



Pragmatic GeoAI: Geographic Information as Externalized Practice

Simon Scheider¹ · Kai-Florian Richter²

Received: 11 November 2022 / Accepted: 21 December 2022 / Published online: 20 January 2023
© The Author(s) 2022

Abstract

Current artificial intelligence (AI) approaches to handle geographic information (GI) reveal a fatal blindness for the information practices of exactly those sciences whose methodological agendas are taken over with earth-shattering speed. At the same time, there is an apparent inability to remove the human from the loop, despite repeated efforts. Even though there is no question that deep learning has a large potential, for example, for automating classification methods in remote sensing or geocoding of text, current approaches to GeoAI frequently fail to deal with the pragmatic basis of spatial information, including the various practices of data generation, conceptualization and use according to some purpose. We argue that this failure is a direct consequence of a predominance of *structuralist* ideas about information. Structuralism is inherently blind for purposes of any spatial representation, and therefore fails to account for the intelligence required to deal with geographic information. A *pragmatic turn* in GeoAI is required to overcome this problem.

Keywords AI for geographic information · Purpose · Practice of geographic information · Explainable AI

1 Introduction

Based on technological developments in Machine Learning (ML), Artificial Intelligence (AI) has made impressive progress in imitating, substituting, and challenging human competencies in various domains of human culture that have long been difficult to handle by computers [59]. An often-cited example is deep learning methods in self-driving vehicles that effectively recognize moving objects and street signs. The recent AI wave also affects central methods of geographic information and the geosciences, including distant reading, harvesting, and georeferencing of spatial information in natural language texts [33, 77], or geographic question-answering (geoQA) and geo-information retrieval (GIR) [9, 57, 72], as well as recognition and retrieval of objects in geo-referenced images, in particular, remote

sensing imagery and point clouds [90]. Today, AI methods are being used for modelling geospatial phenomena across all spheres of the earth [78]. The trend of applying AI methods to such problems is sometimes called *GeoAI* [43].

Since AI affects or even substitutes the more traditional approaches to all these tasks mentioned above, the underlying sciences (e.g., geosciences, life or social sciences making use of geographic information) are modified, too. Geoscientists are beginning to put their trust into generalized, automated learning to accomplish tasks they previously tackled with tailored methods because machine learning methods often accomplish higher accuracy on test data [78]. Furthermore, there is a corresponding trend to regard the practices within these sciences less as a precious methodological heritage requiring careful reconsideration, but rather as a burden that needs to be overcome, to make place for the next, “geographic” variant of a universal, AI fuelled data science [74]. In the latter, geographic information practices are either becoming obsolete or are being transformed into mere applications of (what is essentially) ML to a specific kind of georeferenced data. Such substitution may appear beneficial for any information practices. We know that human cognition is biased and error-prone [46], so isn’t the prospect of such substitutions desirable? Looking at the problem more closely, however, reveals that this argument is seriously flawed.

✉ Simon Scheider
s.scheider@uu.nl

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

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

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

For one, there is an *apparent inability* to remove the human from the loop [20, 44, 59]. For example, image recognition may fail miserably in unforeseen ways when images are manipulated or for unforeseen visual scenes¹. It is currently highly doubtful whether human intervention will ever become obsolete for self-driving cars [20]. To prevent errors, humans frequently need to intervene at discrete moments in AI algorithms, even if they were in principle designed to substitute human practice. This applies also to OpenAI's novel dialogue system *ChatGPT*². Therefore, in reality, the human role is not substituted but rather *shifted*. We seem to run into the so-called *paradox of automation* [4]. The more developed an automated system is, the more attentive and skilled human controllers need to be. However, since human practices are substituted with ML algorithms, human operators lack practice and are, thus, becoming less and less skilled [20].

Furthermore, it becomes more and more apparent that data-driven models *fail to capture human-level intelligence* in important respects required to solve certain problems. This includes both human capacities of conceptualisation [60] as well as the use of data [20, 54]. Yann LeCun, one of the leading researchers in the field of deep learning, has recently started questioning a purely data-driven approach to AI [54]. He argues that general (artificial) intelligence requires the ability to account for *levels of detail* and for *simulations using world models*³. According to LeCun, the two main approaches of data-driven AI, namely supervised (deep) learning and reinforcement learning, are dead ends, at least when pursuing the goal of general AI. Correspondingly, cognitive scientists such as Gerd Gigerenzer [20] have warned that the flexibility of ML solutions, which is the main reason for their success and which was meant to prevent bias errors [31], comes at a high price. In contrast to cognitive models, unbiased ML models become *brittle* and, thus, cause errors in unforeseen situations, which go unnoticed in the data [22]. As a consequence, human heuristics can outperform complex models precisely because of their inherent bias [21]. Thus, in contrast to what one might assume, human cognition and its inherent biases play an essential role in handling a particular kind of uncertainty that goes unnoticed by current AI models. From the viewpoint of information science, this uncertainty manifests itself in a *lack of knowledge about the usefulness of data for a given purpose* [70]. It has at least three dimensions [20, cf. chapter 5]:

1. *our world* described by the data is *unstable* and thus concepts need to change.
2. *human conceptualizations* or *good theories* needed to *interpret the data* are not well understood or remain unclear.
3. *procedures* used to generate data (*provenance*) and the *purposes* of using data remain intransparent, and, thus, the data's *quality* cannot be assessed.

In this discussion article, we argue that the uncertainty about data quality and use is a consequence of a lack of pragmatic models. What is largely missing in today's (geo)AI is what philosophers call *pragmatic knowledge*, i.e., knowledge of information practice [42]. Current data-driven GeoAI therefore becomes blind to exactly those practices that it tries to substitute. As a result, accuracy measures become misleading, because they assume these practices without mastering them. To address this blind spot, GeoAI needs to become *pragmatic*. It needs to put possibilities, purposes and procedures underlying geographic information in the centre of modelling, at least to a degree that allows controlled interactions between GeoAI and human experts.

In the remainder, we first illustrate this challenge using a typical example of geographic information practice. We then argue that the reasons for the discussed blind spot emerge from a particularly popular background philosophy of AI, which we call *structuralist AI*. We discuss the shortcomings of the structuralist approach to GeoAI in terms of 8 challenges. This is followed by a definition of what GeoAI should be, and a sketch of an alternative methodical approach we call *pragmatic GeoAI*. At its core, it is based on an information-theoretic action model, which explains what kinds of knowledge need to build on each other. Finally, we discuss how such an approach might be used to handle the previously identified challenges.

2 The Problem in a Nutshell

The inability of substituting information practices by machine learning due to missing concepts in the data, and due to a corresponding lack of understanding of the purposes underlying data, is a recurrent topic in AI. Take the well-known example of the *Bongard* problems (Fig. 1), discussed in [32].

The task is to find a principle that distinguishes the images on the left from those on the right. In the illustrated case, the concept in question is *convexity*, but there are unlimited further possibilities. As [60] explains, current ML approaches not only fail to learn solutions from data beyond any specific principles, they also rely on (large amounts of)

¹ Cf. e.g. Uber's fatal experiment with a self-driving car in 2016, and numerous other examples collected in [20, 59].

² <https://openai.com/blog/chatgpt/>.

³ By the way, this is an insight gained decades ago in cognitive science and linguistics [5, 45].


