KI - Künstliche Intelligenz (2023) 37:91–97
https://doi.org/10.1007/s13218-022-00790-6

INTERVIEW

GeoAI and Beyond

Interview with Krzysztof Janowicz

Simon Scheider[1] · Kai-Florian Richter[2] · Krzysztof Janowicz[3,4]

Published online: 20 January 2023

Krzysztof Janowicz is a professor for Geoinformatics at the University of Vienna and the University of California, Santa Barbara. His research focuses on how humans conceptualize the space around them based on their behavior, focusing particularly on regional and cultural differences, with the goal of assisting machines to better understand the information needs of an increasingly diverse user base.

✉ Simon Scheider
s.scheider@uu.nl

Kai-Florian Richter
kai-florian.richter@umu.se

Krzysztof Janowicz
krzysztof.janowicz@univie.ac.at

1 Department of Human Geography and Spatial Planning, University Utrecht, Princetonlaan 8a, 3584 Utrecht, CB, Netherlands

2 Department of Computing Science, Umeå University, 901 87 Umeå, Sweden

3 Department of Geography and Regional Research, Vienna University, Universitätsstraße 7, Vienna 1010, Austria

4 Center for Spatial Studies, University of California, Santa Barbara, 4830/2610, Ellison Hall, Santa Barbara, US

1 Personal questions

KI: Krzysztof, in 2000 you started a company centred around Linux networking and IT training in Germany. In 2003 you received a diploma in geo-ecology at the University of Münster, Germany and a PhD in geoinformatics in 2008 from the same institution. Afterwards, you made an impressive career as a geographic information scientist in the AI-related field of geo-semantics, first at the Pennsylvania State University (PSU) and then at the University of California Santa Barbara (UCSB). In 2014 you received the Harold J. Plous Award, one of UCSB's most prestigious faculty honours, and in 2016 you received your department's outstanding mentor award. Since 2020 you are the director of UCSB's Center for Spatial Studies. Currently, you are a professor at the University of Vienna. How has this biography, especially its origin, shaped your views on the field of AI?

If you read through my body of work, I guess you would see a lot of ecology and even evolutionary biology shine through. Cognitive science has also influenced a lot of my thinking, for instance my work on similarity. Earlier, I was interested in philosophy, this is reflected in my research on information science ontologies. When it comes to my background in industry, I guess it has also shaped the way I look at research in general. I remain very involved with the industry and technology transfer. I am just returning to Europe from the US after 12 years. While in the US, I have always provided a European perspective in terms of critical thinking, the style of research, topics such as privacy, and mentorship. Now that I'm in Europe, I will try to do the opposite. I will take an American perspective, for instance, looking for opportunities first and problems second. It's the combination of these perspectives that matters.

KI: Your work has been located in a field between geographic information science (GIScience) and AI. In my view, there have been at least three important technological developments in AI which particularly influenced GIScience: (1)

The adoption of the Semantic Web and Linked Data principles, which recently led to W3C's "Spatial data on the web best practices" [1]*for example, and the development of geographic knowledge graphs. (2) Second, the increased availability of geo-referenced texts from geo-social media and natural language processing (NLP), which improves geographic information retrieval (GIR). And (3) there's this large field of deep learning based image recognition used in the context of remote sensing. How do you think your career has echoed these developments, or even driven some of them?*

I would not say *driven*, but I think it's fair to say that I've been *involved* in all three. Your overview makes a lot of sense, but let me provide an additional perspective: You could look at it from a temporal viewpoint and at the geography in which this was playing out. In the early 80s we saw the first really substantial push of AI technology in geography. These were the "pre-GIScience days". It's fair to say that this was, in part, driven by UCSB and the US. Helen Couclelis and Terry Smith come to mind, and on the European side, certainly Stan Openshaw. Then starting in 1998 and peaking around 2010, there was the geo-ontologies and geo-semantics push. I believe this mainly took place in Europe, including Italy, Germany, Austria, and the UK. I am thinking of people like Andrew Frank and Werner Kuhn here, and the people around them that came from similar schools of thought, even though they might have done part of the work in the US, like Thomas Bittner and Max Egenhofer, for example. And of course David Mark from Buffalo. Around 2017, GeoAI took off as part of the broader Spatial Data Science development. It is not unreasonable to say that UCSB had a vital role to play in this as well. Many of our former students who are now professors are among the leading minds behind the ongoing GeoAI push: Song Gao, Yingjie Hu, Grant McKenzie, WenWen Li, who was a postdoc at UCSB, come to mind and more recently Gengchen Mai and Rui Zhu.

