

How to talk to each other via computers

Semantic interoperability as conceptual imitation

Simon Scheider and Werner Kuhn

Abstract What exactly does interoperation mean in the context of information science? Which entities are supposed to interoperate, how can they interoperate, and when can we say they are interoperating? This question, crucial to assessing the benefit of semantic technology and information ontologies, has been understood so far primarily in terms of standardization, alignment and translation of languages. In this article, we argue for a pragmatic paradigm of interoperability understood in terms of *conversation* and *reconstruction*. Based on examples from geographic information and land cover classification, we argue that semantic heterogeneity is to a large extent a problem of multiple perspectives. It therefore needs to be addressed not by standardization and alignment, but by *articulation and reconstruction of perspectives*. Concept reconstructions need to be grounded in shared operations. What needs to be standardized is therefore not the perspective on a concept, but the procedure to arrive at different perspectives. We propose *conceptual imitation* as a synthetic learning approach, and *conceptual spaces* as a constructive basis. Based on conceptual imitation, user and provider concepts can be checked for perspectival correspondence.

Key words: semantic interoperability, conceptual spaces, conceptual construction, perspectivity

1 Semantic interoperability: from convention to conversation

Metadata and semantic technology can be seen as cornerstones of a working information society¹. Increasingly, they are being recognized as such. The dominant search engines on the Web, Google, Yahoo, and Bing, have incorporated light weight ontologies and present facts semantically related to a search result². The Linked Data Web helps libraries, governments and museums provide seamless access to their archival data, which have remained information silos for decades (Hyvönen, 2012). Bioinformatics³, Environmental Science (Madin et al, 2007), as well as Geoinformatics (Janowicz et al, 2012) have realized some

Simon Scheider

Institut für Geoinformatik, Westfälische Wilhelms-Universität Münster, Heisenbergstraße 2, D-48149 Münster, Germany e-mail: simon.scheider@uni-muenster.de

Werner Kuhn

Institut für Geoinformatik, Westfälische Wilhelms-Universität Münster, Heisenbergstraße 2, D-48149 Münster, Germany, e-mail: kuhn@uni-muenster.de

¹ Compare the role of metadata in Gray et al (2005).

² <http://schema.org/>

³ <http://zacharyvoase.com/2009/11/02/bioinformatics-semWeb/>

time ago that their data sets of diverse origin and uncontrolled authorship require semantic metadata for successful sharing.

Semantic computing draws on methods from ontology engineering, computational inference, information retrieval, and machine learning of semantic concepts. The purported goal is to help computer systems, and, indirectly, people, interoperate across cultural, domain and application boundaries. However, despite this development, there still remains a central blind spot: *What exactly does interoperation mean?* Which entities are supposed to interoperate, how can they interoperate, and when can we say that they are interoperating? Skeptics have not ceased to question the usefulness of explicit knowledge representations, and for very good reasons (Shirky, 2009). Semantic interoperability as a research goal remains largely underexposed. Central underlying concepts, such as the relationship of information and interoperation have remained obscure and incomprehensible. This leaves the benefit of semantic technology blurry, and negatively impacts its credibility.

In related work on data semantics (Janowicz and Hitzler, 2012; Janowicz, 2012; Scheider, 2012), it has become apparent that semantic engineering of digital information may be regarded as a matter of dynamic *machine mediated communication* between data providers and users, with data (as well as semantic meta-data) being the explicit top of a pyramid of implicit acts of interpretation, observation and construction.

In this article, we argue for a pragmatic shift in semantic interoperability. The proposed paradigm of conversation and reconstruction holds that data has meaning in a pragmatic communication context which implies a *perspective*. The problem is that this perspective is lost, in one way or another, under the conditions of digital communication. Semantic engineering therefore needs to support the communication and reconstruction of this perspective, not necessarily its standardization or translation. The goal is to support users in comparing perspectives and in judging their usefulness for their own (often diverging) goals. We propose to use *conceptual spaces* for this purpose. In this article, we make the case for a new paradigm which is meant to provoke future research into this direction.

In the remainder, we will first motivate the new paradigm by analyzing examples of land cover categories. We will then suggest how information sharing in human-machine-human conversations can be based on *conceptual imitation* and conceptual spaces. The last section demonstrates how conceptual imitation can be used to reconstruct and compare land cover classes.

2 Semantic interoperability revisited

A shift from convention to conversation in semantic technology mirrors the historic shift in linguistics and philosophy of logic which has led to the development of *speech act theory* (Austin, 1953). In his seminal paper, Grice (1975) argued that the debate between formalists and informalists in philosophical logic may be overcome if formal logics could be made to pay adequate attention to the conditions that govern *conversation*. Conversation depends on perspective. For example, the sentence “he is in the grip of a vice” may refer equally well to one’s bad habit or to the fact that part of one’s body is caught in a tool. The problem is that any standardized, conventional meaning of the word “vice” may be just not what was meant by the speaker, as this meaning shifts with speech situations. A system of inference which is based on conventional or default word meanings is therefore bound to fail in human communication. Insofar as data is a vehicle for human communication, we may be confronted with a similar situation.

2.1 Semantic interoperability as a problem of multiple perspectives

The first example we discuss is taken from an early article on interoperability of geospatial data (Harvey et al, 1999; Riedemann and Kuhn, 1999). Suppose we had to represent a sportsground like that depicted on the aerial photograph in Figure 1 in a cartographic map for the purpose of noise abatement planning. Our goal is to identify objects on the sportsground that emit noise.



Fig. 1 Sportsground Roxel near Münster in an aerial photograph. Taken from (Harvey et al, 1999).

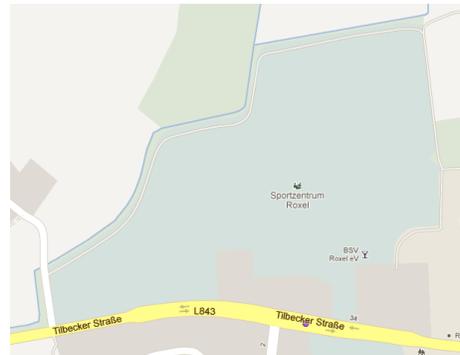


Fig. 2 The sportsground on Google maps.

Which information resource about the sportsground will provide us the required information? Figures 3 and 4 show two very different kinds of maps. In the cadastral map (Figure 3), the most prominent feature, the track and soccer field area, is not shown. Similarly, the tennis courts are left out. This can be explained if we recall that a cadastral map is about land parcels and ownerships, and that the distinction between a soccer field and its surrounding area is not one of ownership. Google maps seems to stick to a similar cadastral world view in Figure 2.



Fig. 3 Cadastral map (ALK) of the Sportsground Roxel (Harvey et al, 1999).