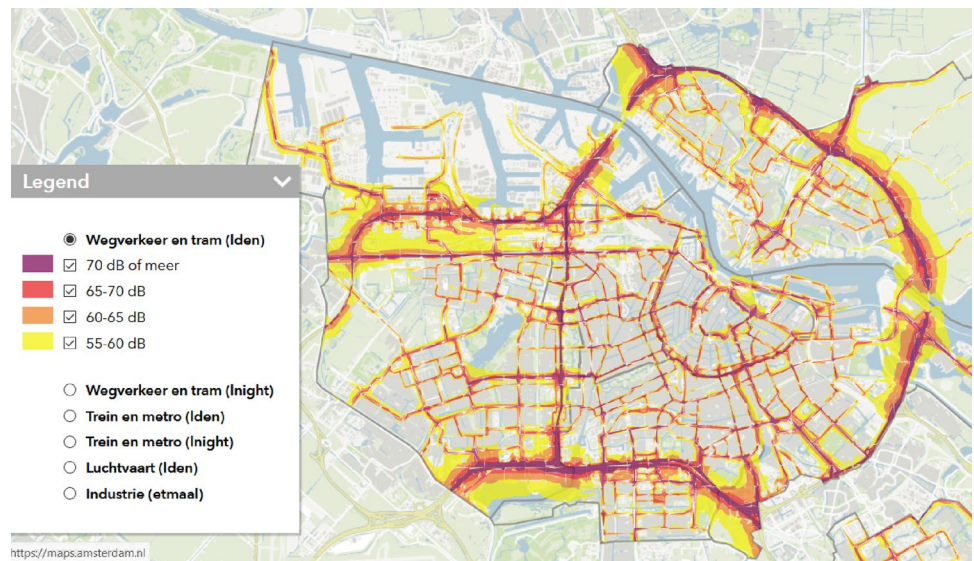


Fig. 1 Bongard problem: the left side shows only convex shapes, as opposed to concavity shown on the right side

Fig. 2 Map of noise contours in Amsterdam. Source: Amsterdam municipality, <https://maps.amsterdam.nl/geluid>



Bongard problems are similar to the problem of handling geographic information in terms of *maps*. Maps require *interpretation*, *generation* and *composition* from diverse origins [63, 89]. In principle, there is potential for GeoAI methods to make geographic information better accessible and usable by automating the interpretation, retrieval and composition of maps [14, 30, 50, 81]. However, it remains unclear how precisely machines could substitute the interpretation and transformation skills of GIS analysts based on data-driven methods. For example, to quantify the effect of noise on the health of citizens, we may need to retrieve and transform a noise contour map (Fig. 2) into a statistical summary to derive the *proportion of the area covered by 70dB noise in Amsterdam*. The latter information serves

data examples⁴. In contrast, human solvers can figure out new concepts based on seeing a single example. This suggests that the practice people use to solve these problems is different from general-purpose pattern recognition [59]. It involves a *repertoire of concepts that can be applied* to the given sketches by problem inventors and solvers alike, and which can be searched for equivalences. Note that this repertoire is not in any way “contained” in the data. Instead, it is part and parcel of the practice of any competent human interpreter of geometric figures. Furthermore, the repertoire is extensible, as there are unlimited *possibilities* for coming up with Bongard problems. Finally, the underlying concepts, such as convexity, can be considered the *purposes* of Bongard problems, namely to learn what distinguishes sides.

the purpose of assessing the health conditions of living in Amsterdam. Deriving this information is a typical GIS task. How should a map like the one in Fig. 2 be interpreted for this purpose? Which transformations are possible and meaningful for this purpose? And how can we know that they satisfy this purpose?

Possible practices of solving this problem manually with GIS software are illustrated in Fig. 3. These practices reveal that *having the data* or the *analytic algorithms* is not even close to being sufficient for solving this task. More precisely, it is not enough to know that the map is in (vector) polygon format covering certain kinds of regions in space, nor that noise values consist of integers within a certain range. Instead, we would need to know that the map can be interpreted as a *spatially continuous field* represented as contours, and not as a *collection of objects* [51]. Yet whether a map can be interpreted in this way is neither contained in the data nor is it generally known (and thus could be retrieved

⁴ While it is possible to train an ML model to recognize a *particular* principle, Bongard problems require picking from an *extensible repertoire*.

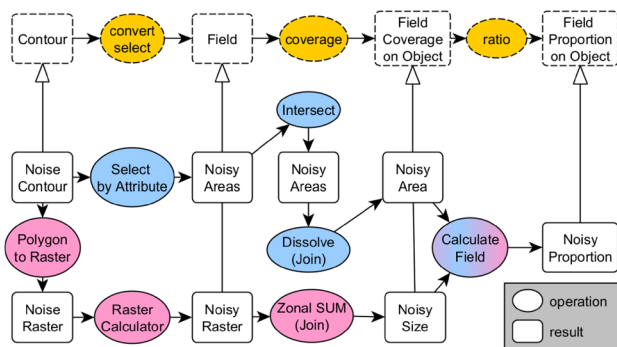


Fig. 3 How to transform a noise contour map into a measure of the proportion of noisy area? The dotted *transformations* are on a *conceptual level*, in parallel to different computational procedures on different data structures (blue: vector version, red: raster version). Procedures are schematized versions of real GIS workflows (color figure online)

as a fact) [71]. At the same time, conceptualisations of geographic quantities [80] inform us that aggregations cannot be counts or densities of objects, but should be field integrals or field coverages, measuring the area covered by some interval over the noise field [71] (cf. dotted boxes in Fig. 3). We furthermore need to know that such aggregations could be implemented in terms of both raster (zonal map algebra) as well as vector overlay (combining, e.g., intersect and dissolve methods) (cf. schematized workflows in Fig. 3). In summary, the knowledge needed to solve this task goes well beyond data, data structures, and algorithms: On the one hand, it goes beyond what and how data is encoded in a map, by drawing on a *repertoire of concepts* that require interpretation. On the other hand, it goes beyond the knowledge of map algorithms, by drawing on a *space of possibilities for transforming concepts*. Finally, data-driven approaches fail because *gathering data about these practices* is very hard [63]. It requires expertise, and thus presupposes exactly those skills that we intend to substitute. We thus run into a severe *cold start problem*.

There is a lesson in modesty underlying this. If AI is supposed to “revolutionize” the geosciences, it should be able to deal with this simple example of everyday GIS practice. However, current approaches to GeoAI fail to handle not only the required interpretations of maps in terms of information concepts, such as *fields* and *proportions*, but also corresponding transformation practices [72], and as a consequence, cannot deal with analytic purposes.

3 Structuralist AI

These apparent limitations of data-driven AI are reflected in certain legacies in thinking about information in general. These legacies still influence and constrain our modern understanding of AI. The subsequent explanations largely follow the argumentation of Janich, cf. [42].


3.1 Structuralism and Related Legacies

For one, there is the legacy of *naturalism*. In this background philosophy, information appears as a quantity occurring in nature extractable with scientific means. In Shannon’s information theory [73], e.g., information is treated as if it was obtainable from signals using thermodynamic entropy, flowing like energy through nature from a sender to a receiver. Yet, any modern-day talk about information *in* nature is nothing more than a metaphor. In scientific practice, instead, information *about* nature is always dependent on the technological practice used to obtain it [39, 41]. Our understanding of nature is therefore a sophisticated abstraction of culture via technical means, e.g., via standardized measurement scales and repeatable experiments [39].


A second and related legacy is (scientific or epistemic) *structuralism* [75]⁵. This background philosophy assumes that theories could be reduced to the structure and relation of symbols abstracted from their content. The old idea that scientific knowledge could be handled as a purely formal structure without any grounding [12, 75], in terms of a “self-stabilizing” network of concepts with free interpretations (cf. the critique in [18] and [29]), finds its modern-day equivalent in the idea of *knowledge graphs* [15] (see below). Yet, structuralism underestimates culture and practice as a source of scientific knowledge, in particular, as a source for establishing reference, trans-subjectivity and reproducibility of information [42].

A third legacy is the *mechanization* of communication. The undoubted success of the mechanization of signal transmission may lead us to think that the human communication process constitutive of information can be similarly mechanized. However, in doing so, we tend to abstract away dimensions critical for the success of communication [69], such as the underlying *intentions* and *purposes* as well as the *actions* and *practices* which are means to such ends. For instance, the map in Fig. 2 was designed to determine the spatial distribution of a noise field. What we can do with the map, e.g., summarize noise, needs to *correspond* to this purpose. These aspects, falling under the term *information pragmatics*, typically get lost in the course of mechanization.

