other*. Palaeontologists consider the past is the key to the future, while geologists assume the present is the key to the past. Model training basically operates on such a uniformitarian principle that uses data collected in the past to optimize parameters for improved predictions with data new to the model. However, with climate change, many geographic factors, relationships, and processes will change; in other words, the present may no longer be the key to understanding the past, nor may the past be the key to understanding the future. GeoAI models trained with and learned from the past data are unlikely to predict the present or the future well. After all, understanding, not prediction, is at the core of science. The overwhelming attention of GeoAI research on prediction does not advance science. Noteworthy, computational learning theories, such as VC theory by Vapnik and Chervonenkis (2015) and algorithmic learning theory by Gold (1967) set the foundation for

learnability in machine learning. What are theoretical foundations for machines to learn geography? Merely learning from data, GeoAI can quickly reach the limits of predictions in geography. Second, *the cause may no longer occur before the effects*. When looking at causal relationships, we often think that the cause has to occur before the effects. However, effects can be observed before the cause, known as backward causation in philosophy or retrocausality in quantum physics. Neuroscientists have also observed that the downstream neurons fire before the upper stream neurons. How about causality in geography? Understanding mechanistic and functional relationships among variables, processes and geographic contexts is essential to affirm causality. Can GeoAI progress from finding facts to explaining geography and discerning geographic causality? Finally, *seeing may no longer be believing*. This is probably rather straightforward given the many fake images and fake videos generated by AI nowadays. As distrust grows with increases in fake data, science is at risk. What kinds of GeoAI algorithms may help detect

fake geographic data or estimate the level of data validity? I hope that these three statements can help us think more creatively about future GeoAI research.

2.4. Thoughts from Orhun Aydin

My thoughts on the progress of GeoAI are more from the perspectives of computer science theory and geoscience applications. In my view, the major advances of GeoAI in the past five years or so can be seen in three dimensions: data representation, statistical learning, and deep learning (Figure 2a). In terms of data representation, there are big wins powered by GeoAI. For example, sensor fusion has been instrumental for unlocking many remote sensing workflows; data encoding has enabled many downstream prediction tasks; and knowledge graphs have helped enrich representing spatial data in spatial learning tasks. There is also progress in exploratory methods, which I view as the intersection of statistical learning and data representation. These exploratory methods enable us to find patterns and also anomalies in data, and may

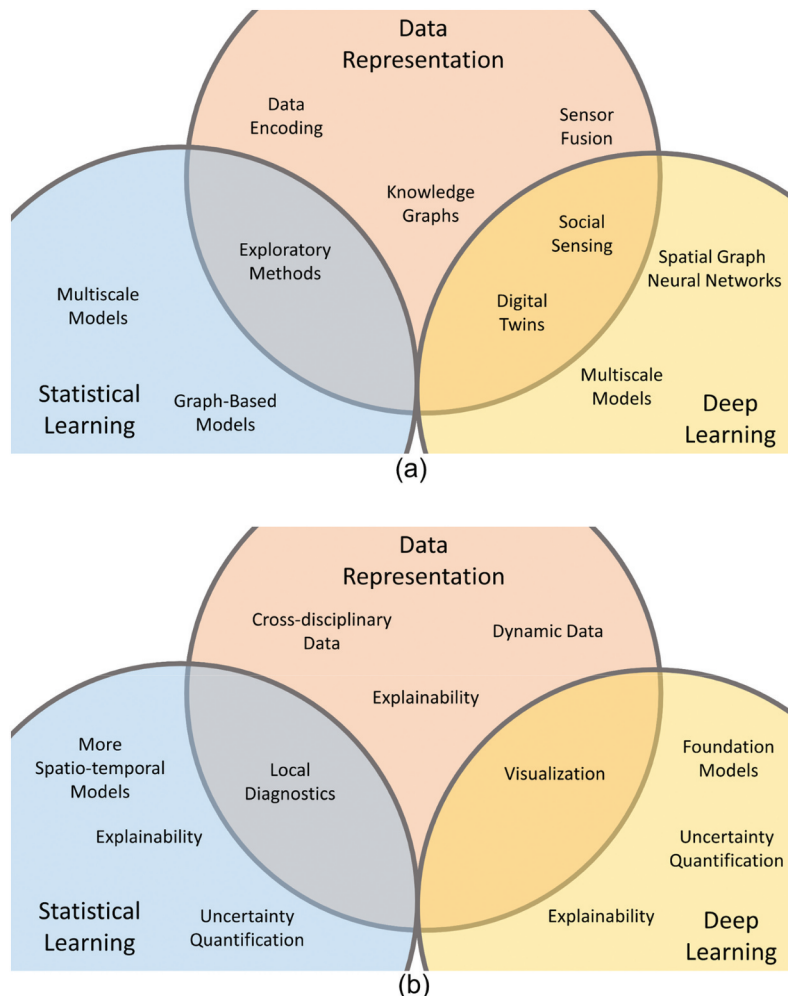


Figure 2. Advances and limitations in GeoAI research in the view of Orhun Aydin: (a) advances; (b) limitations.

motivate us to collect more data for our studies. In the area of statistical learning, there are remarkable advances in multiscale models, especially in the realm of geographically weighted regression (Fotheringham, Yang, and Kang 2017; Z. Li et al. 2020). Much progress was also made in graph-based models for prediction, clustering, and classification tasks. Graph models provide a rigorous mathematical and statistical framework to represent spatial relations based on geographic distances, virtual relationships, and others. Meanwhile, the development of spatially explicit models (Janowicz et al. 2020) makes space and spatial relations an inherent part of graph models. For the area of deep learning, there has been much progress in convolutional neural networks and image classification beyond the past five year time-frame of our focus here. Within the past five years, impressive progress was also made in multiscale models, spatial graph neural networks, and their applications in digital twins and social sensing.