Under the hood, there's also another interesting aspect. Geographically speaking, many of the currently prominent GeoAI researchers are from Asia, specifically China. One may argue that they mostly work at US universities or that the key breakthroughs in these fields still come from a very small set of US universities and industry players. I experience it differently; there is no creativity or innovation gap. I would rather say that one of the most shocking insights for me, in hindsight, was this: when I came to the US, the competition in the fields of geo-semantics, big data, data science, GIScience, and so on was mainly between the United States and Europe. Many Europeans like me came to the US to work on these topics. Today, this has vastly shifted to a competition between the United States and China. Sometimes it feels like Europe has been reduced to a mere market, as far as adopting results from AI or data science is concerned. This is both interesting and problematic because current developments in the broader field of AI, whether applied to geo-technologies or not, are redefining the fabric of society.

KI: Is this something you would like to change, as you recently moved to Vienna?

Well, I would like to make a tiny contribution. In general, I think that science should not be a zero-sum competitive game. It's best when done in teams and when the different cultures and schools of thought bring in what they are best at in order to create a global and openly available body of knowledge.

KI: You're one of the founders and editors in chief of the journal "Semantic Web - Interoperability, usability, applicability". This is a high-quality open access journal founded in 2010, which quickly advanced to one of the leading journals in the field of the Semantic Web. As a GIScientist working in the field of the Semantic Web, in which communities do you feel at home? How do you think these communities link?

I feel at home where people do interesting work, especially at the intersection of geography or spatial sciences, computer science, and cognitive science. The most exciting phase is always the initial spark when new ideas arise at the intersections of domains, fueling rapid progress, new collaborations, or shining a light on the same old problems from very different perspectives. In general, I like to think in analogies. These days, my attention span gets shorter and shorter. Long phases in research where we try to tweak our work for a tiny delta often tire me. This is both my strength and my weakness, depending on how you like to look at it.

You also mentioned the journal. I've been involved in open source and open content for more than 20 years now. But before 2010, I never had the chance to also bring some of these ideas into the process of academic publishing. The Semantic Web Journal has an open and transparent review and management policy, and at that time this was very radical. So everything that happens in the journal, uploading a paper, the assignment of editors and reviewers, the reviewing, an editor's decision, everything plays out in the open. Papers are visible, reviews are visible, and the decisions in between are visible. In the end, the reviewers and editors are attributed in the header of the paper. I guess Pascal Hitzler (the other co-founder) and I feel it's better to foster an atmosphere of collaboration and joint ownership in science instead of blind competition.

[1] https://www.w3.org/TR/sdw-bp/.

2 About AI

KI: Let's talk a bit about AI in general. One question that I have is related to what you said about cognition and evolutionary biology. Much of the research on spatial information has been on spatial cognition. In AI, there's this term of cognitive plausibility, which includes the ability to communicate in the ways humans do, to take into account human perspectives, including context, intention, and beliefs, and to provide solutions that are not only accurate but also acceptable according to human ethical and moral standards. Do you think that AI, including GeoAI, should be cognitively plausible to be able to operate within human society?