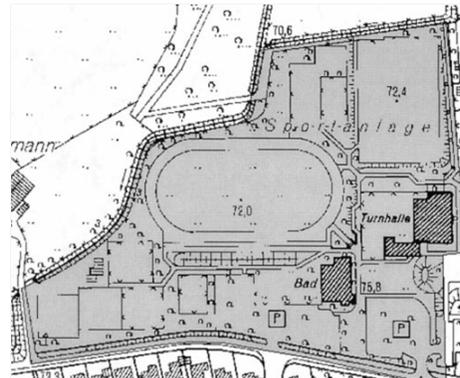


Fig. 4 Topographic map (DGK) of the Sportsground Roxel (Harvey et al, 1999).

In contrast, the sportsground itself is represented in the cadastral map since it can be distinguished from its surroundings precisely based on ownership. The topographic map in Figure 4, on the other hand, shows the track area and the tennis courts but leaves out the soccer field. This can also be explained if we recall that a topographic map is about ground

surface features. There is a distinction in surface texture between the elliptic track area or the tennis courts on one hand and the lawn on the other.

However, since our goal is to identify sources of noise, we are interested in identifying the soccer field. Surface texture does not allow to distinguish it from its embedding lawn. Rather, we need to switch perspective on the sportsground and regard it in terms of *social affordances* (Scarantino, 2003). The latter are conventionally established dispositions to play a game, indicated by linear signs on the lawn⁴. Under the new perspective, some aspects of the surface, such as the signs, become relevant, while others, such as the texture, become irrelevant.

It seems that there is large variability in mapping a single area, to the extent that even individual geographic phenomena, not only their semantic categories, appear and disappear depending on the kind of abstraction one is imposing on an observed environment. Moreover, it is not the environment, but rather the purpose, and with it the kind of observation and abstraction employed, which allows to distinguish the cadastral from the topographic and the game perspective. These perspectives are based on different underlying concepts, namely ownership, surface texture and gaming affordance.

Many geoscientific categories have an intrinsic multi-perspectival nature. They are situated, i.e., their interpretation depends on space and time, and sometimes on a group of researchers (Brodaric and Gahegan, 2007). This frequently causes interoperability problems (Brodaric, 2007). A further example is the assessment of the environmental impact of land use on our climate. *IPCC⁵ land cover classes* serve to quantify land use change based on a transition matrix such as in Figure 5. The matrix is used to classify land area change (with every matrix element standing for a type of transition) in an attempt to estimate the impact of land use on our climate.

	Forest land	Grass-land	Crop-land	Wet-lands	Settle-ments	Other land
Forest land	FF	FG	FC	(FW)	FS	FO
Grass-land	GF	GG	GC	(GW)	GS	OS
Crop-land	CF	CG	CC	(CW)	CS	CO
Wet-lands	(WF)	WG	WC	WW	WS	WO
Settle-ments	(SF)	(SG)	(SC)	(SW)	SS	(SO)
Other land	OF	OG	OC	(OW)	OS	OO

Fig. 5 Land cover categories of the IPCC in a transition matrix used to determine land use change for climate impact assessment. Each element denotes a transition class. Source: Global Forest Observation Initiative (GFOI).

The problem is that the IPCC land cover classes allow for a *large degree of freedom of interpretation* and *do not distinguish incommensurable perspectives*. The continuous transition of surface texture from forest to grassland allows for arbitrary category boundaries. This causes an explosion of more than 800 locally adopted forest definitions (Lund, 2012), each of which has a specific ecological (and political) context. Furthermore, and more importantly, some classes such as cropland and forest draw on incommensurable perspectives on a land surface. From an ontological standpoint, a *crop* is not a kind of plant, *it is a role played by a plant in the course of human cultivation*. Oil palm plantations, for instance, can be considered croplands or forests, depending on whether one takes a vegetation or a cultivation perspective (Lund, 2006). This makes the two categories end up on orthog-

⁴ Similar to traffic locomotion affordances (Scheider and Kuhn, 2010).

⁵ International Panel on Climate Change, <http://www.ipcc.ch/>

onal conceptual dimensions, instead of being mutually exclusive and uni-dimensional, as required by the IPCC (Lund, 2006).

The examples demonstrate that semantic interoperability is to a significant extent a problem of multiple perspectives. This frequently causes fruitless debates about how phenomena should best be represented “in general”, as documented, e.g., by the object-field dualism (Couclelis, 1992). The existence of multiple purpose dependent perspectives on a domain suggests a pragmatic approach to semantics, as proposed, e.g., by Brodaric (2007) and Couclelis (2010).

Multiperspectivity is a basic trait of human culture, cognition and perception. The multiple perspectives underlying language were put sharply into focus by Quine⁶ and Wittgenstein (2003). Analogously, human perception was recognized early by Gestalt psychologists as being *multi-stable*, i.e., allowing to switch between different geometric interpretations of the same scene (Köhler, 1992). This is demonstrated, e.g., by the Necker cube or the face/vase illusion in Figure 6. It seems therefore that we need to pay closer attention to the different pragmatic techniques⁷ that deal with perspective. In how far do current semantic engineering approaches take account of this?

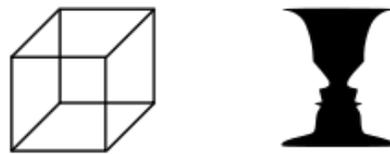


Fig. 6 Necker cube and face/vase illusion can be perceived in terms of two contradicting 3-D interpretations by the subconscious visual routines (Lehar, 2003).

2.2 Paradigms of semantic interoperability

What is it that makes semantic interoperability a challenging problem? In essence, we seem to underestimate multiperspectivity by assuming that it can be solved through standardization or translation.

- The problem cannot be reduced to *labeling*, i.e., to establishing standard terms or term mappings for given concepts. A labeling problem is a *multi-term/single perspective* problem. It implies that the inventory of concepts, i.e., individuals and categories, is fixed, and interoperability problems arise only because we use different names for the same concepts.
- Interoperability is also not a *translation problem*. Translation is a standard approach in the Semantic Web and ontology mapping (compare chapter 2.4 in Stuckenschmidt and van Harmelen (2003)). In the translation approach, we assume that the sets of concepts underlying different datasets may be different but can be mutually translated based on a *shared ontology*. The latter is a theory where each concept is definable so that a new concept has a translation in the ontology (Stuckenschmidt and van Harmelen, 2003). For example, a new concept may be a hyponym of an existing one and thus may be related to it by subsumption⁸. This approach requires, however, that a shared ontology exists.

⁶ Quine based his famous arguments of indeterminacy of theory (Quine, 1951) as well as reference (Quine, 1974) on this.

⁷ Better captured by the German term “Kulturtechniken”.

⁸ Compare Chapter 6.1 in Stuckenschmidt and van Harmelen (2003).