In terms of what we can do better, we can still look at the three main dimensions, i.e. data representation, statistical learning, and deep learning (see Figure 2b). Model explainability is a critical issue and I have put it in multiple places in this figure. This is a big challenge, and increasing model explainability also requires model visualization. For data representation, the data in GeoAI may involve not only quantitative data but also data from other disciplines (e.g. qualitative data). Handling various types of data in mixed approaches may be an area of growth in the coming years. Also, uncertainty quantification is a challenge for both statistical and deep learning. We know that uncertainty exists in our data and models; yet, predictions from many machine learning models, such as random forests, are single values. We need better abilities to quantify, visualize, and communicate the uncertainties involved in those predictions output by our models.

In terms of exciting future directions, I think uncertainty is one important direction to pursue. When we start from raw observations (e.g. surveys and satellite imagery), we start with direct, but noisy, measurements. However, we often further process this raw data, e.g. by creating geographic indexes from the data, and the uncertainty increases as the data go through different levels of processing. The processed data is used in spatial analysis via numerical methods. Models also have their own uncertainties and assumptions. When the uncertainties of data meet the uncertainties of models, uncertainty propagates. If we are working on grand societal challenges, we also need to couple models of different systems, which further propagates the uncertainty. I believe that GeoAI plays vital roles in addressing today and tomorrow's societal challenges, and it is

important to improve our ability to quantify, visualize, and communicate uncertainty involved in our data, models, and processes.

Finally, the recent development of ChatGPT has triggered both excitement and concerns. The building block of ChatGPT is the transformer model (Vaswani et al. 2017) which has a clear structure with open-source code. GPT-1 still has a quite transparent model structure (Radford et al. 2018), but the latter models are more opaque, such as GPT-2 (Radford et al. 2019), GPT-3 (Brown et al. 2020), and GPT-4 (OpenAI 2023), probably due to their high business values to their companies. Such opaqueness is dangerous for scientific research, and we should try to keep our GeoAI models transparent while avoiding the GPT-4 route. I would like to conclude with a comparison between state-of-the-art NLP models such as ChatGPT-4 and an early model from robotics known as the Vaucanson's Duck. Jacques de Vaucanson created this mechanical duck that could walk and eat like a real duck, a very early version of artificial intelligence. He created the duck in the 1700s, but did not reveal his design of the duck to anyone. An inventor who lived about 100 years after Vaucanson realized that the duck was built significantly different from how Vaucanson advertised it. I hope that the artificial intelligence that has seen massive development in the past decade due to its strength in open science does not go in a full circle, and the future GeoAI research will not be like the new Vaucanson's duck.

2.5. *Thoughts from Budhendra Bhaduri*

There is an anecdote about the name of 'GeoAI'. In the mid-2010s, I had a distinguished research fellow, Dr. Huina Mao, working at the [anonymous organization]. She was passionate about AI and did nice dissertation research on predicting the stock market using Twitter (Bollen, Mao, and Zeng 2011). She, along with Yingjie Hu, Song Gao, and several other colleagues, wrote a workshop proposal on AI and geospatial, and submitted it to the ACM SIGSPATIAL conference in early 2017. The proposal came back with a positive review but the organizers said they felt that the workshop name was not catchy enough. We had a discussion on this name issue, and we eventually decided to put 'Geo' in front of 'AI' to indicate the integration of geospatial and AI. That is how the first GeoAI Workshop and the name 'GeoAI' came out in the ACM SIGSPATIAL conference in November 2017 (Mao et al. 2018), and many of you in this community have made nice contributions to that workshop.

Regarding the successes of GeoAI research in the past five years, one major success is the production of the

incredible amount of rich geographic data created at the global scale, which did not exist before. Those rich data can be largely attributed to the development of GeoAI. We have done remarkable work in object recognition and functional mapping based on remote sensing images. We have also explored temporal sequential data to uncover landscape processes, such as agricultural and urbanization trends. Another success of GeoAI is that it brings the geography community closer to the computing community. We have seen good integrations of geographical research with high-performance computing, including some of the largest machines in the United States such as Blue Waters and Roger. We have also seen trends in industry that integrate computing and geography, such as Google Earth Engine and Microsoft Planetary Computer. Many applications, including disaster response and human security, have significantly benefited from having the rich geographic data and the ability to process those data at a fast pace.