⁵ Not to be confused with structuralism in social science.



(a) Pyramid of layers of information from a structuralist standpoint.



(b) Pyramid of layers of information from a pragmatic standpoint.

Fig. 4 Knowledge pyramids with layers of information

To summarize, hoping that syntax might constitute semantics, and that semantics might be used to handle pragmatics (giving rise to the information pyramid as depicted in Fig. 4a, cf. [61]), has demonstrably not worked in practice [11]. We believe this is a misconception due to the legacies of naturalist, structuralist, and mechanistic thinking about information. To make progress, we need to take pragmatics seriously. We cannot dispense with the pragmatic dimension as the most fundamental level of information, from which other aspects such as semantics and syntax need to be abstracted. This turns the pyramid upside down (Fig. 4b). Note that this new ordering reflects how information is constituted, not the sequence of its occurrence. In practice, we are frequently moving up and down in the pyramid so that symbols might feed into actions and back. Yet, the inverted pyramid implies that we can do so only because action possibilities constitute any semantic concept denoted by a symbol, even if only implicitly. And as the pyramid gets thinner towards its top, the following challenges appear as mere consequences of a pragmatic deficit.

3.2 Pragmatic Challenges for AI

The structuralist view of information has become predominant in AI and data science. Paraphrasing this view, we might say that *information is a structure observed in nature. Data is a way to share such structures.* Therefore, information about nature *can be read out of structural patterns in data.* In consequence, *communication is just data exchange,*

and *intelligence becomes just pattern recognition.* For example, in the *representation learning paradigm* [6], “information” is a pattern of vectors learned from examples such that the distances between vectors are tweaked to capture conceptual similarity in a highly context-sensitive manner. Yet, while these methods can be immensely useful to automate classifications and retrieval, we are struggling with their information theoretic problems:

Abundance of data, but lack of shared procedural and pragmatic knowledge

Challenge 1 The information needed to make use of data is not contained in this data. It consists of a (tacit) repertoire of shared information practices, including procedural knowledge, purposes and requirements (pragmatic knowledge).

This represents a genuine logical problem for structuralist AI: to serve a given purpose, *trying to extract* the missing information from the data is not feasible, because the data misses essential concepts needed to interpret the data, including its purpose. Yet, *trying to add* the missing information by learning meta-data is running into the same problem. Therefore, current AI approaches, though themselves very useful, largely fail to account for the usefulness of data. When we ask what kind of knowledge is missing, we need to refer to *repertoires of shared information practices* as the bottom of our knowledge pyramid. Yet the core of structuralist AI is formed by data or formal structures devoid of the underlying generating procedures. This also explains why we need to distinguish *externalized* and *internalized* sources of knowledge. Only part of the procedural knowledge underlying the use of data is made explicit in data or in accompanying documentation, a lot of it is internalized, tacit knowledge. However, though implicit, this knowledge can still be shared within a community. For example, people can understand data produced by another person by recapitulating its generation procedures. Furthermore, they can account for the usefulness of this data by understanding the purposes of procedures.

Lack of reflection and imagination

Challenge 2 Reasoning with procedures requires reflection, i.e., procedures for reasoning with other procedures, their purposes and requirements.

Challenge 3 To assess the possibilities of data, it is necessary to understand purposes and to reason with procedures that would satisfy such a purpose (imagination).

To model the use of data, we need to reflect on what is possible, not on what was done. We need to reason with information procedures, not simply with data. Furthermore, we need to be able to simulate such procedures

without performing them. Both require making use of procedures themselves without committing to their application. For example, in Fig. 2, figuring out that there are two ways to generate the goal map requires reasoning over GIS procedures. This corresponds to reasoning with higher-order functions (functions that take other functions as input). Since structuralist AI models data and not the procedures that generate its structures, it is difficult to simulate possibilities and, thus, to deal with reflection and imagination (“what would happen if I do this?”).

Lack of reproducibility, reusability and interoperability

Challenge 4 The reproducibility crisis is due to a lack of understanding of the provenance of results. Yet knowledge of provenance requires procedural knowledge to be machine readable/reproducible.

Challenge 5 Missing reusability and interoperability indicate a lack of understanding of the possibilities, and, thus, of the requirements and purposes of procedures.

One of the weaknesses of current approaches to AI is related to the reproducibility/reusability crisis [34]. Models (especially ML models, but also computational models) are rarely reused because it is unknown whether they can be reused or trusted in a different context. This problem has triggered various reparation initiatives [85]. However, at its core, there seems to be a fundamental pragmatic problem.

Lack of modularity and composability

Challenge 6 Modularity and composability of workflows require understanding the possibilities, and, thus, both purposes and requirements of procedures within and between modules.

For similar reasons, the current approaches to AI make it hard to modularize and compose software solutions. Currently, models and software solutions are distributed over millions of developments forming redundant resources. How is it possible to know that a certain model or computational function solves a particular problem such that it can be reused in combination with others in a workflow for a certain purpose [70]? Structuralist AI inherently lacks data on procedures and purposes. Furthermore, even if there was data about procedures available, how to make sure it is not biased according to certain habits and, thus, misses out on options that were never tried? Again, the way to deal with the challenge requires an in-depth understanding of what is possible, beyond what was done.

Lack of quality, value and validity

Challenge 7 To discover the quality and value of data, it is necessary to model the requirements and results of the procedures of data generation and use and to compare them with purposes. Furthermore, the validity of scientific claims requires quality criteria defined with respect to purposes of data analysis.

Concepts of *data quality* (such as completeness, accuracy, and level of detail) form requirements and are realized as results of the procedures of data generation and use. For example, resolution is a result of measurement, whereas accuracy and completeness are a result of comparing measurements [62]. Furthermore, data is *of value* only as a requirement for a specific purpose. For example, the noise contour map is of value for health exposure assessment because it satisfies the requirement of representing an environmental factor for health. Finally, the value of data for a purpose is the very basis for judging whether scientific claims based on data analysis methods can be considered valid. *Validity* is a core notion of the scientific method which is dependent on criteria defined relative to purposes implied by claims [38]. For example, aggregating a noise contour map might be a valid method for supporting claims about the noise exposure of citizens living in a particular building. Yet, such claims are only valid if the resolution, accuracy, the spatial and temporal extent of the map correspond to this purpose. Structuralist AI lacks notions of both purpose and requirement, yet without these notions it becomes very difficult to handle data quality, fitness for use, as well as validity in a data-driven way.

Lack of ethics

Challenge 8 Effective data ethics requires understanding and restricting the potential (mis-)uses of data. This is only possible if we can licence purposes and procedures.

Finally, one of the most pressing current problems of AI, namely data ethics and privacy, becomes problematic in structuralist AI exactly for the same reasons. Data privacy is, in essence, not threatened by data (despite what the term may suggest), but rather by *what can be done with data*, in particular by what goes against one’s interest. For example, a person’s right to affordable health insurance is not threatened by their personal data, but by the usage of this data to estimate their health risks. To effectively protect a person’s privacy, we, therefore, need to restrict potential misuse of the data [84].

4 Pragmatic Framework

Pragmatic GeoAI can build on pragmatic theory which may be reused as a framework. We discuss such sources in the following.

4.1 Theoretic Predecessors

With the "linguistic turn," language moved into the centre of interest of philosophy in the 20th century [86]. However, while the failure to reconstruct the language of science from empirical observations with formal means [65] resulted in methodical relativism and modern structuralism [75], it is easily overlooked that pragmatics formed the starting point of the philosophy of language [3].