- In addition, interoperability cannot be assured by *standardization*. Standardization of names solves only labeling problems. And standardization of concepts (in terms of a formal theory) requires a common perspective which serves as a denominator for all the different perspectives. This, however, is precisely what can not be expected under multiperspectivity.

Paradigm	Main idea	Strategy	Required means	Basic assumption
HOLISTIC STANDARDIZATION	Term-meaning standard	Heterogeneity resolution	Ability to subscribe to a standard	Term-meaning can be standardized
TOP-LEVEL ONTOLOGY ALIGNMENT	Alignment with core standard	Heterogeneity avoidance	Ability to align with core standards	Core term-meaning can be standardized
PLURALIST PEER-TO-PEER TRANSLATION	Term similarity and translation	Heterogeneity mitigation	Ability to translate between similar terms	Term-meanings are comparable because concepts overlap
RECONSTRUCTION AND CONVERSATION	Term-meaning regeneration	Heterogeneity articulation	Ability to reconstruct concepts and to act on information	Concepts can be reconstructed and term-meanings can be communicated

Table 1 Paradigms of semantic interoperability

Table 1 lists different paradigms that semantic engineering has gone through so far, together with a new one, called reconstruction and conversation. The paradigms are ordered by their tolerance with respect to semantic heterogeneity. While holistic standardization tries to resolve heterogeneity by applying a technical standard (an example would be an ISO⁹ standard), top-level ontology alignment seeks to avoid heterogeneity by sticking only to a top-level standard, as proposed, e.g., in Masolo et al (2003). Pluralist peer-to-peer translation, in contrast, does not require an ontology standard. It acknowledges heterogeneity and at the same time tries to mitigate its negative effects by establishing translations between ontologies based on similarity (Euzenat and Shvaiko, 2007). The ability to translate between ontologies implies at least an overlap in central concepts (Stuckenschmidt and van Harmelen, 2003). Concept similarity, however, is only a necessary and not a sufficient condition of concept identity. What if concepts are similar but at the same time untranslatable, because they correspond to different perspectives on the same matter, as in the case of landcover classes?

Semantic interoperability rather seems to be a *multi-term multi-perspective* problem. This means that perspectives vary in fundamental ways with respect to their ontological commitment which makes them to some degree *untranslatable* (Quine, 2001). That is, each dataset comes with concepts which are not present in another perspective, since they have entirely different origins. For example, think about translating ownership terms into vegetation surface terms on the same sportsground.

For the purpose of comparison, however, it is still possible to expose different concept origins. We suggest therefore that interoperability should be based on *reconstruction* and *conversation*. Reconstruction involves knowing how information was obtained, i.e., which observations were performed and how abstractions were generated. It enables users to understand differences and to regenerate conceptualizations. Conversation validates this knowledge in a *peer-to-peer* fashion, i.e., with respect to a certain data producer. This is in analogy to Grice's conversationally fixed meaning. Following Janowicz and Hitzler (2012), and Janowicz (2012), the strategy therefore should not be to resolve, avoid or mitigate heterogeneity, but to articulate it. This implies that semantic differences need to be understood

⁹ <http://www.iso.org>

and be focused on, not leveled. For this purpose, term-meanings need to be (re)generated and communicated, not standardized, aligned or translated, see Figure 7.

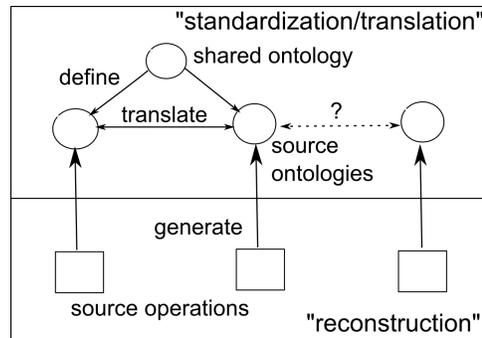


Fig. 7 Reconstruction vs. standardization/translation paradigm. In the reconstruction paradigm, the impossibility of translations can be assessed based on reconstructing concepts with source operations.

The principal problem with the existing paradigms is their weakness in *distinguishing perspectives*, i.e., in detecting basic conceptual differences, and thus translatability in the first place. On what grounds do we assume that there is a shared ontology? And vice versa, how could we possibly know that no such ontology exists? It seems that the current paradigms do not provide good answers to these two questions. Also, their focus on ontology alignment and translation risks technically enforcing some kind of unification on superficial grounds.

One should note that the paradigms in Table 1 are not meant to substitute each other. Conceptual standards and translations are neither useless nor wrong, merely insufficient to detect multiple perspectives. Instead of standardizing perspectives, we suggest to *standardize procedures to arrive at different perspectives*.

3 Purpose, design and sharing of information in a machine context

In order to address perspectivity by pragmatic means, we suggest to take a pragmatic view on information as such. Janich (2006) has argued that one of the most far-reaching errors of modern science is the fairy tale of “information as a natural object” which is supposed to exist independently of cultural techniques¹⁰. The problem with this idea is that it tends to blur exactly those origins of information that we need to expose. Namely, that information is *designed for a purpose*, and that the *sharing of information* is simply a function of this design and the capabilities involved.

Design makes information a product of (communication) culture, not nature (Couclelis, 2010). Therefore, the distinction between pragmatics and semantics of information, as well as between situated and non-situated concepts (Brodaric, 2007), seems spurious. The sharing of digital information is a matter of its design in a human-machine-human conversation.

¹⁰ The history of this fundamental misunderstanding can be traced back to Morris’ naturalized semiotic process and Shannon and Weaver’s mechanistic information theory, and can be currently studied in terms of modern nonsense about “information” allegedly being “transferred and understood by machines, computers, and DNA molecules” (Janich, 2006).

3.1 Information in human-machine-human conversations

The problem is not that machines cannot communicate, but that humans misunderstand each other when communicating via machines (Scheider, 2012). At first glance, digital machines have greatly increased the efficiency, speed and range of human interaction. They have successfully substituted humans in searching, filtering, transforming and validating information, to the extent that their role appears to be negligible. The Semantic Web can be seen as an effort to lift this to the level of meanings.

However, the latter view neglects that *information* is basically a *derivative of human speech acts*, and that the success of any information-processing machine therefore needs to be measured with respect to human capacities of speech (compare Janich (2006): Chapter 5). Information appears if people inform other people about something. The acts involved turn data into information. This remains true even if human speech is technically encoded, transmitted and extended by computers and technical sensors (Figure 8). The limits of modern information technology seem therefore less defined by the technical efficiency of symbol processing, but by the *substitution of human speech acts by equivalent technical capacities*. Regarding the latter challenge, modern information technology turns out to be surprisingly weak.