Regarding limitations in GeoAI, one limitation in my view is that a majority of our research is still focused on image processing. This limitation is probably due to data availability, i.e. there exists a vast amount of satellite and airborne imagery (e.g. the Landsat archive) covering the entire planet. So far, we have not seen much research on vector data but there is potential. A second limitation is that our GeoAI models are largely inflexible. They are often sensitive to human labelling, and we have made only limited progress in automated data labelling. It is difficult for a model to generalize from one source of images to another, which often have different spectral features. Models are also fragile to the varied image

quality from different data sources. Third, current GeoAI models lack higher level reasoning to help us interpret the results. This is related to model explainability. Here is one example from our own human settlement mapping work at ORNL. Figure 3(a) shows different types of neighbourhoods derived by a GeoAI model from satellite images where different neighbourhoods seem to have general correspondence with size of structures. However, Figure 3(b) shows one confusing result output by the trained model. After a careful examination of the result, we still could not explain this result since little difference in settlement structures or texture is observed between the two regions. It is difficult to know whether this separation is correct at this moment. It could be a misclassification due to the limited generalizability of the trained model, or the classification could be reasonable given the existence of some more structured buildings in the upper right of the blue region. As the current model does not provide an explanation of how it made such a classification, it is difficult to assess and interpret this result.

For future research opportunities, there has been much interest in foundation models recently. A GeoAI foundation model may be pre-trained with large amounts of images and other data from different sources (e.g. different satellites and sensors), which may mitigate expensive sampling requirements for particular tasks. In addition, we may also want to further understand uncertainties in data and models, and address adversarial threats such as fake images generated by AI. Future GeoAI research may also focus on edge deployment, i.e. deploying GeoAI models not in




Figure 3. An example of the difficulty in interpreting the results of human settlement mapping: (a) pink and blue colors representing two different types of neighborhoods where boundaries seem to follow building morphology; (b) a confusing separation boundary output by the trained model with no apparent change in settlement structures, which makes the result challenging to explain.

lab computers but in devices with limited-communication and low-power settings (e.g. drones). In this direction, our group at ORNL recently started experimenting with deploying object detection models in drones to check the status of electric poles after a hurricane. Another future opportunity is in forecasting and generating anticipatory scenarios. For example, can we accurately forecast the land use and land cover in the next 20 and 30 years? Finally, another possible opportunity that I have been thinking about is whether we can develop GeoAI models that can go beyond our human's sensing capability and our currently best sensor technologies. For example, can new GeoAI models predict the existence of certain geographic phenomena that cannot be captured by our human eyes or the best sensor technologies of our current time?

3. After-panel thoughts from the organizers

After the panel session, the organizers also shared their thoughts on important issues raised by the panellists. One important issue that has been commented on by all panellists is model explainability. While progress has been made in explainable AI, existing explainability research often focuses on explaining the prediction results of AI models, e.g. why an image classification model classifies an image to *cat* not *dog*, rather than the associations or causal relationships between input and output variables that much scientific research cares about. Furthermore, machine learning models often ingest a large number of input variables and disregard potential correlations among these variables. While this is probably fine from a prediction perspective, explainable AI may need to unpack the dependency among input variables. Fulfilling these needs may require efforts to explore existing explainable AI methods (e.g. rule-based, gradient-based, layer-based, spectral-based explanations), improve these methods, and eventually increase model explainability. Some nice efforts have been made by researchers in this direction (Cheng et al. 2021; Z. Li 2022; Hsu and Li 2023; Xing and Sieber 2023). In addition, we may also need to demonstrate the benefits of explainable GeoAI to the wider research community to increase community adoption. Why should researchers and decision makers use more complex GeoAI models, instead of the simpler and more widely accepted models (e.g. regression models)? Or what new insights can GeoAI models bring in for explanation beyond what existing models can offer? Answering these questions is necessary for demonstrating the value of explainable GeoAI and increasing its adoption.

Both methodological developments in GeoAI and using GeoAI methods to address societal challenges

have been commented on by the panellists. These two parts of GeoAI research should go hand-in-hand in the coming years. It seems that methodological developments, e.g. incorporating spatial structures and developing spatially explicit models, have already received good attention from the GeoAI community. Meanwhile, using GeoAI to address societal challenges is sometimes simply referred to as 'GeoAI applications'. When done properly, such research goes beyond merely 'applying' a model: it not only helps address real-world challenges but also reveals limitations of current GeoAI methods and informs future methodological development. In the recent initiative of NSF to establish National AI Institutes, the agency emphasizes the roles of both *foundational research* (focusing on developing methods and theories) and *use-inspired research* (focusing on situating methods under use cases that benefit society). We hope that GeoAI research follows a similar two-prong framework that forms a virtuous cycle in developing spatially explicit models (*foundational GeoAI research*), using the developed models to address real-world challenges and advance knowledge (*use-inspired GeoAI research*), and bringing the outcomes of use-inspired research back to inform methodological development.

Producing rich geographic data has been considered as both a success and a limitation of existing GeoAI research. On the one hand, research on producing data has important societal values for supporting decisions and enabling more research. Many U.S. federal agencies, such as the U.S. Census Bureau and the U.S. Geological Survey (USGS), have a mission to produce high-quality data, such as the socioeconomic data from the Census and digital elevation model (DEM) data from the USGS. Without such data, it will be highly difficult to carry out research in many disciplines and to inform decisions and policies based on scientific evidence. The Oak Ridge National Lab also maintains an important global population distribution dataset, LandScan, and GeoAI research played an important role in enhancing its data quality (ORNL 2022). On the other hand, GeoAI research should not be limited to data production but needs to further identify and fill knowledge gaps. This links back to the *use-inspired GeoAI research* discussed earlier, which has a direct connection to theories and knowledge gaps related to the use cases. For research that mainly focuses on methodological development, there typically also exist motivating use cases, such as those in urban studies, disaster resilience, public health, and transportation. We may need to carefully identify relevant theories and research questions in these motivating use cases,

and articulate how our research can contribute to advancing these theories, answering research questions, and informing new methodological development.