In pragmatics, the signs of a language appear only as a superficial structure of underlying *possibilities for actions*. For example, [67] discovered that the competence of language speakers to formulate meaningful sentences requires more than the knowledge of grammar. It implies knowledge of actions. This insight leads to the foundation of *speech act theory* [3, 26], which continues to play an important role in modelling the meaning of sentences in natural language processing (recently e.g. in terms of rational speech act models (RSA) [25]), as well as in the modern theory of mind [17]. Ryle [67] also realized that pragmatics forms a much broader basis underlying any kind of shared human competence. For example, knowing how to play chess requires more than knowing the rules of chess. Knowing the rules is not sufficient to claim that a person is a competent chess player. Ryle suggested that explicit knowledge (*knowing that*) is always grounded in *knowing how*⁶. According to Ryle, *knowledge*, and likewise *intelligence*, occurs in the implicit form of know-how, i.e., it consists of operational dispositions that allow people to do things competently. Furthermore, such dispositions become *trans-subjective* when taught and learned by agents. They can be specified in terms of *procedures*, i.e., instructions built based on such dispositions. In chess, it is only when the game is taught to other players that this know-how is (partially) turned into an *explicit form*. A chess player might explain a strategy in terms of configurations that go beyond the rules of chess, such as moving a bishop to force the opposing queen to cease to threaten one's king. Yet, since know-how generally surpasses the rules of the game and the conditions of its mechanics, it is an error to reduce the former to the latter. So there is always a loss of understanding involved when going from know-how to knowing that.

⁶ Note: not the other way around.

⁷ Not to be confused with *post-structuralist* philosophers.




Fig. 5 Pragmatic action model according to Janich (cf. [38])

In the second half of the 20th century, similar insights motivated *methodical constructivist* scholars⁷ to put action dispositions of humans at the core of their epistemology [37, 56, 82]. Among these, the "Erlangen" school of constructivism is particularly interesting, because it focused on the methodical reconstruction of scientific language starting from elementary operations that can be learned in introductory lessons. From these, more abstract terminology including mathematical concepts [56] can be established based on equivalence operations [47]. Similarly, other researchers embarked on the pragmatic justification of methods in the sciences, e.g., of Physics, Chemistry, Biology, and Psychology, by making explicit the craftsmanship underlying them [41]. Janich's "proto-physics" and "proto-geometry" [35, 36], e.g., reveal procedures for generating artefacts needed for practical geometry and time measurement, without presupposing existing technical measurement standards. We will not go into the details of such methodical reconstructions, but rather introduce basic ideas taken from Janich's pragmatic theory [38] and his theory of information [39, 42]⁸.

4.2 A Pragmatic Model of Know-how and Information

In the following, we explain the basic terminology underlying Janich's *logic-pragmatic propedeutics* [38] (Fig. 5). Regarding the influence and state of discussion of Janich's theory, cf. [8]. An important difference compared to traditional action theory or cognitive architectures, such as ACT-R [2], is that at its core there is not an individual agent, but the *practice of people cooperating in a community*. Another key difference is that it allows for defining what it means for *information to be useful* in this community.

In a community of acting persons, people can *attribute*⁹ actions to people. Furthermore, people can also *direct* actions at one another. Since people can take sides in this, they can train each other to obtain the *capacity to perform actions*. For example, when teaching children how to throw

⁸ As the original texts are in German, a certain uncertainty in our translation of the terminology is inevitable. This is why we add German originals to each technical term.

⁹ 'zurechnen'

a ball, adults may take turns in throwing and *prompting*¹⁰ children to throw the ball.

The capacity to act is called an *action schema*¹¹. As illustrated by the ball-throwing example, such schemas can be shared among members of a community, and thus allow different people to repeat “the same” actions and to recognize others doing so. This is called *trans-subjectivity*. When we attribute actions to people and recognize the schema that is involved, we call this action an *actualization* of this schema.

Action schemas can have *purposes*¹², which are simply further schemas that this schema points to. When people attribute schemas to the actions of a person, they assume that the purposes of those schemas are followed by this person. For example, when we assume a person throws a ball at another person, then we also assume that the former wants the latter to catch the ball. If the latter person catches the ball, then this catching is an action which is an actualization of a catching schema. At the same time, it is a *result* of the throwing. The catching schema, in turn, is the *purpose of the throwing schema*. In case all these four relations hold (the right-hand side in Fig. 5), then we can say that the throwing was *successful*. Similar to purposes, action schemas can have *requirements*, which are schemas that need to be actualized before an action can take place. An action that occurs *before* another action is called a *condition* for this action. Being both an actualization of a requirement and a condition of the actualization of the schema turns a condition into an *explanation*. For example, a teacher can explain why the throwing was unsuccessful by noticing that one of its requirements, namely the absence of heavy wind, was not actualized before throwing the ball.

Success can be defined as the result of an action that corresponds to an actualization of the purpose schema that was attributed to this action. Furthermore, success can be *explained* by actualized requirements in an analogous way. However, action schemas do not necessarily have purposes, and therefore do not always lend themselves to talking about success. Consider the making of art, where such a purpose often remains unclear. Janich distinguishes actions that succeed (‘erfolgreich’) and actions that *work* (‘gelingen’). An action that does not work is one that simply cannot be actualized, regardless of any purpose. For example, throwing does not work if the ball is too heavy for the child, regardless of whether the child had the purpose of throwing the ball to somebody else.

Furthermore, some entities are not actions or action schemas. For example, actualized purposes do not need to be actions, they can also be *artefacts*. When a carpenter

successfully teaches an apprentice how to make a table, and the apprentice successfully performs this action schema, then the result of the actualization will be a table (an artefact). Again, the teaching is successful precisely when this table actualizes the purpose of being a table of a certain quality.

This *mechanism of testing and explaining success* based on checking whether conditions satisfy requirements and results satisfy purposes takes on a central role in human learning (Fig. 5). Being able to explicitly recognize artefacts, actions, and their schemas, as well as the way they are related in terms of purposes, requirements, conditions, results, and attributions to people, is a capacity which can be taught by competent speakers of a language community by giving feedback on success (and switching roles in this), in the same way that they can teach each other to become carpenters or to throw a ball. It is not only the mechanism by which members of a language community can make sure their schemas become trans-subjective. It is also the way how they organize cooperation and, thus, how *information* comes into the world [42]. Information derives from the verb *to inform*. More precisely, *speech*¹³ is an action attributed to some person, which is directed at other people, and which *prompts* some schema. Speakers *understand* each other when this prompting is successful, in which case we can also say they successfully *informed* each other. For example, when we inform someone about a meeting “at 17:00 at home,” we are prompting the use of a particular calendar and a particular place (schema), and we know our prompting was successful when the right schema was used and the person shows up in time. A community of speakers and their success in performing such acts are the objects by which any theory of information must prove itself useful, whether or not some of them are externalized in the form of information artefacts¹⁴ (written text, computer code, gestures, speaking robots, etc).

To summarize, Janich’s theory provides a way to explain the usefulness of information for purposeful action. In a nutshell, *information is useful if it is successful in prompting actions, which would generate some result that actualizes some purpose schema*.

5 Pragmatic GeoAI

In the following, we discuss what a non-structuralist alternative of GeoAI might look like, which builds on geographic information practice and is centred around the action model sketched above.

¹⁰ ‘auffordern’

¹¹ ‘Handlungsschema’

¹² ‘Zweck’

¹³ This includes not only audible speech but all forms of utterances.

¹⁴ Janich calls this “technical substitution of communicative competence” [40].

5.1 What is GeoAI?

In current approaches to GeoAI [43], the value of a-priori knowledge is considered obsolete given more data¹⁵. As a result, and despite efforts towards “spatially explicit modelling” [43], GeoAI is seen largely as AI on georeferenced data [55]:

Definition 1 *GeoAI is the practice of AI applied to georeferenced data as used by geoscientists.*

In turn, the practices of geoscientists are seen as a tradition standing in the way of data-driven innovation (“just give us your data, we are not interested in your practice”). We believe that Definition 1 is truly misleading because it falls prey to all the structuralist problems discussed above. So what would be a better way to define GeoAI?