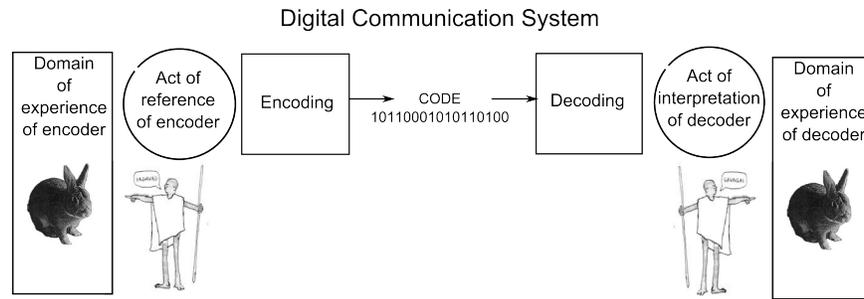


Fig. 8 The problem of reference in human-machine-human conversation. Acts of reference are encoded in a digital form, which makes it hard for a decoder to interpret the code in terms of his or her own domain of experience.

For example, there is a fundamental problem regarding the encoding of reference (Scheider, 2012), which is documented, e.g., by the debate about the meaning of URI¹¹ references in the Semantic Web (Booth, 2006; Halpin, 2013). What does a URI refer to, over and above the Internet address to which it can be resolved? Digital information is basically a product of encoding *acts of reference* to a domain experienced by an encoder (Figure 8). In terms of visual indices to jointly perceived scenes, acts of reference play a central role in learning and sharing a language (Tomasello, 1999; Quine, 1974). However, a language standard for encoding resource reference, such as RDF¹², does not ensure that the decoder understands and shares referents.

The role of information technology is to substitute speech acts by equivalent capacities in order to increase the range and efficiency of human speech. The challenge is to *talk to each other via computers*, even if communicators are not present in the same scene (Janich, 2006). What is needed for this purpose?

We know that the efficiency and range of human speech can be effectively extended beyond natural boundaries using signs. This is demonstrated, for example, by a traffic light which prompts a driver to stop and move even in the absence of a traffic policeman. Similarly, the turn signal on a vehicle indicates a turn without the driver being visible (Kamlah

¹¹ Uniform resource identifier.

¹² Resource Description Framework, <http://www.w3.org/RDF/>

and Lorenzen, 1996). However, in order to understand such signs, knowledge about the underlying human capacities is essential, as they are the methodological ground for understanding. Correspondingly, traffic lights and turn signals are meaningless in the absence of knowledge about prompted stopping and turning actions.

In a similar sense, data sets can be regarded as manifest speech acts (Scheider, 2012). Speech acts *prompt an intended kind of reaction* from a listener. It is this reaction that the listener needs to understand. For example, in Figure 8, the encoder draws the attention of the decoder to a rabbit in his or her domain of experience. More precisely, the encoder prompts a focusing of attention with a symbol like “Gavagai” (reusing the famous example from Quine (2001)). Furthermore, the encoder makes a statement about this rabbit, e.g., that it was observed at a location at a time. If this statement is digitally encoded as a sentence, then we obtain *data*. However, only if the drawing of attention was successful is reference to the rabbit shared among agents, and the data set then becomes meaningful for the decoder. Otherwise, the decoder remains unsure about what was meant. Thus, we need to ask how we can make users *learn the reactions intended by the providers of a data set*.

3.2 The law of uphill analysis and downhill invention

The challenge we face reflects a more general one in AI and robotics, first formulated by Valentino Braitenberg, a cybernetician, in his famous book about “vehicles” (Braitenberg, 1986). Braitenberg presented a neat collection of very simple construction plans for technical devices made from analog circuitry (“vehicles”) that are able to generate a complex set of human-like behaviour, such as love, fear, thought and optimism. For example, by a simple analog circuit and a light intensity sensor, a vehicle can be made to move in such a way as to avoid light sources.

Braitenberg’s thought experiment shows that it is easy to create an interesting repertoire of complex behaviour based only on simple processes with feedback¹³. This is “downhill invention”. However, from the outside, without opening the black box, an observer of the vehicles’ behaviour has almost no chance of correctly guessing how they were built. The content of the black box may look as if it were beyond the grasp of human reason, just as love and thought continuously appear to human analysts. This is “uphill analysis”. There are several reasons why analysis is harder than synthesis. First, many different potential mechanisms could realize the same behaviour. And second, induction is harder than deduction because one lacks the constructive basis, i.e., “one has to search for the way, whereas in deduction one follows a straightforward path” (Braitenberg, 1986, p.20).

Braitenberg suggests that *analysts* of a mechanism therefore tend to overestimate its complexity because they only have access to complex behaviour, not to the constructive basis. They first *describe* complex behaviour (which is complex), and then try to guess how it may be generated. However, once analysts know which basis is the correct one, things start to become easy, and they become *synthesists*.

This insight applies also to semantic interoperability. If we take Braitenberg’s law seriously, then *synthetic learning needs to take on an indispensable role in the learning of meanings*¹⁴. This explains why it is hard to handle semantic interoperability under the familiar paradigms. Current approaches to data integration are either based on “black box” descriptions of concepts (*ontologies*), or on *machine learning* of semantic concepts. The latter is an automatization of guessing a function based purely on observations of external behaviour (Hastie et al, 2001). The former are a specification of purported behavioural and conceptual constraints in a top-down fashion (Guarino, 1998). However, both approaches

¹³ For a proposal how central spatial concepts can be based on Braitenberg vehicles, see Both et al (2013).

¹⁴ Compare also the arguments given in Chapter 3 of Scheider (2012).

are basically analytic, not synthetic. That is, they rely on descriptions of concept behaviour and do not involve any information about how a concept was generated and on which constructive basis¹⁵. This may be one reason why semantic interoperability is hard with analysis tools, while human language learning is almost effortless for children.

What does it take, then, to interoperate with information? First, “interoperation” should be taken literally, i.e., it consists of speech acts that intertwine and prompt reactions, even if part of these acts are performed by machines. And second, the prompted kinds of reactions include those that need to be learned in a synthetic fashion. Synthetic learning is what makes information a *designed* product. And this is the tough part of interoperability, because it requires the correct constructive basis.

3.3 Synthetic learning requires imitation

How do we learn in a synthetic rather than analytic fashion?

If we follow somebody in synthesizing something, then we *imitate* this person based on our own capacities. For example, if a child is taught how to build a castle, it learns by imitating the construction (even if the result may look different, and even if its competences are slightly different). In the same sense, *sharing meaning requires imitation*, and as such seems to be the fundamental mechanism for knowledge transmission, as argued by Tomasello (1999). For example, the robot on the hand left side of Figure 9 learns the concept of “holding something in front” by recomputing the meaning of this notion in terms of its own body frame of reference. That is, it does two things simultaneously: it observes a speech act behaviour of its opposite and then recomputes a concept based on its own sensori-motor system. Note that the latter system largely constrains the constructive basis.