Uncertainty and model generalizability are two other areas that have been commented on by the panellists. Uncertainty has been long recognized as an important issue in GIScience (J. Zhang and Goodchild 2002), and researchers generally agree on the existence of uncertainties in both data and models. Unlike many statistical models, AI models often output a single value without providing a range of uncertainty. This result can be partly attributed to the fact that AI models often treat training data as simple data records, rather than samples drawn from distributions as most statistical models do. While some uncertainty quantification methods have been proposed for deep learning in the literature (Abdar et al. 2021), we may need to think about how to integrate and improve such methods in GeoAI research and how to measure the varying uncertainty across different geographic locations. For model generalizability, we typically want the trained models to be able to generalize to other geographic regions via little or no model adaptation (or *few-shot learning* and *zero-shot learning* using machine learning terminology). However, it is very challenging for GeoAI models to perform out-of-distribution learning tasks given the spatial heterogeneity and temporal variations of geographical phenomena (Goodchild and Li 2021). Self-supervised learning and foundation models have been suggested by the panellists as techniques to help improve model generalizability, and it is great to see early work on GeoAI foundation models coming out recently (Mai, Cundy, et al. 2022). Meanwhile, we also need to be wary about the opaque nature of foundation models, which may limit our ability to advance knowledge and pursue open science. How to leverage these and other recent AI techniques for social good and to advance theories in geography? That is an important direction for GeoAI research to explore in the future years.

4. Conclusions

The Annual Meeting of AAG in 2023 marked a five-year milestone since the first AAG GeoAI Symposium in 2018. In this context, we organized a panel and invited five panellists to share their thoughts on the advances and limitations in GeoAI research. This article summarizes the presentations from the panellists and also provides after-panel thoughts from the organizers. To conclude, we synthesize these thoughts to answer the four questions posed before the panel.

- Which areas do you think GeoAI research has excelled in over the past five years?

Three main areas of successes have been identified by the panellists. The first area is the development of spatially explicit GeoAI models, as commented on by Zhu, Yuan, and Aydin. These spatially explicit models take into account the special characteristics of geographic phenomena by incorporating spatial representations and can more effectively accommodate the spatial heterogeneity associated with most geographic phenomena. The second area is the production of large-scale geographic datasets, as commented on by Yuan and Bhaduri from two different perspectives and followed up in our after-panel thoughts. These large-scale geographic datasets (e.g. global building footprint data) might not have existed before and thus have practical value in enabling new research. The third area is the use of GeoAI models to address a wide range of real-world challenges, as commented on by Bhaduri and Yuan and also briefly discussed in the Introduction section of this article. Researchers have already used machine learning and deep learning models in many domains, from disaster response and national security to agricultural management and traffic forecasting. While these real-world problems are often formalized into classification or prediction tasks (as related to Goodchild's comment), GeoAI does play a positive role in helping address them. In addition to the three main areas of successes, Zhu also commented on the successful use of VGI to supplement training data, and Bhaduri commented on the closer interactions between the geography and computer science communities and GeoAI-related developments in industry.

- Which areas do you think GeoAI research has fallen short or could benefit from further improvement?

Three main areas of limitations have also been identified by the panellists. The first area is the lack of engagement with theories in geography, as commented on by Goodchild and Yuan. Existing GeoAI research often focused on classification or prediction tasks, or on deriving more data products from original data. While these research efforts have their value, there is a need to further engage theories in geography, advance knowledge, and improve our understanding of geographic phenomena. The second area is the limited explainability of GeoAI models, as commented on by all five panellists and further expanded in our after-panel thoughts. More research is needed to make the results of GeoAI

models more interpretable to human experts, and to demonstrate the new explanations provided by GeoAI models beyond what existing and more widely accepted models (e.g. regression models) provide. The third area is the limited generalizability of GeoAI models, as commented on by the panellists from different perspectives. Zhu and Yuan commented on the difficulty of generalizing the models trained using data from one geographic region to another due to spatial heterogeneity. It is necessary to differentiate this limited generalizability from the development of spatially explicit models discussed previously. While a spatially explicit model may effectively accommodate spatial heterogeneity within a study region and a time frame (e.g. a city in one year), it may not generalize well to a new study region (e.g. a totally different city with distinct environment and infrastructure) or a different time period due to the temporal variations of geographical phenomena and therefore shows limited generalizability. Bhaduri commented on the difficulty of generalizing the models trained on images from one data source to those from other data sources, due to different image quality and spectral features. In addition to these three main areas of limitations, Bhaduri also commented on the limited research on other types of geospatial data beyond images (e.g. vector data), and Aydin commented on the need for data representation methods for a mixture of quantitative and qualitative data.

- What advancements in GeoAI research would you like to see in the next five years?