Following the inverted pyramid (Fig. 4b), we should instead hook our definition of GeoAI into the levels of practice, and then work our way towards the data. Georeferenced data is not the basis, but rather requires cartographic and semantic interpretation, e.g., in terms of a theme, resolution and extent. To leave open whether we talk about digital representations or analogue maps, and whether for the purpose of visualisation or computation, we just use the term *map* in the following:

Definition 2 *GeoAI consists of all geographic information practices mastered by geoscientists using maps, which a computer cannot (yet) master.*

While this definition does justice to the fundamental role of practice, it constrains GeoAI to the practices of the past and of a particular group (geoscientists). However, GeoAI enables us to invent entirely new practices in disciplines outside of the geosciences. To account for this, we simply adopt *potential practices* in our definition:

Definition 3 *GeoAI consists of all geographic information practices that make use of maps to satisfy some information purpose, but which a computer cannot (yet) master.*

Note that information purposes and possibilities become central here. Furthermore, since practice (as a form of higher-order knowledge) is always rooted in other practices, it follows that substitution must always be partial:

Remark 1 *Since geographic information practices are always rooted in other human practices, GeoAI can only*

¹⁵ “GeoAI research will have to make a case for spatially explicit models”.




Fig. 6 Core action model for GeoAI

partially substitute human practice. Human practice is therefore never obsolete, yet it can take on new roles.

In this article, we suggest discarding Definition 1 and adopting Definition 3 with Remark 1.

5.2 Predecessors of Pragmatic GeoAI

Some earlier work has followed a pragmatic approach to GeoAI and geographic information. For instance, Brodaric [7] recognized that many concepts depicted in geological maps are dependent on the historical context in which they evolved, such as “Baker Brook basalt.” Similar to *place* in human geography, this makes semantic concepts in geology *situated*, i.e., dependent on history. Yet, here, pragmatic information is considered from a naturalist viewpoint (evolution vs. design), and only as an addition to semantic information. In contrast, Helen Couclelis puts the notions of *purpose* and *design* at the center of her theory of geographic information [10]. Yet, purposes and practices still form the highest level of abstraction in this theory [11]. Thus, both approaches appear to commit to a knowledge pyramid of a structuralist flavour (Fig. 4a). Pragmatics has played a larger role in the context of wayfinding research [16, 79, 83], where it was applied to the specific case of route instructions. In summary, the question remains how a general model of the pragmatics of spatial information should look like, how it would apply to our example from Sect. 2, and how this could help address the challenges discussed above.

5.3 A Pragmatic Model of Knowledge in GeoAI

Following the idea of the inverted knowledge pyramid and Janich’s action model introduced in Sect. 4.2, GeoAI requires the explicit modelling of *actions of informing someone about the geography of some phenomenon*, e.g., about the geography of noise in Amsterdam. Actions are performed on particular artefacts, called *maps*, to satisfy this purpose. Knowledge of geographic information therefore simply means *knowing how to transform*¹⁶ *maps according*

¹⁶ With the term transformation, we denote any derivation of map artefacts, not just coordinate system transformations.

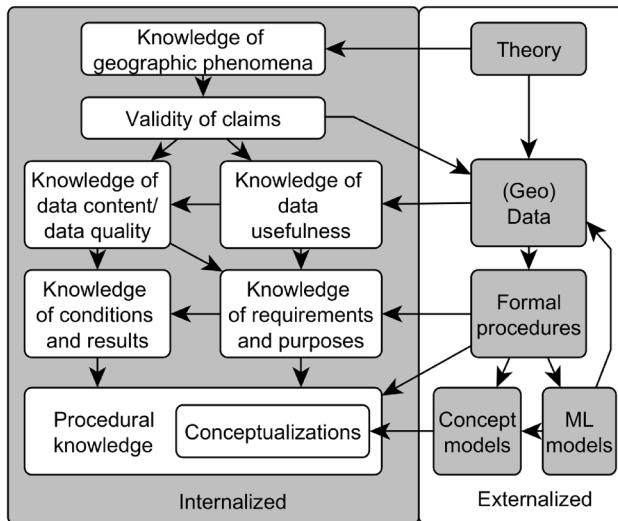


Fig. 7 Pragmatic knowledge sources in GeoAI. Arrows denote dependencies between different forms of internalized/externalized knowledge. Since procedural knowledge is specific for geographic information, GeoAI is more specific than general AI. ML appears as part of a particular externalization strategy that relies on data

to some purpose (see Fig. 6). This implies that in GeoAI, action schemas become *transformation schemas*, artefacts become *maps*, actions become *tool applications* (because tools implement transformations), and schema actualizations become *map interpretations*. Furthermore, a *successful transformation* implies that the resulting map is an actualized purpose of a transformation schema. In essence, *pragmatic GeoAI* means to be able to reason over transformations of maps.

Based on different kinds of reasoning, we can highlight different forms of knowledge including their dependencies (Fig. 7).

The first and most important distinction is the one between *internalized* and *externalized* knowledge. This acknowledges the fact that the most relevant form of geographic knowledge is tacit, and thus not externalized in terms of map symbols or other representations [71]. For example, the fact that the noise map in Sect. 2 is in polygon vector format is explicitly represented in the meta-data, but the fact that these polygons represent contours of a noise field is usually implicit knowledge.

Let us first highlight the different forms of *internalized knowledge* in this example. *Conceptualizations* (Fig. 7) correspond to actualizations of schemas in map artefacts, e.g., the interpretation of the noise contour map in terms of the core concept *field*. With *procedural knowledge*, we denote knowledge of *transformation schemas* that may be concatenated to form larger schemas. For example, “measuring the spatial coverage of a field” is a transformation schema that

consists of many steps transforming a field into a measure of the space it covers.

On the next level, *knowledge of results and conditions* as well as of *requirements and purposes* depend on knowledge of conceptualizations and transformation schemas (Fig. 7). More precisely, transformation schemas can be actualized (performed) and then give rise to *tool applications* which have *conditions* and *results*. For example, “measuring the spatial coverage of a field” can be performed using the tool *zonal map algebra*, which has some raster map as a condition and some vector map as result, or else by some equivalent transformations on vector maps. By connecting conceptualizations to transformation schemas, we can refer to their *requirements* and *purposes*. The requirement of the transformation schema “measuring the spatial coverage of a field” is that the condition represents a field (e.g. of noise), and its purpose is a resulting map interpreted as the “area covered by this field within some object.”

Based on this kind of knowledge, we can introduce knowledge of *data content* and *data quality* on the next level. Knowledge of data content means knowing that a result was generated by a particular procedure starting from a particular condition. For example, since we know that the result in the workflow of Fig. 3 is generated from a noise contour map, we know it is a map of noise. More precisely, procedures can be used to *concatenate* results, and, thus, to *parameterize* result schemas with input schemas. In this way, we can infer, e.g., that the resulting map depicts the “area covered by noise > 70 dB within postcode areas” because the input field of the transformation can be conceptualized as “70 dB noise” and the input objects as “postcode areas.” If schemas are *higher-order* they allow us to reason across transformations. Likewise, knowledge of how this map was produced gives us knowledge of its data quality. For example, to determine its resolution, we need to know which cell size was used when generating the raster map, or which vector regions were used when aggregating noise.

If a transformation result cannot be interpreted in terms of the purpose, then the transformation is not successful. This way we can learn whether the resulting *map is useful for this purpose*. This level of reasoning is based on comparing both data quality and data content with a given purpose or requirement. For example, if our purpose is a map of noise coverage for city blocks, but our result is on the level of postcode areas, then we can *explain* that our result is not useful because it misses the requirement of city blocks. Finally, if we can assess the usefulness of a map transformation for a purpose, we can define *criteria for the validity* of claims [38] made based on this map, and thus justify whether such claims can be considered *geographic knowledge*. Such claims may be about static geography, but may also involve geographic causality [1, 23] and explanations of geographic phenomena in terms of constituting processes [19].