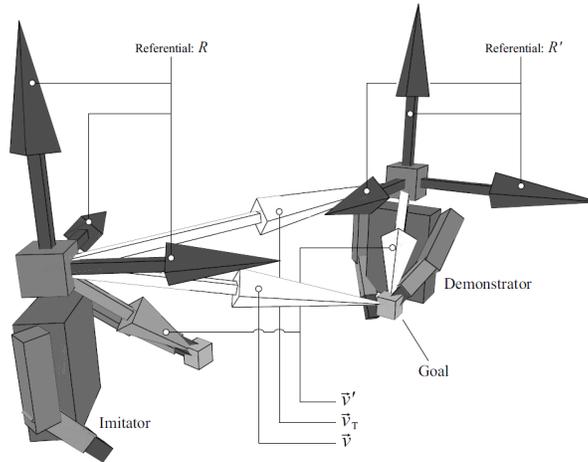


Fig. 9 Robot imitation learning of “holding something in front” based on recomputing it with respect to body reference frames (Sausser and Billard, 2005). The imitator robot needs to translate vector v' relative to the demonstrator referential frame into a corresponding vector of its own referential frame based on vectors v, v' and v_T .

The system basically consists in a vector calculus grounded in body postures. The grounding leaves no question which constructive basis to choose, and thus learning is synthetic.

¹⁵ Machine learning is analytic in the sense that it prescribes a constructive basis (e.g., in terms of a vector calculus in support vector machines (Hastie et al, 2001)) or automatically selects it based on observed behaviour (as in Bayesian model selection).

3.4 *Synthetic learning needs to be grounded*

The insight that synthetic learning and imitation needs to proceed from a *common embodied ground* is central in embodied AI and robotics (Steels, 2008). We argue that this is also the case for synthetic learning in data semantics, because the grounding disambiguates the constructive basis for a concept (Scheider, 2012).

Luc Steels has demonstrated that robots can construct simple perceptual distinctions and share them via *discrimination games* (Steels, 1997). This is an example of synthetic learning of concepts. Based on this, robots can exploit social mechanisms of language learning (Steels, 2003) which allow them to establish term-meanings and their own languages via so-called *language games* (Steels, 2002).

Grounded symbol systems exist also in human communities and play an immensely useful role in technology, namely in the form of *reference systems*. There are reference systems for measurements (measurement scales), for locations (geodetic coordinate systems) as well as for time (calendars), allowing us to refer to phenomena in a very precise manner. It is important to understand, however, that not every formal symbol system is grounded. In particular, grounding is not implied by formal ontologies (Scheider, 2012). For many ontological concepts, reference systems and groundings are currently lacking. Kuhn (2003) therefore proposed to build new (*semantic*) *reference systems* for arbitrary kinds of concepts, and Scheider (2012) discussed how grounding of such systems might be done.

As it turns out, there is a large variety of constructive calculi as well as primitive perceptual operations that may be used for this purpose. Beule and Steels (2005), for instance, have proposed Fluid Construction Grammar as a formalism. The latter follows the idea of Johnson Laird's procedural semantics (Johnson-Laird, 1997), where the meaning of a phrase is taken to be the execution of a (perceptual) program. For humans, a rich set of perceptual operations is available for any reproducible phenomenon that can be perceived, including surfaces, actions and action potentials (Scheider, 2012, Chapter 5). Measurement procedures and technical sensors simply increase the range of human perception. From the viewpoint of logic, constructive calculi range from sets of constructive rules (as in intuitionist logic) (Lorenzen, 1955), over recursive functions in higher-order logic (HOL) (as in functional programming) (Nipkow et al, 2002), to the formation of logical definitions in ordinary first-order logic (FOL) (Scheider, 2012, Chapter 4). They may, in particular, involve geometric construction, as suggested by Gärdenfors' *conceptual spaces*, where concepts are constructed as convex regions (Gärdenfors, 2000). A *grounding level* therefore is not a static thing, and it is essential that it has been established, either by convention and standardization or by practice, in order to start synthetic learning.

4 Interoperability through conceptual imitation

In this section, we discuss tools that support conceptual imitation. We propose to use Gärdenfors' conceptual spaces (Gärdenfors, 2000) as a convenient constructive basis for synthetic learning. We regard a conceptual space as a *grounded multidimensional space* equipped with vector space operators, following Adams and Raubal (2009), where the grounding can be secured either in terms of conventionalized human perception or in terms of measurement. Concepts are generated in this space by all operations that generate convex regions. Conceptual imitation requires to establish conceptual spaces and to describe and reconstruct concepts in terms of them. Reconstructed concepts can then be compared, and it can be assessed whether they are interoperable or not. We demonstrate this on the basis of land surface category examples discussed in Section 2.

4.1 Conceptual imitation tools and conceptual spaces

Which tools can be used to support agents in synthetic learning? We focus on the problem of describing concepts in terms of *predicates*, i.e., in terms of concept symbols (F) with slots for instances (a) which can be used to state that the former symbols apply to the latter ($F(a)$). Firstly, we need tools to support the *reconstruction of predicates in some grounded language*. For this purpose, we need some reference theory (Scheider, 2012). This is a formal system of symbols that comes with a fixed operational interpretation of the following kind (compare also the left hand side of Figure 10 and Table 2):

1. A set of *grounded instances* (the domain). These may consist of *foci of attention* of an observer (Scheider, 2012), or of *technical foci* of sensors according to their discrimination unit (Frank, 2009)¹⁶. They may also consist of conventionally established *objects*. Furthermore, predicates of interest (see below) may become instances at a higher level of abstraction, standing for classes of instances on the lower level (Scheider, 2012).
2. A set of *grounded primitive predicates*, i.e., symbols with slots for instances whose meaning is fixed by standard procedures, as described below.
3. For each primitive predicate, a *known instantiation procedure*, which may be embodied and therefore not part of the formalism, and which allows to *decide* whether the predicate applies to grounded instances (i.e., it allows to instantiate the predicate) or not. Examples may be measurable sensor properties such as altitude and temperature, but may also involve perceptual distinctions of humans such as the discrimination of objects and their properties.
4. A *constructive calculus*¹⁷ on predicates, i.e., a set of syntactic rules which allow to generate (define) new predicates in terms of primitive predicates. Examples are the syntactic rules of a logic which may be used to generate definitions of new predicates, or geometrical operations which may be used to construct new regions in a conceptual space. A particular construction is described by a *constructive procedure* or a *constructive definition* (definiens). The constructive calculus also involves semantic rules which prescribe how a defined predicate must be interpreted into instances.
5. A set of *predicates of interest* (definiendum) generated by constructive procedures.
6. An *instantiation procedure* for all predicates of interest, i.e., a sequence of inference steps that allows to decide whether these predicates apply to grounded instances, or not. The instantiation procedure is not (always) identical with the constructive procedure. For example, the syntactic pattern of a FOL definition does not imply a decision procedure. Only *decidable languages* have a decision procedure for every possible predicate of interest. If definitions are recursive, then inference procedures need to include a fixpoint operator. Note also that the instantiation procedure for predicates of interest may be distinct from that for primitives. Once the primitive predicates are determined, for example based on measurements, instantiation may proceed purely syntactically for defined predicates.
7. Instantiation procedures give rise to a set of *ground sentences*¹⁸ about instances. They allow to compute if a predicate (whether primitive or not) applies to instances, or not. Ground sentences are considered *data*.