In my dissertation about semantic similarity, I worked on some of these questions: What is cognitively plausible? What is cognitively adequate? How do they differ? One distinction is whether AI should mimic how humans reason, or whether we can relax the condition so that only results align with human thinking. That's a substantial difference because in the latter case, you only have to produce outputs that are less surprising to a human or that share cognitive biases. In the other case, you need to make sure the process by which we derive results is taken into account, too. But humans are irrational agents, as we know from Kahnemann and Tversky [9], and others. So you are essentially asking the question: How irrational do we want AI to be? Would an AI or recommender system as an interface between AI and humans benefit from human-like approaches to intelligence? In the 21st century, we are faced with the very high cost of humans being irrational agents. There's this irrational fear of AI, while millions are dying due to human stupidity. Whether it's our struggle with COVID, our role in the changing climate, the war in Ukraine, or even social media, humans are not really humane. Artificial general intelligence, if possible and measurable at all, will be nothing like human intelligence, and I guess this is good news.

KI: But if we are going to develop an AI that is capable of doing this, what do you think is needed to deal with intentions, beliefs, and context? So that these systems can satisfactorily communicate with humans? Do you think a mix of connectionist AI and symbolic approaches could substitute human experts?

I think the most interesting work happens at the intersection between top-down, symbolic or declarative knowledge representation and reasoning, and neural or connectionist architectures [6]. Ling Cai, for instance, looks at representing and reasoning over topological relations purely from a representation learning perspective [1]. One could say this is learning without prior knowledge. But, of course, that's not entirely true, since the architecture for learning embeddings is heavily driven by RCC8 theory [11]. For instance, we need to ensure that reflexivity or transitivity can be learned in principle. So in the end, it needs both. Learning excels at handling noise and incomplete knowledge at the cost of requiring a lot of training data. Outcomes may be difficult to predict and to transfer, and there is always a temporal lag as data is about the past. Classical deduction complements this very well.

In a recent discussion I had with Renee Sieber, who is thinking critically about some of the present GeoAI and ML work [13], she made an interesting argument based on a paper she read that studies how neural architectures classify dogs, cats, and so on. She said that the authors found that the *Cat* category was not only identified, as they expected, by their body features, but also by the background, since more cats were photographed inside and dogs outside. Renee used this as a basis for her critique. For me, I would argue that this is actually a strength of such connectionist approaches. It may be adequate to learn that cats appear more in an indoor context and dogs more in an outdoor context. More broadly speaking, meaning and the ways we assign category membership is an emerging, context-dependent property. This is one of the surprising insights from both the cognitive science and machine learning literature. It is worth noting that AI may use entirely different sets of features for classification and other tasks than we presently do.

KI: Human learning is indeed very context-dependent. It is also connectionist in the sense that a lot of it is dependent on a network of concepts. There are studies at least from the 80s in Psychology showing this. However, does this really imply that machine learning is the way to account for this? Maybe you know the work of Gerd Gigerenzer [3, 5]. He has investigated the failures and biases of complex machine learning models compared to simple human heuristics. Gigerenzer illustrates how complex models cause large amounts of problems compared to human judgments because they are unstable. ML models are usually highly sensitive to the training data quality. This severely limits their robustness, as illustrated by, for example, the failures to recognize traffic signs, which can be easily hacked, or by models that are trained to detect traffic objects but confuse pedestrians with road infrastructure because they cannot deal with unforeseen situations. In the case of Uber in 2016, this problem caused a fatal accident with a self-driving car [10]. What do you think could be done to overcome this blind spot?

The problems you're describing are related to the "bias-variance tradeoff". They should also be quite visible in human learning. I think the way humans work, reflected in the literature you mentioned, is that we are very good at heuristics. What saves us is a high bias and a relatively low variance. But keep in mind that we are also very forgiving of ourselves. We can hold incoherent beliefs, we filter and forget what we don't like, and everyday behavior does not require any optimal solutions. Instead, we find good enough solutions most of the time. We also often ignore that human learning is not a one-person task. It is influenced by society. Our decisions and beliefs should not be viewed in isolation. They are shaped by people and places around us. Dan Montello, for instance, is interested in *collaborative navigation* [2]. The way we learn to achieve something, e.g., to recognize a traffic sign, is heavily based on collaboration.