The problem of data-driven GeoAI is that only small parts of these sources of knowledge are *externalized*, and, thus, available for supervised learning. We can externalize procedural knowledge in terms of *formal procedures*, e.g., algorithms. In our example, the transformation schema *measuring the spatial coverage* can be implemented in terms of a *zonal map algebra function*, which can be executed on raster data. Furthermore, maps are *artefacts* externalized in terms of vector or raster data and as such depend on measurement procedures. The conceptualisation of a map as a noise field can be externalized as a measurement or estimation procedure for noise [13]. For example, we could estimate the amount of noise using a supervised ML model. Whether this externalisation is successful depends on whether we would come to similar conclusions when using measurements. Our model then depends not only on an externalization of the noise concept (a *concept model*) but also on data produced according to some formal procedure. Finally, *theories* are artefacts that are dependent on data artefacts. For example, a theory of core concepts can be used to annotate data examples and can be externalized as a *knowledge graph* [71]. However, note that in our model, theories also depend indirectly on knowledge of the usefulness of data for the purposes given by the theory. In consequence, they depend on the entire implicit knowledge pyramid, including the full stack of conceptualizations and procedural knowledge which underlies them.

6 Discussion and outlook

Finally, we discuss to what extent the proposed core model of pragmatics may be capable of addressing the problems presented in the introduction and the challenges of Sect. 3.2. We use examples from GIScience and AI to illustrate possible technical approaches as well as areas of future work.

Challenge 1 implies that the information needed to make use of geodata is not in the data. Therefore, it cannot be extracted by a data-driven GeoAI. From the viewpoint of pragmatic GeoAI, this is rather a theorem. It follows from the fact that knowledge on which GeoAI relies is mainly internalized know-how of informing somebody about the geography of some phenomenon. Thus, solving this problem requires not starting with geodata but to *externalize this procedural knowledge* as far as possible. This requires pragmatic modelling techniques for both conceptualizations and transformations.

Modeling conceptualisations *Interpretability*, *consistency* and *limited labeling* appear among the top challenges of deep learning for the geosciences [66]. Papadakis et al. [64] recently requested a “clear reasoning path from data to conclusions” for explainable GeoAI (X-GeoAI), similar to traditional geo-analysis, by combining concepts with

statistics (cf. also [87]). Unfortunately, we are still far from knowing enough about the concepts needed for modelling purposes of geographic information [51, 52], at least to an extent which would enable GeoAI. The role of core concepts of spatial information for map selection was recently studied with an online experiment in [63]. To implement *conceptual models* [28] *formal ontologies* may be used. We have recently tested an OWL ontology of core concept data types for annotation and data retrieval [71]. However, we currently face two methodological challenges. One concerns the standardization and automation of the annotation of map resources with such concepts. Annotation instructions are needed to generate high-quality annotations (with a high inter-annotator agreement) for such concepts [63]. Based on these, supervised methods, such as *vector embeddings* [27] on texts or graphs can be used to learn and expand the annotation of documents and data. A second challenge relates to the large *variability* of pragmatic concepts, which is a challenge for (manual) ontology design. Take, e.g., the variability of a concept like “the area covered by this field within some object” applied to objects ranging from buildings to countries and to fields from temperature to noise. This can be addressed using transformation models as a *vehicle for generating conceptual possibilities*, i.e., the space of possible conceptualizations, as described below. Based on this idea, we have recently developed a grammar that can be used to interpret geo-analytical questions as concept transformations [88].

Modeling transformations In the past, *scripts* and *frames* were proposed in AI to capture knowledge in terms of storylines [68]. More recently, *process models* have been proposed to describe business workflows or development processes [58]. However, such models target stereotyped storylines or actual workflows, not transformations of (geographic) information concepts according to some purpose. One way to model the transformation of geographic information would be to *simulate* it with a learned model of the map artefact. This comes close to the suggestions made by [45, 54], however, we currently lack any methods that would use simulations for assessing the space of possibilities provided by geodata. The relevance of process simulations for AI models in the earth sciences has recently been discussed in [66]. Another option is to use *planning* theory, which can be used to infer *possible transformations* based on reasoning. *Loose programming* [53] may be useful in this respect, which is a way to search for GIS workflow graphs satisfying some goal, given a model of transformation steps specified by classes of inputs and outputs taken from a semantic taxonomy [48]. This method has already been used successfully

for modelling possible GIS transformations [50]. To address the challenge of conceptual variability, we have developed and tested a concept transformation algebra¹⁷ for combining and transforming core concepts [76]. The underlying *bounded parametric polymorphic type system* can be used for concept inference¹⁸ by propagating information about concept types through transformation graphs. This technology can be used to construct concepts, to reason over possible conceptual transformations, and to query GIS workflows [76].

Challenges 2 (about reflection on procedures) and 3 (on imagination of data possibilities) are closely related to how procedural knowledge is modeled. If we can reason over *possible conceptualizations* based on *possible transformations* in GIS, then we obtain a way to imagine what is possible with a geodata set, i.e., which questions can be answered in more or less direct ways [72]. Furthermore, as mentioned above, the needed reasoning may be realized based on parametric polymorphism [76]. This gives us a way to reflect on (reason over) procedures. Thus, powerful transformation models can provide ways of dealing with the first three challenges at the same time.

Handling challenges 4, 5 and 6 (reproducibility, interoperability and composability/modularity) might be addressed using higher level pragmatic knowledge as defined in our model. Reproducibility of maps requires procedural models, both for modelling data provenance as well as conceptual theory, as acknowledged in recent work [85]. Assessing the interoperability of maps requires models for generating potential workflows that make use of those maps. Furthermore, recent work [89] on composability of maps stresses the role of purpose-oriented refinement of mashups, as opposed to strict compliance with standards. Thus, models of purposes would allow assessing how a map could be reused in a different context.

Regarding challenge 7 about data quality, we argue that a pragmatic model of spatial data quality would be more general than current approaches. Spatial data quality is usually conceived in terms of *uncertainty*, e.g., in terms of statistical error models or fuzzy values. Yet, we know that a model of provenance would give us the advantage of analysing error in terms of *uncertainty propagation*, as illustrated, e.g., in the work on *measurement-based GIS* [24]. This insight does not only apply to accuracy but also to other spatial data quality dimensions, such as resolution and completeness. In general, assessing spatial data quality is best modelled in pragmatic terms since quality is a consequence of data generation according to some purpose.

Challenge 8 about geoprivacy and geodata ethics was approached in the past mainly based on licensing or obfuscating sensitive data [49]. Yet, as argued in [84], a more effective and less restrictive way of protecting location privacy would be to restrict the potential use of private location data instead. This requires licensing (and thus modelling) purposes of data use depending on whether they run against one's interests. To realize such a model of privacy and ethics requires procedural and pragmatic knowledge of potential data use.

7 Conclusion

Pragmatic GeoAI addresses major issues regarding the use of geographic information: the procedural knowledge needed to account for data provenance and data possibilities, the purposes and requirements of map transformations, and the conceptualizations needed to interpret maps. These are blind spots of current GeoAI, due to a legacy of misleading background philosophies including structuralist, naturalist, and mechanized views of information. Correspondingly, ML models, such as vector embeddings of knowledge graphs and texts, are based on structural similarities in the data, and, thus, need to handle both pragmatics and semantics based on syntax. In contrast, pragmatic GeoAI turns the knowledge pyramid upside down, by building data on concepts and concepts on actions. This is needed to overcome brittle models and to handle uncertainty with the necessary biases. Though pragmatic GeoAI cannot directly address unstable worlds, it contains knowledge on *how* conceptual changes can be accommodated for. We have suggested a core action model for pragmatic GeoAI and discussed geo-information examples showing how the challenges of structuralist GeoAI might be addressed. Though ML and georeferenced data will still play an important role in externalizing conceptualisations of maps, they are less central in this approach, while transformation models become essential. Future work should concentrate on further developing and testing pragmatic models of geographic information, instead of throwing out the baby of practice with the bathwater.