Table 2 lists these imitation tools with some examples, drawn from FOL and conceptual spaces. What differentiates the schema from FOL is the distinction between instantiation procedures based on syntactical inference (3) and instantiation procedures based on grounding (6). This distinction is not made in ordinary logic. Furthermore, it turns out that conceptual spaces are simple and straightforward examples of conceptual imitation

¹⁶ The “instantaneous field of view” (IFoV) of a satellite is an example for the latter.

¹⁷ According to Lorenzen, a calculus is a set of rules used to generate “figures from other figures” (Lorenzen, 1955).

¹⁸ These are sentences without variables.

tools. We illustrate this based on the categories *mountain* or *forest*, which can be defined in terms of convex regions in a space with dimensions based on remote sensors (compare also Section 4.3 and Adams and Janowicz (2011)): Grounded instances are spatio-temporal granules (*instantaneous fields of view* (IFoV)), which correspond to pixels in a satellite image. Primitive predicates include ground altitudes and tree heights, which correspond to particular values on some measurement scale. Measurement scales correspond to single dimensions of a conceptual space, and combinations of values of different dimensions correspond to points in this space. The constructive calculus is a vector calculus, which allows to form algebraic expressions over points. Constructive procedures are algorithms for generating convex regions in this space, such as polytopes. Constructive definitions are definitions of convex regions, e.g., lists of points for polytopes. Predicates of interest are particular convex regions which account for a concept, e.g., mountain or forest polytopes. And instantiation procedures for predicates of interest are point-in-polygon tests, i.e., algorithms which determine whether a certain point lies within some region that defines a concept. For a more precise formulation of conceptual space operators, see Adams and Janowicz (2011); Adams and Raubal (2009), and for further illustrations regarding this example, see Section 4.3.

All of the tools discussed above need to be learned in order to use them in conceptual imitation. Some tools, such as the constructive definition of a polytope, may be easily communicated, while others, such as the measurement procedure underlying tree height, may be more difficult to acquire.

Note that we assume the operations that allow *instantiating* primitive predicates to grounded instances to exist outside the reference theory, and thus outside of any computer. In the example above, these are measurement procedures of a satellite, but they could also be human observations. The requirement is not that these operations are *computable*, but that they are *shared and repeatable* in an inter-subjective way. Note that the acknowledgment of decision procedures of different origin (embodied/analog/non-deterministic as well as algorithmic/digital) distinguishes our approach from many others.

Imitation tool	Examples	Conceptual space example
Grounded instances	Foci of attention / sensor ranges / perceptual phenomena	IFoV of some satellite
Grounded primitive predicate	Measurable or perceptual qualities	Ground altitude / tree height
Constructive calculus	Logical syntax/algebra	Vector calculus
Constructive procedure	Formation algorithm	Convex region generator
Constructive definition	Definiens	Convex region
Predicate of interest	Definiendum	Mountain/Forest
Instantiation procedure	Inference (predicate satisfaction) algorithm	Point-in-polygon test

Table 2 Synthetic conceptual imitation tools

We also assume that predicates of interest come with a decision procedure. In our example, these are point-in-polygon tests for polytopes. In a more general setting, which may be based on the flexible FOL syntax, the ordinary FOL syntactic calculus needs to be either restricted, such as in the Semantic Web, or handled with some care to ensure that there is a decision procedure for each constructed predicate¹⁹.

Besides such synthetic tools, one may also use analysis tools, e.g., machine learning (Figure 10). However, the latter come with a bias in their construction calculus (Mitchell,

¹⁹ Note that we do not require a decision procedure for the entire constructive calculus, only for the predicates of interest. This allows to use unrestricted FOL or HOL as the most flexible syntactic standard, but comes at the price of caring about the computation of decisions on a case-to-case basis.

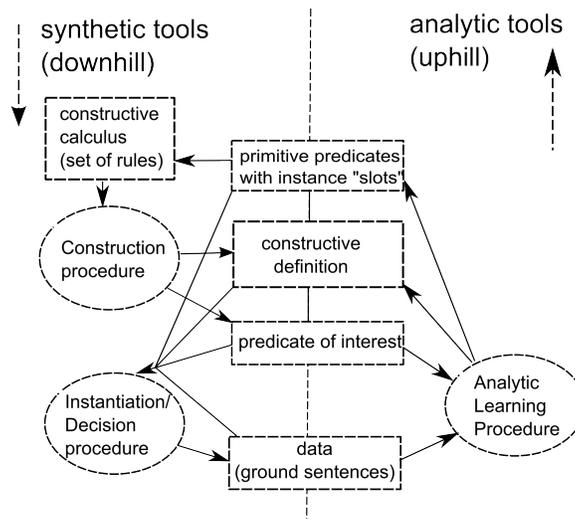


Fig. 10 Kinds of conceptual imitation tools and their role in learning. Synthetic tools start with a constructive basis, which consists of primitives and a calculus as well as an instantiation procedure for generating data, while analytic tools try to guess the constructive basis starting from the analysis of data.

1980) and they are analytic, i.e., they lack information about the constructive basis. This leaves learners with the problem of choosing a constructive basis on their own, or of delegating the problem to some standard algorithm.

4.2 Measuring perspectival correspondence

Once a foreign concept has been learned synthetically, i.e., is reconstructed in terms of conceptual imitation, it becomes possible to measure the extent to which it corresponds to a concept in question or rather belongs to a different perspective. In principle, one can distinguish several levels of perspectival correspondence of concepts, depending on whether they share primitive predicates (with underlying operations) and correspond in terms of construction procedures (and constructive definitions), and whether they correspond to each other synthetically or analytically:

1. *Cocomputable*: The concepts are equivalent in the sense that they are constructed by the same procedure in terms of a calculus of deterministic primitive predicates.
2. *Codescendant*: The concepts are similar in the sense that they are constructed by the same non-deterministic predicates. This corresponds to drawing samples from the same random process.
3. *Comparable*: The concepts draw on the same primitive predicates, but may correspond to different procedures. This allows *translating* them into each other.
4. *PartiallyComparable*: Primitive predicates overlap. This allows projecting concepts onto common dimensions, but not translating them, since there are missing parts.
5. *Incomparable*: The concepts draw on different primitive predicates.
6. *Coincident*: Concepts apply to the same grounded instances.
7. *Overlapping*: Concepts overlap with respect to grounded instances.