In addressing the previous panel question, our panellists have already identified limitations as opportunities for future advancements. Thus, topics discussed previously, such as engaging theories in geography, increasing model explainability, and enhancing model generalizability, are all important directions and opportunities for future GeoAI research. In addition, two

more areas have also been discussed by the panellists. The first area is quantifying uncertainty, as commented on by Aydin and Bhaduri and followed up in our after-panel thoughts. Given the existence of uncertainty in data and models and the multiplication effect when different models and data are used together to address grand challenges, how can we best quantify, visualize, and communicate uncertainties in GeoAI research? The second area is leveraging recent AI methods and techniques, as commented on by Yuan and Bhaduri. For this area, it is exciting to see many GeoAI researchers are already exploring cutting-edge techniques, such as foundation models, and it is probably reasonable to believe that this trend will continue in the near future.

- Any other thoughts you would like to share?

The panellists also shared some stimulating thoughts. Goodchild posed the question: why did we witness the ‘GeoAI phenomenon’ despite much early research on AI and geography in the 1980s and 1990s? Is the recent surge of GeoAI research due to the coin of the term ‘GeoAI’, the availability of massive data, or something more fundamental? Yuan shared three statements for helping think out-of-the-box. These statements serve as reminders that training data in the past may not predict the future, effects may occur before the causes, and data may not be real. Finally, Aydin shared the story of Vaucanson’s Duck and encouraged transparency and open science in future GeoAI research.

These four questions and their answers are further summarized more concisely in [Table 1](#):

To sum up, this article compiles and synthesizes the thoughts from the panellists and organizers from the GeoAI panel at the 2023 AAG Annual Meeting. We hope that these thoughts can help stimulate new ideas in the field, and we look forward to new advancements in GeoAI research in the next five years and beyond.

Table 1. The four questions and their answers summarized as itemized points.

Question	Answers
<i>Which areas do you think GeoAI research has excelled in over the past five years?</i>	<ul style="list-style-type: none"> • Development of spatially explicit GeoAI models • Production of large-scale geographic datasets • Use of GeoAI models to address a wide range of societal challenges
<i>Which areas do you think GeoAI research has fallen short or could benefit from further improvement?</i>	<ul style="list-style-type: none"> • A lack of engagement with theories in geography • Limited explainability of current GeoAI models • Limited generalizability to new geographic areas and data sources
<i>What advancements in GeoAI research would you like to see in the next five years?</i>	<ul style="list-style-type: none"> • Improving over the limitations identified above • Better quantifying, visualizing, and communicating uncertainties • Exploring and leveraging emerging AI methods
<i>Any other thoughts you would like to share?</i>	<ul style="list-style-type: none"> • The ‘GeoAI phenomenon’ and why it happened • The intricate relations among data, predictions, causes, effects, and the real world • Transparency and open science in GeoAI research

Acknowledgements

Yingjie Hu acknowledges support from the U.S. National Science Foundation (NSF) under Grant No. BCS-2117771. May Yuan's effort is based upon work supported by (while serving at) the NSF. Wenwen Li acknowledges support from the NSF under Grant No. BCS-1853864 and OAC-2230034. Song Gao acknowledges support from the NSF under Grant No. OAC-2112606. The statements included in the article are solely the authors' and do not necessarily reflect the view of the NSF.

This manuscript has been co-authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<https://energy.gov/downloads/doe-public-access-plan>).

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by the U.S. National Science Foundation [BCS-2117771, BCS-1853864, OAC-2230034, and OAC-2112606].

Author contributions

Conceptualization: all authors; panel organization: YH, SG, WL, DL, OA, and SN; writing-original draft: YH, MG, AZ, MY, OA, and BB; figures: MY, OA, BB, and YH; writing-editing and revision: all authors; communication and coordination: YH.

Consent

All authors have read and approved the final manuscript.