Regarding the bias-variance trade-off, I could now toss along terms like regularization or ensemble methods. But I think the problem runs deeper, and my intuition is that it is mostly one of expectations. For instance, coming back to your specific example: Would you rather sit with me as a driver or with the Tesla Autopilot? My answer is: In Las Vegas I would go with the Tesla every day. Here in Vienna, I may still be the better driver.

KI: Well, if we didn't expect any other traffic participants, autopilots would be just fine. But it's the pesky cats and dogs and bicycles which throw these systems off. By the way, Gerd Gigerenzer argues that under these circumstances, what he calls the "unstable world", heuristics can turn out to be far better than any more sensitive models [4]. It is therefore actually not the case that human learning over small sample sizes, which is based on such heuristics, is worse or falls for the same kind of problems. On the contrary, there is evidence that it can work better, even if non-optimal.

As a meta-remark: Part of the problem may also be due to how we do research. There is a lot of small-delta, leaderboard-style research based on the same, often synthetic datasets, without enough attention to transferability, complexity and energy efficiency. But keep in mind that I only represent a tiny fragment of the AI field.

3 About GeoAI

KI: Let's move on to the specific area of GeoAI. I would be interested to learn how you see GeoAI as a field. How would you define it?

You could say that GeoAI is AI applied to geographic problems or geo-data. But this is not my view. It leads to the same old narrative in which the geosciences are either the receivers of technology, i.e., the application domain, or merely the data providers. It is more interesting to look at our contributions to the broader foundational AI literature, and I believe GeoAI makes such contributions. One example is research on "spatially explicit models" from Gengchen Mai and Bo Yan [12]. I think that our community has something interesting to offer to the broader field. This is not just because of the usual argument that spatial is special and that many interesting questions have a spatial component. Due to the highly disciplinary nature of our field, researchers in our domain were often among the early adopters or even early contributors. Examples include work on information ontologies, classification and categorization in remote sensing, looking into issues of representativeness, coverage, bias, and FAIRness, and so on before many others did. So, what is GeoAI? GeoAI applies insights from AI to better address fundamental questions in geography and the earth sciences, but it is also a field that contributes new insight to AI as far as the spatial and temporal properties of knowledge, behavior, and intelligence are concerned.

KI: Do you think there's a community evolving around GeoAI?

Yes, I do, and many scientists from the broader AI, ML and KDD fields are aware of it. However, it's unfortunate that fields such as geoinformatics and GIScience sometimes insist on using their own terminology. This hinders communication and the reuse of our contributions. For example, we should have cast a lot of our work, such as spatial autocorrelation, in terms of compression, information content, and entropy. This does not take away from the fact that we require terms such as neighborhoods, stationarity, interaction, and so on. In fact, this way, we would add to the broader information theoretic perspective instead of developing in parallel to it.

KI: GeoAI has the syllable "geo" in it. You're a professor of "geographic" information science and you come from a geography department. Do you think there's a special role for "geo" other than just being "spatial AI"? Should GeoAI be part of spatial AI?

I would be careful interpreting "geo" as "geographic". Of course, I'm currently in a geography department, and a lot of my colleagues and my thinking is influenced by geography. I learned a lot about knowledge representation and semantics from geography. But "geo" should be broadened to the geosciences. For good reasons, geoscientists do not feel fully represented by the word geographic, since the latter has a substantially stronger descriptive and human component. Then again, the term "geo" could be broadened to "spatial". I'm currently the director of the Center for Spatial Studies at UCSB. Underlying this center is the idea that there is something common across, e.g., astronomy, the life sciences, chemistry, biology, geography, earth science, architecture, and so on, in terms of spatial and temporal characteristics. Yet, calling the field just "SpatialAI" wouldn't be good. For instance, prior work has shown that human cognition of places can overwrite the otherwise typically metric properties of space. Take the example of Tobler's First Law, which states that observations close in geographic space are likely to be also close in some attribute space, say elevation. This law isn't isotropic. It is heavily influenced by the ways humans shape space and assign meaning to it. Suppose you travel the same distance north or south from San Diego. In that case, your experience will differ radically, for instance, due to abrupt changes in cultural and social norms and legal frameworks. If GeoAI should serve humans (and here we are back to the discussion about cognitive plausibility), then it needs to be about place. And the notion of place is really at the core of geography. I think I went full circle now ...