Acknowledgements We are very thankful for the feedback from two anonymous reviewers. This article also owes a lot to the work of Peter Janich and to contemporary investigators of biases, such as Gerd Gigerenzer. Important ideas were developed in discussions at the COSIT 2022 conference, and within the QuAnGIS project which is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 803498).

Data availability There is no data coming with this publication.

¹⁷ <https://github.com/quangis/transforge>.

¹⁸ <https://github.com/quangis/cct>.

Declarations

Competing interest The authors declare that no conflict of interest is involved with this publication.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Akbari K, Winter S, Tomko M (2021) Spatial causality: a systematic review on spatial causal inference. *Geograph Anal.* <https://doi.org/10.1111/gean.12312>
- Anderson JR, Matessa M, Lebiere C (1997) Act-r: a theory of higher level cognition and its relation to visual attention. *Human-Comput Interact* 12(4):439–462
- Austin JL (1975) How to do things with words. Oxford University Press, Oxford
- Bainbridge L (1983) Ironies of automation. In: Johannsen G, Rijnsdorp JE (eds) Analysis, design and evaluation of man-machine systems. Elsevier, pp 129–135
- Barsalou LW (1999) Perceptual symbol systems. *Behav Brain Sci* 22(4):577–660
- Bengio Y, Courville A, Vincent P (2013) Representation learning: a review and new perspectives. *IEEE Trans Pattern Anal Mach Intell* 35(8):1798–1828
- Brodaric B (2007) Geo-pragmatics for the geospatial semantic web. *Trans GIS* 11(3):453–477
- Capurro R (2002) Menschengerechte Information oder informationsgerechter Mensch? In: B. Markscheffel (ed.) Wege zum Wissen—Die menschengerechte Information, pp. 271–287
- Chen W, Fosler-Lussier E, Xiao N, Raje S, Ramnath R, Sui D (2013) A synergistic framework for geographic question answering. In: 2013 IEEE seventh international conference on semantic computing, pp. 94–99. IEEE
- Couclelis H (2009) The abduction of geographic information science: Transporting spatial reasoning to the realm of purpose and design. In: International Conference on Spatial Information Theory, pp. 342–356. Springer
- Couclelis H (2019) Unpacking the “i” in gis: Information, ontology, and the geographic world. In: The Philosophy of GIS, pp. 3–24. Springer
- Da Costa N, French S (2000) Models, theories, and structures: thirty years on. *Philos Sci* 67:S116–S127
- de Kluijver H, Stoter J (2003) Noise mapping and gis: optimising quality and efficiency of noise effect studies. *Comput Environ Urban Syst* 27(1):85–102
- Feng Y, Thiemann F, Sester M (2019) Learning cartographic building generalization with deep convolutional neural networks. *ISPRS Int J Geo-Inform.* <https://doi.org/10.3390/ijgi8060258>
- Fensel D, Şimşek U, Angele K, Huaman E, Kärle E, Panasiuk O, Toma I, Umbrich J, Wahler A (2020) Introduction: what is a knowledge graph? In: Fensel D et al (eds) Knowledge Graphs, pp 1–10 Springer
- Frank AU (2003) Pragmatic information content-how to measure the information in a route. In: Duckham M, Goodchild FM, Worboys M (eds) Foundations geograph information science. Taylor & Francis, pp 47–70
- Frank CK (2018) Reviving pragmatic theory of theory of mind. *Aims Neurosci* 5(2):116
- Frigg R (2006) Scientific representation and the semantic view of theories. *Theoria. Revista de Teoría, Historia y Fundamentos de la Ciencia* 21(1):49–65
- Gahegan M (2020) Fourth paradigm GIScience? Prospects for automated discovery and explanation from data. *Int J Geogr Inf Sci* 34(1):1–21
- Gigerenzer G (2022) How to stay smart in a smart world: why human intelligence still beats algorithms. Penguin UK
- Gigerenzer G, Brighton H (2009) Homo heuristicus: why biased minds make better inferences. *Top Cogn Sci* 1(1):107–143
- Gigerenzer G, Hertwig R, Pachur T (2011) Heuristics: the foundations of adaptive behavior. Oxford University Press, Oxford
- Glymour M, Pearl J, Jewell NP (2016) Causal inference in statistics: a primer. John Wiley & Sons, New York
- Goodchild MF (2002) Measurement-based GIS. In: Shi W, Fisher P, Goodchild MF (eds) Spatial data quality, vol 5. CRC Press
- Goodman ND, Frank MC (2016) Pragmatic language interpretation as probabilistic inference. *Trends Cogn Sci* 20(11):818–829
- Grice HP (1975) Logic and conversation. In: Cole P and Morgan JL (eds) Speech acts. Brill, pp 41–58
- Grohe M (2020) word2vec, node2vec, graph2vec, x2vec: Towards a theory of vector embeddings of structured data. In: Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, pp 1–16
- Guarino N, Guizzardi G, Mylopoulos J (2020) On the philosophical foundations of conceptual models. *Inform Model Knowl Bases* 31(321):1
- Harnad S (1990) The symbol grounding problem. *Physica D* 42(1–3):335–346
- Harrie L, Oucheikh R, Nilsson Å, Oxenstierna A, Cederholm P, Wei L, Richter KF, Olsson P (2022) Label placement challenges in city wayfinding map production—identification and possible solutions. *J Geovisualization Spatial Anal* 6(1):16. <https://doi.org/10.1007/s41651-022-00115-z>
- Hastie T, Tibshirani R, Friedman JH, Friedman JH (2009) The elements of statistical learning: data mining, inference, and prediction, vol 2. Springer, New York
- Hofstadter DR (1979) Gödel, Escher, Bach: an eternal golden braid, vol 13. Basic books New York
- Hu Y, Adams B (2021) Harvesting big geospatial data from natural language texts. In: Handbook of big geospatial data. Springer, pp 487–507
- Ioannidis JPA (2005) Why most published research findings are false. *PLoS Med.* <https://doi.org/10.1371/journal.pmed.0020124>
- Janich P (1976) Zur Protophysik des Raumes. In: Boehme G (ed) Protophysik. Fuer und wider eine konstruktive Wissenschaftstheorie der Physik. Suhrkamp
- Janich P (1997) Das Maß der Dinge: Protophysik von Raum, Zeit und Materie. Suhrkamp
- Janich P (1997) Methodical constructivism. In: Issues and Images in the Philosophy of Science. Springer, pp 173–190
- Janich P (2001) Logisch-pragmatische Propädeutik. Velnbrueck Wissenschaft
- Janich P (2006) Die Naturalisierung der Information. In: Kultur und Methode. Suhrkamp, pp 213–255
- Janich P (2006) Technische Substitution kommunikativer Kompetenz? In: Kultur und Methode. Suhrkamp, pp 275–288