Inside a conceptual space, concept similarity can be measured based on spatial proximity of concepts (Gärdenfors, 2000). This is only possible if concepts are at least *comparable* or *partially comparable*; only in this case can we project concepts onto common dimensions. Our classification scheme for correspondence, however, covers also the case of incomparable concepts, where concepts are orthogonal to each other. In this case, similarity is not

meaningfully captured by distance. However, there may still be equivalences on the level of types of phenomena that different conceptual spaces are covering, such as altitude, speed or temperature²⁰.

Note that correspondences are not mutually exclusive. For example, some concepts may be coincident and incomparable, such as the concepts of heat and red, and may then be called *proxies*. Also, every codescendant is also comparable. Some relations may be refined gradually, such as Comparable, or Overlapping (the latter in terms of precision and recall). We discuss examples of these concept relations in the following.

4.3 Conceptual imitation of land cover categories

We illustrate the use of conceptual imitations by reconstructing land cover categories, as illustrated in Baglatzi and Kuhn (2013), and Adams and Janowicz (2011).

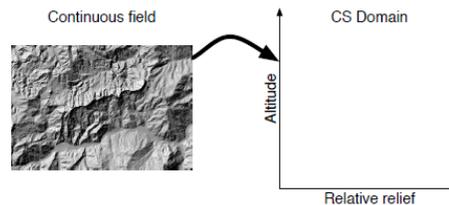


Fig. 11 A conceptual space of continuous altitude measures (Adams and Janowicz, 2011). A continuous field of remote sensor pixels is mapped into a conceptual space of altitude and relief measures.

Suppose we have a vector space of *spatially continuous relief properties*, such as proposed by Adams and Janowicz (2011). Dimensions of this space may be altitude and relief measures as well as location, as depicted in Figure 11. Each vector of this space represents a particular measurement in each of the metric dimensions at some measured focus in space (and time).

An elevation that is called a “mountain” in Scotland, such as Ben Lomond (974 m), would hardly be called the same in Asia. A relief space can be used for teaching and distinguishing local mountain concepts, for example *English mountains* and *Asian mountains*, and thus different perspectives on the concept mountain. The relief space comes with a set of primitive predicates, namely the set of values of relief measurements. There is an unambiguous decision procedure for each value in terms of the well known relief measurement procedures. The set of grounded instances is simply the finite sample of measurements taken. These instances can be projected into a set of vectors in Figure 12. Note that measurement instances cannot be identified with points in this conceptual space, because the projection need not be injective, so that different measurements map into the same point. The vector calculus can be used to *define concepts as convex regions* in this space. The latter function as predicates of interest. Convex regions for the concept “mountain” can be efficiently computed based on convex polytopes (see Figure 13). Point-in-polygon tests decide whether some instance lies in a polytop, and thus whether it falls under a specific mountain concept, or not.

In this way, one can find out to what extent English and Asian mountain concepts are different (compare Figure 13); concepts are *non-overlapping*, yet defined in terms of the same grounding level, and thus *comparable and translatable*. In a similar way, one could

²⁰ This aspect of similarity is based on experiential equivalence and is not discussed in this article.

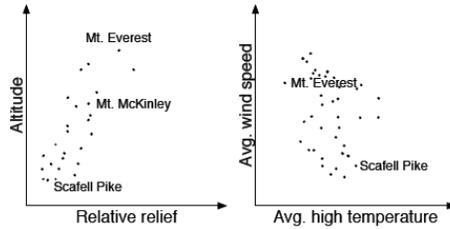


Fig. 12 Grounded mountain instances and their location in two conceptual spaces (Adams and Janowicz, 2011).

define the meaning of hydrological object classes such as pond, lagoon and lake as regions in a space of size, naturalness, and marshiness, following the ideas of Mark (1993).

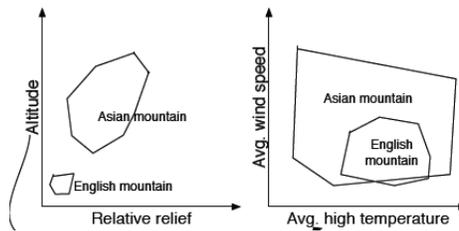


Fig. 13 Local mountain concepts synthetically learned in terms of convex regions (Adams and Janowicz, 2011).

In analogical fashion, one may learn the IPCC land cover classes *cropland* and *forest*, as introduced in Figure 5. The multitude of local forest definitions, as listed in Lund (2012), may be mapped into a conceptual vector space, provided observation procedures are established for the so called “Marrakesh” variables as dimensions. The latter include tree height, crown cover, and minimum area. They also imply an operational definition of the concept “tree”, which needs to be established first, as well as expectation values of measured values for the future (Lund, 2006). In such a conceptual space, one can discover that several different forest concepts overlap, see Figure 14. According to our scheme, they are therefore *comparable*.

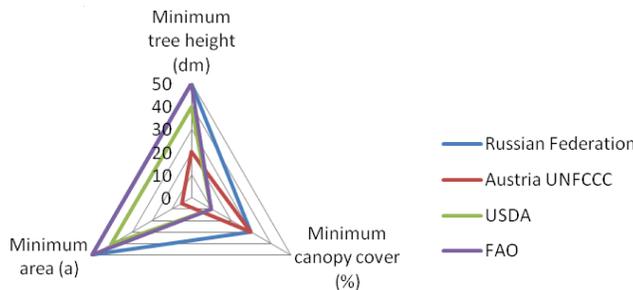


Fig. 14 Local forest definitions projected as triangles into a 3 dimensional vector space of “Marrakesh” variables. Thresholds were taken from Lund (2012).

With the construction of the concept *cropland*, the situation is different. One first needs to ground a separate primitive predicate, e.g. *Cultivated*, that captures whether the land cover is used for growing crops (Lund, 2006), i.e., whether it is an object of human cul-

tivation. This predicate is conceptually orthogonal to the Marrakesh variables, which are not based on perceiving agricultural actions. Thus forest and *Cultivated* are not mutually exclusive²¹. However, since we now know that *Cultivated* is defined on an entirely different operational basis than forest, we can conclude that both classes are based on incommensurable perspectives, they are *incomparable*. It may or may not be the case that forests tend to be non-cultivated. If not, the two concepts are also *overlapping*.