KI: So we are back with geography, then?

Well, maybe you made me realize some things I haven't thought through yet *[laughs]*. But I believe it should be "GeoAI" without necessarily implying geographic.

KI: Currently, do you see a particular kind of approach in the field of GeoAI which inspires you most? Work which you would consider leading the way also in the future in this field?

Many of the recent breakthroughs come from within machine learning. But keep in mind AI is not just that. Still, I find these breakthroughs inspiring because they give me a new toolkit to think about old problems. An old problem that I am passionate about is knowledge representation and especially its local characteristics. How regionally variant are human conceptualizations of space? How many cultural, social, spatial, and "platial" differences exist in how humans perceive the environment and communicate about it? Ten years ago, we started work on *semantic signatures* to study this. Can we develop shareable, pre-trained, reusable, openly published libraries of signatures, similar to spectral signatures in remote sensing, but based on human behavior, human descriptions of places, their spatial layout, and the times people interact with these places? Can we characterize places such as neighborhoods or individual establishments purely using such data-driven approaches and then reason about similarity or provide summaries? Years ago, these ideas fueled several PhD topics. But we always missed a unified way to represent these semantic signatures. For instance, they were all glued together by essentially concatenating spaces that had nothing in common. Then suddenly, around 2015, my students introduced me to "embeddings" and *representation learning* as a sub-field of machine learning, and I was very impressed. Not only because the first approaches in the *TransE* family were so geometric, but because they were strikingly beautiful to me, despite all of their shortcomings. And they opened up entirely new avenues for thinking about semantic signatures, similarity, and even analogy more broadly.

KI: In the past, you have written several papers in which you highlighted the autonomous role of semantics (e.g. "Why the data train needs semantic rails" [8]), including knowledge domains like geoscience and geography, for the advancement of AI. I was wondering, with this recent work, which is basically machine learning, how are you looking at your old arguments? How do you currently see the role of semantics and geospatial knowledge for the advancement of AI, when machine learning has taken over large parts of your agenda?

I am only a visitor to this field; I do not see myself as a machine learning researcher. For me, this is yet another interesting way to study the spatial, temporal, and cultural dimensions of geographic feature types. So, yes, semantics remains close to my work. Given the success of foundational models such as GPT-3 or DALL-E, I do ask myself what geo-foundational models would look like and how others could use them as building blocks for an entire zoo of models and downstream tasks. In the end, at its core, this is also a knowledge representation question, isn't it?

KI: Would semantics then be about "overseeing" this zoo, or bridging these models, for example?

Yes, that is a nice view. These days, the majority of my semantics work so to speak is about knowledge graphs, like the KnowWhereGraph project [7]. These knowledge graphs are interesting in many regards, for instance because they can act as a common data structure underlying and integrating highly heterogeneous data. For instance, it would be interesting to see whether we can draw on them as fact repositories for language models.

KI: In 2021 you wrote an article about GeoAI called "GeoAI: spatially explicit artificial intelligence techniques". In this article, you discuss the role of symbolic AI and the so-called "spatially explicit models" for geographic information. These are human-understandable models of space as opposed to neural models that learn spatial representations from labelled data. In this paper, you ask what portion of a data set has to be spatial to justify spatially explicit models. And you conclude, "those that favour domain-specific models will have to justify why developing more complex models is superior to providing more labelled data". Can you explain these statements?

There is a brilliant keynote by Frank van Harmelen from 2011[2] in which he asked: Have we learned something (like a fundamental law) from designing all the Linked Data, from all the knowledge graphs, about the foundational properties of *knowledge*? For instance, for me, one of the most striking insights about AI, in general, is that intelligence does not require consciousness. What Frank asked in my recollection was something like, "is knowledge a graph?". This is the type of issue we discussed in the editorial you mentioned.