References

- Abdar, M., F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, et al. 2021. "A Review of Uncertainty Quantification in Deep Learning: Techniques, Applications and Challenges." *Information Fusion* 76:243–297. <https://doi.org/10.1016/j.inffus.2021.05.008>.
- Abler, R. F. 1987. "The National Science Foundation National Center for Geographic Information and Analysis." *International Journal of Geographical Information System* 1 (4): 303–326. <https://doi.org/10.1080/02693798708927819>.
- Bollen, J., H. Mao, and X. Zeng. 2011. "Twitter Mood Predicts the Stock Market." *Journal of Computational Science* 2 (1): 1–8. <https://doi.org/10.1016/j.jocs.2010.12.007>.
- Boulos, K., N. Maged, M.-P. Kwan, K. El Emam, A. Lai-Ling Chung, S. Gao, and D. B. Richardson. 2022. "Reconciling Public Health Common Good and Individual Privacy: New Methods and Issues in Geoprivacy." *International Journal of Health Geographics* 21 (1). <https://doi.org/10.1186/s12942-022-00300-9>.
- Boulos, M. N. K., G. Peng, and T. VoPham. 2019. "An Overview of GeoAI Applications in Health and Healthcare." *International Journal of Health Geographics* 18 (1): 1–9. <https://doi.org/10.1186/s12942-019-0171-2>.
- Brown, T., B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, et al. 2020. "Language Models are Few-Shot Learners." In *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, 1877–1901. Vol. 33. Red Hook, New York: Curran Associates, Inc.
- Cheng, X., J. Wang, H. Li, Y. Zhang, L. Wu, and Y. Liu. 2021. "A Method to Evaluate Task-Specific Importance of Spatio-Temporal Units Based on Explainable Artificial Intelligence." *International Journal of Geographical Information Science* 35 (10): 2002–2025. <https://doi.org/10.1080/13658816.2020.1805116>.
- Chiappinelli, C. 2022. "Think Tank: GeoAI Reveals a Glimpse of the Future." *Esri's WhereNext Magazine*.
- Dobson, J. E. 1983. "Automated Geography." *The Professional Geographer* 35 (2): 135–143. <https://doi.org/10.1111/j.0033-0124.1983.00135.x>.
- Fotheringham, A. S., W. Yang, and W. Kang. 2017. "Multiscale Geographically Weighted Regression (MGWR)." *Annals of the American Association of Geographers* 107 (6): 1247–1265. <https://doi.org/10.1080/24694452.2017.1352480>.
- Gaikwad, S., S. Iyer, D. Lunga, T. Yabe, X. Liang, B. Ananthabhotla, N. Behari, S. Guggilam, and G. Chi. 2022. "Data-Driven Humanitarian Mapping and Policymaking: Toward Planetary-Scale Resilience, Equity, and Sustainability." In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, Washington DC, 4872–4873.
- Gao, S. 2021. "Geospatial Artificial Intelligence (GeoAI)." *Oxford Bibliographies*. <https://www.oxfordbibliographies.com/view/document/obo-9780199874002/obo-9780199874002-0228.xml>.
- Gao, S., Y. Hu, W. Li, and L. Zou. 2023. "Special Issue on Geospatial Artificial Intelligence." *Geoinformatica*. Springer, 1–4.
- Gold, E. M. 1967. "Language Identification in the Limit." *Information & Control* 10 (5): 447–474. [https://doi.org/10.1016/S0019-9958\(67\)91165-5](https://doi.org/10.1016/S0019-9958(67)91165-5).
- Goodchild, M. F. 2007. "Citizens as Sensors: The World of Volunteered Geography." *GeoJournal* 69 (4): 211–221. <https://doi.org/10.1007/s10708-007-9111-y>.
- Goodchild, M. F., and W. Li. 2021. "Replication Across Space and Time Must Be Weak in the Social and Environmental Sciences." *Proceedings of the National Academy of Sciences* 118 (35): e2015759118. <https://doi.org/10.1073/pnas.2015759118>.
- Gupta, J., C. Molnar, Y. Xie, J. Knight, and S. Shekhar. 2021. "Spatial Variability Aware Deep Neural Networks (SVANN): A General Approach." *ACM Transactions on Intelligent Systems and Technology* 12 (6): 1–21. <https://doi.org/10.1145/3466688>.
- Hsu, C.-Y., and W. Li. 2023. "Explainable GeoAI: Can Saliency Maps Help Interpret Artificial Intelligence's Learning Process? An Empirical Study on Natural Feature Detection."