41. Janich P (2015) *Handwerk und Mundwerk: über das Herstellen von Wissen*. CH Beck
42. Janich P (2018) What is information?, vol 55 of *Electronic mediations*. University of Minnesota Press, Minneapolis
43. Janowicz K, Gao S, McKenzie G, Hu Y, Bhaduri B (2020) GeoAI: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond. *Int J Geogr Inf Sci* 34(4): 625–636
44. Janowicz K, Van Harmelen F, Hendler JA, Hitzler P (2015) Why the data train needs semantic rails. *AI Mag* 36(1):5–14
45. Johnson-Laird PN (1983) *Mental models: towards a cognitive science of language, inference, and consciousness*. No. 6. Harvard University Press, Cambridge, MA
46. Kahneman D (2011) *Thinking, fast and slow*. Macmillan
47. Kamlah W, Lorenzen P (1984) *Logical Propaedeutic Pre-School of Reasonable Discourse*. University Press of America, Lanham Maryland
48. Kasalica V, Lamprecht AL (2020) Workflow discovery with semantic constraints: The SAT-based implementation of APE. *Electronic Communications of the EASST* 78
49. Keßler C, McKenzie G (2018) A geoprivacy manifesto. *Trans GIS* 22(1):3–19
50. Kruiger JF, Kasalica V, Meerlo R, Lamprecht AL, Nyamsuren E, Scheider S (2021) Loose programming of GIS workflows with geo-analytical concepts. *Trans GIS* 25(1):424–449
51. Kuhn W (2012) Core concepts of spatial information for trans-disciplinary research. *Int J Geogr Inf Sci* 26(12):2267–2276
52. Kuhn W, Hamzei E, Tomko M, Winter S, Li H (2021) The semantics of place-related questions. *J Spatial Inform Sci* 23:157–168
53. Lamprecht AL, Naujokat S, Margaria T, Steffen B (2010) Synthesis-based loose programming. In: 2010 Seventh International Conference on the Quality of Information and Communications Technology, pp. 262–267. IEEE
54. LeCun Y (2022) A path towards autonomous machine intelligence. Tech rep. <https://openreview.net/forum?id=BZ5a1r-kVsf>. Accessed 11 Nov 2022
55. Li W (2020) Geoai: where machine learning and big data converge in giscience. *J Spatial Inform Sci* 20:71–77
56. Lorenzen P (2013) *Einführung in die operative Logik und Mathematik*, vol 78. Springer-Verlag
57. Mai G, Janowicz K, Zhu R, Cai L, Lao N (2021) Geographic question answering: challenges, uniqueness, classification, and future directions. *AGILE* 2:1–21
58. Mendling J, Reijers HA, van der Aalst WM (2010) Seven process modeling guidelines (7PMG). *Inf Softw Technol* 52(2):127–136
59. Mitchell M (2019) *Artificial intelligence: a guide for thinking humans*. Penguin UK
60. Mitchell M (2021) *Abstraction and Analogy-Making in Artificial Intelligence*. arXiv preprint [arXiv:2102.10717](https://arxiv.org/abs/2102.10717)
61. Morris CW (1938) *Foundations of the Theory of Signs*. In: Neurath O (ed) *International encyclopedia of unified science*. Chicago University Press, Chicago, pp 1–59
62. Morrison JL (1995) Spatial data quality. In: Guptill SC and Morrison JL (eds) *Elements of spatial data quality*, vol 202. Elsevier Science inc., New York, pp 1–12
63. Nyamsuren E, Top EJ, Xu H, Steenbergen N, Scheider S (2022) Empirical evidence for concepts of spatial information as cognitive means for interpreting and using maps. In: 15th International Conference on Spatial Information Theory (COSIT 2022). Schloss Dagstuhl-Leibniz-Zentrum für Informatik
64. Papadakis E, Adams B, Gao S, Martins B, Baryannis G, Ristea A (2022) Explainable artificial intelligence in the spatial domain (x-geoai). *Trans GIS* 26(6):2413–2414
65. Quine WvO (1976) Two dogmas of empiricism. In: Harding SG (ed) *Can theories be refuted?* Springer, pp 41–64
66. Reichstein M, Camps-Valls G, Stevens B, Jung M, Denzler J, Carvalhais N et al (2019) Deep learning and process understanding for data-driven earth system science. *Nature* 566(7743):195–204
67. Ryle G (1949) *The concept of mind*. Hutchinson, London
68. Schank RC (1983) *Dynamic memory: a theory of reminding and learning in computers and people*. Cambridge University Press, Cambridge
69. Scheider S, Kuhn W (2015) How to talk to each other via computers: Semantic interoperability as conceptual imitation. In: *Applications of Conceptual Spaces*. Springer, pp 97–122
70. Scheider S, Ostermann FO, Adams B (2017) Why good data analysts need to be critical synthesists. *Determining the role of semantics in data analysis. Futur Gener Comput Syst* 72:11–22
71. Scheider S, Meerlo R, Kasalica V, Lamprecht AL (2020) Ontology of core concept data types for answering geo-analytical questions. *J Spatial Inform Sci* 20:167–201
72. Scheider S, Nyamsuren E, Kruiger H, Xu H (2021) Geo-analytical question-answering with gis. *Int J Digital Earth* 14(1):1–14
73. Shannon CE, Weaver W (1949) *The mathematical theory of communication*. University of Illinois Press, Champaign
74. Singleton A, Arribas-Bel D (2021) *Geographic data science*. *Geogr Anal* 53(1):61–75
75. Sneed J (1983) Structuralism and scientific realism. In: Hempel C, Putnam H, Essler W (eds) *Methodology, epistemology, and philosophy of science*. Springer, New York
76. Steenbergen N, Top E, Scheider S, Nyamsuren E (2022) Algebra of core concept transformations. *Procedural meta-data for geographic information*. Tech. rep., Utrecht University. osf.io/j6krv
77. Stock K, Jones CB, Russell S, Radke M, Das P, Aflaki N (2022) Detecting geospatial location descriptions in natural language text. *Int J Geogr Inf Sci* 36(3):547–584
78. Sun Z, Sandoval L, Crystal-Ornelas R, Mousavi SM, Wang J, Lin C, Cristea N, Tong D, Carande WH, Ma X, et al (2012) A review of earth artificial intelligence. *Comput Geosci* 159(105034):1–16
79. Tomko M, Winter S (2009) Pragmatic construction of destination descriptions for urban environments. *Spatial Cognition Comput* 9(1):1–29
80. Top E, Scheider S, Nyamsuren E, Xu H, Steenbergen N (2022) The semantics of extensive quantities within geographic information. *Appl Ontol* 17(3):337–364
81. Touya G, Zhang X, Lokhat I (2019) Is deep learning the new agent for map generalization? *Int J Cartograp* 5(2–3):142–157. <https://doi.org/10.1080/23729333.2019.1613071>
82. Von Glasersfeld E (2013) *Radical constructivism*. Routledge
83. Weiser P (2014) A pragmatic communication model for way-finding instructions. Ph.D. thesis
84. Weiser P, Scheider S (2014) A civilized cyberspace for geoprivacy. In: *Proceedings of the 1st ACM SIGSPATIAL international workshop on privacy in geographic information collection and analysis*, pp 1–8
85. Wilson JP, Butler K, Gao S, Hu Y, Li W, Wright DJ (2021) A five-star guide for achieving replicability and reproducibility when working with gis software and algorithms. *Ann Am Assoc Geogr* 111(5):1311–1317
86. Wittgenstein L (2010) *Philosophical investigations*. John Wiley & Sons
87. Xing J, Sieber R (2021) Integrating XAI and GeoAI. In: *GIScience 2021 Short Paper Proceedings*, UC Santa Barbara: Center for Spatial Studies
88. Xu H, Nyamsuren E, Scheider S, Top E (2022) A grammar for interpreting geo-analytical questions as concept transformations. *Int J Geograph Inform Sci*. <https://doi.org/10.1080/13658816.2022.2077947>
89. Yang C, Fu P, Goodchild MF, Xu C (2019) Integrating GIScience application through mashup. In: Wang S, Goodchild MF (eds)

- CyberGIS for geospatial discovery and innovation. Springer, pp 87–112
90. Zhu XX, Tuia D, Mou L, Xia GS, Zhang L, Xu F, Fraundorfer F (2017) Deep learning in remote sensing: a comprehensive review and list of resources. *IEEE Geosci Remote Sens Magazine* 5(4):8–36