Second, how spatial does the data set have to be before spatially explicit models make a difference given some benchmark dataset? There is the recurring argument that more general models beat more specific models, given better or more training data. Yet this argument overlooks a couple of things. For this to happen we need representative, unbiased, high-quality training data. Yet first of all, most of us do not have access to that. In parts progress in AI is fueled

[2] http://iswc2011.semanticweb.org/keynotes/keynote-speakers/frank-van-harmelen/.

by people who have access to very specific datasets. That's why progress is not uniform: we're making progress in some parts, essentially where big companies collect lots of data, but not so much progress in other parts. Furthermore, do you truly know whether a data set is representative? It may be representative of the past, but how representative of the past is the present, with all the currently changing social norms? If you have very large quantities of high-quality, representative, and current data, congratulations to you! But for the rest of us, we have to be content with something else.

"'If you have very large quantities of high-quality, representative, and current data, congratulations to you! But for the rest of us, we have to be content with something else.'"

And finally: How general are these models really? In the end, the intelligence still sits in front of the screen, creating the architectures behind these models and baking or hardwiring many assumptions into them. Going back to *TransE*-like embedding models: They often cannot deal well with reflexivity, cannot handle 1-to-n or m-to-n relations, nor transitivity. No matter how much data you would throw at them, would they overcome these restrictions and suddenly become suitable to the RCC8 examples we discussed? I don't think so.

"'How general are these models really? In the end, the intelligence still sits in front of the screen.'"

KI: I wonder to what extent this insight conflicts with your argument in the paper. Namely the idea that we could get rid of these sources of knowledge, the more complex, specific models and the background knowledge? If you throw this all away, you will end up exactly with the problem that you need very good data to re-learn what you get for free, so to say, when you use an explicit model.

I guess this brings us back to the idea of geo-foundational models again. I do believe such models are possible. This does not mean that there is no room for theory in designing such models nor that such general, foundational models can be applied irrespective of the specific downstream task or geographic region.

KI: One of the biggest challenges I see regarding all these data-driven GeoAI approaches lies in their dependency on specific high-quality data collections. Starting from human experimental subjects, collecting large amounts of high quality answers to highly specific questions in Geographic Question Answering (GeoQA), over getting enough experts to label maps, obtaining enough user dialogue/interactions, and collecting specific geographic text corpora. These collections usually do not exist and require a lot of effort to generate. Thus if you start learning from scratch you run into a fundamental bottleneck. Although there are of course tons of geo-social media data and data from crowdsourcing available, those sources often lack the required quality for doing something specific. Microsoft and others have large question corpora, but specific questions for geography are hard to find. How to overcome this bottleneck?

One of the bottlenecks is the way how we structure science these days. If you're doing something truly novel, if you are designing a new type of infrastructure, if you are answering different or more specialized questions, as you said, then you are going to run against the fact that you will have fewer data to evaluate against. In consequence, your paper will be more likely to be rejected or not very visible.

As we discussed, breakthroughs are being made in highly specific areas. Think about the stock photography industry. Parts of it will likely disappear in a few years. What about art in 2030? One of the reasons for the success of DALL-E is the availability of millions and millions of pictures that we all publish. But we don't have this luxury for many other topics that we would like to better understand. What about society at large? Isn't this one of science's big frontiers (in addition to the brain and the universe)? We were taken off guard by our response to the pandemic, recent political elections, and so on. It is not only the availability of corpora you mentioned; it is a question about what the past tells us about the future.

There is this criticism that modern-day ML models are all just parrots. So let's assume we would be able to design high-quality parrots and let's say GPT-3 or DALL-E are such parrots (by the way, I think this parrot analogy is a mischaracterization). Wouldn't those parrots be super useful for running experiments about the structure of society? Maybe these foundational models will enable radically new insights and research methods for the social sciences? The argument about parrots is often made in the context of biases and how to mitigate them. Systems turned out to be racist, discriminating against women or other parts of the population. But isn't this a criticism, not of these systems, but rather of society? If you're living in a society with systemic bias and racism, then this is going to be mirrored by the parrots.