- International Journal of Geographical Information Science* 37 (5): 963–987. <https://doi.org/10.1080/13658816.2023.2191256>.
- Hu, Y., S. Gao, D. Lunga, W. Li, S. Newsam, and B. Bhaduri. 2019. "GeoAI at ACM SIGSPATIAL: Progress, Challenges, and Future Directions." *SIGSpatial Special* 11 (2): 5–15. <https://doi.org/10.1145/3377000.3377002>.
- Hu, Y., W. Li, D. Wright, O. Aydin, D. Wilson, O. Maher, and M. Raad. 2019. "Artificial Intelligence Approaches." *The Geographic Information Science & Technology Body of Knowledge* 2019 (Q3): 1–12. <https://doi.org/10.22224/gistbok/2019.3.4>.
- Janowicz, K., S. Gao, G. McKenzie, Y. Hu, and B. Bhaduri. 2020. "GeoAI: Spatially Explicit Artificial Intelligence Techniques for Geographic Knowledge Discovery and Beyond." *International Journal of Geographical Information Science* 34 (4): 625–636. <https://doi.org/10.1080/13658816.2019.1684500>.
- Janowicz, K., R. Sieber, and J. Crampton. 2022. "GeoAI, Counter-AI, and Human Geography: A Conversation." *Dialogues in Human Geography* 12 (3): 446–458. <https://doi.org/10.1177/20438206221132510>.
- Li, W. 2020. "GeoAI: Where Machine Learning and Big Data Converge in GIScience." *Journal of Spatial Information Science* (20): 71–77. <https://doi.org/10.5311/JOSIS.2020.20.658>.
- Li, Z. 2022. "Extracting Spatial Effects from Machine Learning Model Using Local Interpretation Method: An Example of SHAP and XGBoost." *Computers, Environment and Urban Systems* 96:101845. <https://doi.org/10.1016/j.compenvurbysys.2022.101845>.
- Li, Z., A. S. Fotheringham, T. M. Oshan, and L. J. Wolf. 2020. "Measuring Bandwidth Uncertainty in Multiscale Geographically Weighted Regression Using Akaike Weights." *Annals of the American Association of Geographers* 110 (5): 1500–1520. <https://doi.org/10.1080/24694452.2019.1704680>.
- Li, M., S. Gao, F. Lu, and H. Zhang. 2019. "Reconstruction of Human Movement Trajectories from Large-Scale Low-Frequency Mobile Phone Data." *Computers, Environment and Urban Systems* 77:101346. <https://doi.org/10.1016/j.compenvurbysys.2019.101346>.
- Li, W., and C.-Y. Hsu. 2020. "Automated Terrain Feature Identification from Remote Sensing Imagery: A Deep Learning Approach." *International Journal of Geographical Information Science* 34 (4): 637–660. <https://doi.org/10.1080/13658816.2018.1542697>.
- Li, W., C.-Y. Hsu, and M. Hu. 2021. "Tobler's First Law in GeoAI: A Spatially Explicit Deep Learning Model for Terrain Feature Detection Under Weak Supervision." *Annals of the American Association of Geographers* 111 (7): 1–19. <https://doi.org/10.1080/24694452.2021.1877527>.
- Liu, P., and F. Biljecki. 2022. "A Review of Spatially-Explicit GeoAI Applications in Urban Geography." *International Journal of Applied Earth Observation and Geoinformation* 112:102936. <https://doi.org/10.1016/j.jag.2022.102936>.
- Lunga, D., Y. Hu, S. Newsam, S. Gao, B. Martins, L. Yang, and X. Deng. 2022. "GeoAI at ACM SIGSPATIAL: The New Frontier of Geospatial Artificial Intelligence Research." *SIGSpatial Special* 13 (3): 21–32. <https://doi.org/10.1145/3578484.3578491>.
- Lunga, D., H. L. Yang, A. Reith, J. Weaver, J. Yuan, and B. Bhaduri. 2018. "Domain-Adapted Convolutional Networks for Satellite Image Classification: A Large-Scale Interactive Learning Workflow." *IEEE Journal of Selected Topics in Applied Earth Observations & Remote Sensing* 11 (3): 962–977. <https://doi.org/10.1109/JSTARS.2018.2795753>.
- Ma, Y., Y. Hu, G. R. Moncrieff, J. A. Slingsby, A. M. Wilson, B. Maitner, and R. Z. Zhou. 2022. "Forecasting Vegetation Dynamics in an Open Ecosystem by Integrating Deep Learning and Environmental Variables." *International Journal of Applied Earth Observation and Geoinformation* 114:103060. <https://doi.org/10.1016/j.jag.2022.103060>.
- Mai, G., C. Cundy, K. Choi, Y. Hu, N. Lao, and S. Ermon. 2022. "Towards a Foundation Model for Geospatial Artificial Intelligence (Vision Paper)." In *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, Seattle, 1–4.
- Mai, G., Y. Hu, S. Gao, L. Cai, B. Martins, J. Scholz, J. Gao, and K. Janowicz. 2022. "Symbolic and Subsymbolic GeoAI: Geospatial Knowledge Graphs and Spatially Explicit Machine Learning." *Transactions in GIS* 26 (8): 3118–3124. <https://doi.org/10.1111/tgis.13012>.
- Mai, G., K. Janowicz, Y. Hu, S. Gao, B. Yan, R. Zhu, L. Cai, and N. Lao. 2022. "A Review of Location Encoding for GeoAI: Methods and Applications." *International Journal of Geographical Information Science* 36 (4): 639–673. <https://doi.org/10.1080/13658816.2021.2004602>.
- Mao, H., Y. Hu, B. Kar, S. Gao, and G. McKenzie. 2018. "GeoAI 2017 Workshop Report: The 1st ACM SIGSPATIAL International Workshop on GeoAI: AI and Deep Learning for Geographic Knowledge Discovery." *SIGSpatial Special* 9 (3): 25–25. <https://doi.org/10.1145/3178392.3178408>.
- Nguyen, T.-A., B. Kellenberger, and D. Tuia. 2022. "Mapping Forest in the Swiss Alps Treeline Ecotone with Explainable Deep Learning." *Remote Sensing of Environment* 281:113217. <https://doi.org/10.1016/j.rse.2022.113217>.
- OpenAI. 2023. "GPT-4 Technical Report." <https://arxiv.org/abs/2303.08774>.
- Openshaw, S., and C. Openshaw. 1997. *Artificial Intelligence in Geography*. Hoboken, New Jersey: John Wiley & Sons, Inc.
- ORNL. 2022. *Public Release of ORNL Global Population Distribution Data Aids Humanitarian Support*. <https://www.ornl.gov/news/public-release-ornl-global-population-distribution-data-aids-humanitarian-support>.
- Pearl, J., and D. Mackenzie. 2018. *The Book of Why: The New Science of Cause and Effect*. New York: Basic books.
- Peng, B., Q. Huang, J. Vongkusolkrit, S. Gao, D. B. Wright, Z. N. Fang, and Y. Qiang. 2020. "Urban Flood Mapping with Bitemporal Multispectral Imagery via a Self-Supervised Learning Framework." *IEEE Journal of Selected Topics in Applied Earth Observations & Remote Sensing* 14:2001–2016. <https://doi.org/10.1109/JSTARS.2020.3047677>.
- Radford, A., K. Narasimhan, T. Salimans, and I. Sutskever. 2018. "Improving Language Understanding by Generative Pre-Training." https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf.
- Radford, A., J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. 2019. "Language Models are Unsupervised Multitask Learners." *OpenAI Blog* 1 (8): 9.
- Rao, J., S. Gao, Y. Kang, and Q. Huang. 2021. "LSTM-TrajGAN: A Deep Learning Approach to Trajectory Privacy Protection." In *Proceedings of the 11th International Conference on Geographic Information Science* [Online], 12:1–17.

- Vapnik, V. N., and A. Y. Chervonenkis. 2015. "On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities." In *Measures of Complexity: Festschrift for Alexey Chervonenkis*, edited by V. Vovk, H. Papadopoulos, A. Gammerman, 11–30. New York: Springer.
- Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. 2017. "Attention Is All You Need." In *Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS 2017)*, December 4–9, 2017, Long Beach, CA, 30: 1–11. New York: ACM.
- Wang, J., and F. Biljecki. 2022. "Unsupervised Machine Learning in Urban Studies: A Systematic Review of Applications." *Cities* 129:103925. <https://doi.org/10.1016/j.cities.2022.103925>.
- Xing, J., and R. Sieber. 2023. "The Challenges of Integrating Explainable Artificial Intelligence into GeoAI." *Transactions in GIS* 27 (3): 626–645. <https://doi.org/10.1111/tgis.13045>.
- Yan, B., K. Janowicz, G. Mai, and S. Gao. 2017. "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*, November 7–10, 2017, Redondo Beach, CA, 35:1–10. New York: ACM. <https://doi.org/10.1145/3139958.3140054>.
- Yuan, Q., H. Shen, T. Li, Z. Li, S. Li, Y. Jiang, H. Xu, et al. 2020. "Deep Learning in Environmental Remote Sensing: Achievements and Challenges." *Remote Sensing of Environment* 241:111716. <https://doi.org/10.1016/j.rse.2020.111716>.
- Zhang, Y., T. Cheng, Y. Ren, and K. Xie. 2020. "A Novel Residual Graph Convolution Deep Learning Model for Short-Term Network-Based Traffic Forecasting." *International Journal of Geographical Information Science* 34 (5): 969–995. <https://doi.org/10.1080/13658816.2019.1697879>.
- Zhang, J., and M. F. Goodchild. 2002. *Uncertainty in Geographical Information*. Boca Raton, FL: CRC press.
- Zhang, R., S. Newsam, Z. Shao, X. Huang, J. Wang, and D. Li. 2021. "Multi-Scale Adversarial Network for Vehicle Detection in UAV Imagery." *Isprs Journal of Photogrammetry & Remote Sensing* 180:283–295. <https://doi.org/10.1016/j.isprsjprs.2021.08.002>.
- Zhang, G., and A.-X. Zhu. 2020. "Sample Size and Spatial Configuration of Volunteered Geographic Information Affect Effectiveness of Spatial Bias Mitigation." *Transactions in GIS* 24 (5): 1315–1340. <https://doi.org/10.1111/tgis.12679>.
- Zheng, Z., and R. Sieber. 2022. "Putting Humans Back in the Loop of Machine Learning in Canadian Smart Cities." *Transactions in GIS* 26 (1): 8–24. <https://doi.org/10.1111/tgis.12869>.
- Zhou, R. Z., Y. Hu, J. N. Tirabassi, Y. Ma, and Z. Xu. 2022. "Deriving Neighborhood-Level Diet and Physical Activity Measurements from Anonymized Mobile Phone Location Data for Enhancing Obesity Estimation." *International Journal of Health Geographics* 21 (1): 1–18. <https://doi.org/10.1186/s12942-022-00321-4>.
- Zhou, B., L. Zou, A. Mostafavi, B. Lin, M. Yang, N. Gharaibeh, H. Cai, J. Abedin, and D. Mandal. 2022. "VictimFinder: Harvesting Rescue Requests in Disaster Response from Social Media with BERT." *Computers, Environment and Urban Systems* 95:101824. <https://doi.org/10.1016/j.compenvurbysys.2022.101824>.
- Zhu, A.-X. 2017. "Artificial Neural Networks." In *International Encyclopedia of Geography: People, the Earth, Environment and Technology*. Wiley. <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118786352.wbieg0871>.
- Zhu, D., X. Cheng, F. Zhang, X. Yao, Y. Gao, and Y. Liu. 2020. "Spatial Interpolation Using Conditional Generative Adversarial Neural Networks." *International Journal of Geographical Information Science* 34 (4): 735–758. <https://doi.org/10.1080/13658816.2019.1599122>.
- Zhu, X. X., D. Tuia, L. Mou, G.-S. Xia, L. Zhang, F. Xu, and F. Fraundorfer. 2017. "Deep Learning in Remote Sensing: A Comprehensive Review and List of Resources." *IEEE Geoscience and Remote Sensing Magazine* 5 (4): 8–36. <https://doi.org/10.1109/MGRS.2017.2762307>.
- Zhu, A.-X., and M. Turner. 2022. "How is the Third Law of Geography Different?" *Annals of GIS* 28 (1): 57–67. <https://doi.org/10.1080/19475683.2022.2026467>.