KI: Let me come back to GeoAI again. Isn't it nevertheless important to know precisely which kind of know-how and cartographic knowledge is needed for understanding geographic information? I guess you probably would agree that the knowledge required for dealing with spatial information is nontrivial. It goes beyond knowing the geometries or having the geodata, for example, because there's also some map interpretation involved, and this interpretation is not in the data. If this is true, isn't then the reduction of GeoAI to "knowledge extraction from geo-referenced data" generating a severe problem? Doesn't this run the risk of underestimating the difficulty of the task, because it underestimates what it means to interpret geodata?

That's a good question. I agree with you. This makes it difficult to say something interesting or provocative in reply. One possible counterargument I can offer is this: many domains can claim this. People in social sciences could say that our view of society, about elections, policy, and trying to capture all of this computationally, are equally reductionist. The second counterargument could be that the more types of sources we have for the same occurrence of a real-world entity, imagery, video, audio, written text, and geometries, the more systems could learn across representations. So the future could make this argument irrelevant.

KI: What would you advise a young researcher who is just entering the field of GeoAI? What should he or she learn?

Last year I was told by an editor of a top-tier geoinformatics journal that 80% of submissions were somehow related to GeoAI. Whether it is 80%, 50%, or 30%, this is not healthy for a scientific community. In the best case, we are overlooking other important areas of study and their contributions. In the worse case, we are draining these other areas, e.g., by cutting their funding supplies. For instance, think about funding cycles in medical research before and after the pandemic.

KI: Thanks a lot for the interview, Krzysztof!

References

1. Cai L, Janowicz K, Zhu R, Mai G, Yan B, Wang Z (2022) Hyperquaternione: a hyperbolic embedding model for qualitative spatial and temporal reasoning. GeoInformatica 2:1–39
2. Dalton RC, Hölscher C, Montello DR (2019) Wayfinding as a social activity. Front Psychol 10:142
3. Gigerenzer G (2022) How to stay smart in a smart world: Why human intelligence still beats algorithms. Penguin, UK
4. Gigerenzer G, Brighton H (2009) Homo heuristicus: Why biased minds make better inferences. Top Cogn Sci 1(1):107–143
5. Gigerenzer G, Hertwig R, Pachur T (2011) Heuristics: the foundations of adaptive behavior. Oxford University Press, Oxford
6. Hitzler P (2022) Neuro-symbolic artificial intelligence: The state of the art. IOS Press
7. Janowicz K, Hitzler P, Li W, Rehberger D, Schildhauer M, Zhu R, Shimizu C, Fisher CK, Cai L, Mai G et al (2022) Know, know where, knowwheregraph: A densely connected, cross-domain knowledge graph and geo-enrichment service stack for applications in environmental intelligence. AI Mag 43(1):30–39
8. Janowicz K, Van Harmelen F, Hendler JA, Hitzler P (2015) Why the data train needs semantic rails. AI Mag 36(1):5–14
9. Kahneman D, Slovic SP, Slovic P, Tversky A (1982) Judgment under uncertainty: Heuristics and biases. Cambridge University Press, Cambridge
10. Kohli P, Chadha A (2019) Enabling pedestrian safety using computer vision techniques: a case study of the 2018 uber inc. self-driving car crash. In: Future of Information and Communication Conference, pp. 261–279. Springer
11. Randell DA, Cui Z, Cohn AG (1992) A spatial logic based on regions and connection. In: Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning, pp. 165–176
12. Yan B, Janowicz K, Mai G, Zhu R (2019) A spatially explicit reinforcement learning model for geographic knowledge graph summarization. Trans GIS 23(3):620–640
13. Zheng Z, Sieber R (2022) Putting humans back in the loop of machine learning in Canadian smart cities. Trans GIS 26(1):8–24