5 Conclusion

In this article, we have suggested that interoperability should be defined as *correspondence* of conceptual perspectives of communicating agents. Correspondence can be measured in terms of the constructive bases of the concepts in question. A data set of a provider is interoperable with respect to a user if the underlying user and provider perspectives correspond to each other in this sense. The pragmatic paradigm of conversation and reconstruction suggests that interoperability does not require to resolve, avoid or mitigate heterogeneity, but to articulate it. For this purpose, concepts need to be imitated, not standardized, aligned or translated. Correspondingly, digital information should be viewed as an abstraction of human speech acts that prompt imitations. Such imitations include synthetic learning. According to Braitenberg's law, synthetic learning involves knowledge about the constructive basis used to generate concepts, in contrast to analysis which is based solely on imitating concept behavior. We suggested therefore that imitation of semantic concepts should be grounded, and we proposed to do so based on conceptual spaces. This approach allows users to compare their intended concept with respect to a data provider concept, and thus to justifiably assess *the possibility of translations*. We demonstrated the approach by conceptual imitation and assessment of correspondence of land cover classes.

Since this article is an outline of a pragmatic paradigm of interoperability, several questions remain open to future research. First, which grounding levels including calculi and primitive operations other than conceptual spaces are useful for conceptual imitation? Second, how can conceptual imitation be technically realized as computer dialogues, such that users and providers are supported in a their dialogue with the machine. And third, how can correspondence of grounded user and provider concepts be computed?

Acknowledgements

A draft of the ideas in this article was presented at the EarthScienceOntolog session-3 at 10/11/2012²². Research was funded by the International Research Training Group on Semantic Integration of Geospatial Information (DFG GRK 1498), and by the research fellowship grant DFG SCHE 1796/1-1. We thank Helen Couclelis, Benjamin Adams, Krzysztof Janowicz and MUSIL²³ for discussions that helped shape this article.

²¹ This can be gleaned from the fact that Lund (2006) proposes a decision tree which enforces mutual exclusiveness of cropland and forest by defining cropland based on cultivation as well as the logical complement of forest.

²² <http://ontolog.cim3.net/cgi-bin/wiki.pl?EarthScienceOntolog>

²³ musil.uni-muenster.de

References

- Adams B, Janowicz K (2011) Constructing geo-ontologies by reification of observation data. In: Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, ACM, New York, NY, USA, GIS '11, pp 309–318
- Adams B, Raubal M (2009) A metric conceptual space algebra. In: Hornsby K, Claramunt C, Denis M, Ligozat G (eds) Spatial Information Theory, 9th International Conference, COSIT 2009, Aber Wrac'h, France, September 21-25, 2009 Proceedings, Springer, Berlin, pp 51–68
- Austin J (1953) How to Talk. Some Simple Ways. In: Proceedings of the Aristotelian Society, vol 53, pp 227–246
- Baglatzi A, Kuhn W (2013) On the formulation of conceptual spaces for land cover classification systems. In: Geographic Information Science at the Heart of Europe, pp 173–188
- Beule JD, Steels L (2005) Hierarchy in fluid construction grammars. In: KI, pp 1–15
- Booth D (2006) URIs and the Myth of Resource Identity. URL <http://www.ibiblio.org/hhalpin/irw2006/dbooth.pdf>
- Both A, Kuhn W, Duckham M (2013) Spatiotemporal Braitenberg vehicles. In: 21st SIGSPATIAL International Conference on Advances in Geographic Information Systems, SIGSPATIAL 2013, Orlando, FL, USA, November 5-8, 2013, ACM, pp 74–83
- Braitenberg V (1986) Vehicles. Experiments in Synthetic Psychology. MIT Press, Cambridge, MA
- Brodaric B (2007) Geo-pragmatics for the geospatial semantic web. Transactions in GIS 11(3):453–477
- Brodaric B, Gahegan M (2007) Experiments to examine the situated nature of geoscientific concepts. Spatial Cognition & Computation 7(1):61–95
- Couclelis H (1992) People manipulate objects (but cultivate fields): Beyond the raster-vector debate in gis. In: Frank A, Campari I, Formentini U (eds) Spatio-Temporal Reasoning, Springer, pp 65–77
- Couclelis H (2010) Ontologies of geographic information. International Journal of Geographical Information Science 24(12):1785–1809
- Euzenat J, Shvaiko P (2007) Ontology Matching. Springer-Verlag New York, Inc., Secaucus, NJ, USA
- Frank A (2009) Scale is introduced in spatial datasets by observation processes. In: Spatial Data Quality From Process to Decision (6th ISSDQ 2009), CRC Press, pp 17–29
- Gärdenfors P (2000) Conceptual spaces - the geometry of thought. MIT Press
- Gray J, Liu D, Nieto-Santisteban M, Szalay A, DeWitt D, Heber G (2005) Scientific data management in the coming decade. CoRR abs/cs/0502008
- Grice HP (1975) Logic and conversation. In: Cole P, Morgan JL (eds) Syntax and Semantics: Vol. 3: Speech Acts, Academic Press, San Diego, CA, pp 41–58
- Guarino N (1998) Formal ontology and information systems. In: Guarino N (ed) Formal ontology in information systems: Proceedings of the first international conference (FOIS '98), June 6 - 8, Trento, Italy, Frontiers in artificial intelligence and applications, vol 46, IOS-Press [u.a.], Amsterdam
- Halpin H (2013) Social Semantics - The Search for Meaning on the Web, Semantic Web And Beyond, vol 13. Springer
- Harvey F, Kuhn W, Pundt H, Bishr Y, Riedemann C (1999) Semantic interoperability: A central issue for sharing geographic information. The Annals of Regional Science 33:213–232
- Hastie T, Tibshirani R, Friedman J (2001) The elements of statistical learning: data mining, inference, and prediction. Springer-Verlag, New York
- Hyvönen E (2012) Publishing and Using Cultural Heritage Linked Data on the Semantic Web. Synthesis Lectures on the Semantic Web, Morgan & Claypool Publishers

- Janich P (2006) Was ist Information? Kritik einer Legende. Suhrkamp, Frankfurt a. M.
- Janowicz K (2012) Observation-driven geo-ontology engineering. *T GIS* 16(3):351–374
- Janowicz K, Hitzler P (2012) The digital earth as knowledge engine. *Semantic Web Journal* 3(3):213–221
- Janowicz K, Scheider S, Pehle T, Hart G (2012) Geospatial semantics and linked spatiotemporal data - past, present, and future. *Semantic Web* 3(4):321–332
- Johnson-Laird P (1997) *Procedural Semantics*. *Cognition* 5:189–214
- Kamlah W, Lorenzen P (1996) *Logische Propädeutik. Vorschule des vernünftigen Redens*, 3rd edn. J.B. Metzler, Stuttgart, Weimar
- Köhler W (1992) *Gestalt Psychology. An introduction to new concepts in modern psychology*. Liveright, New York
- Kuhn W (2003) Semantic reference systems. *Int J Geogr Inf Science* 17 (5):405–409
- Lehar S (2003) *The world in your head. A Gestalt view of the mechanism of conscious experience*. Lawrence Erlbaum Associates, Mahwah, London
- Lorenzen P (1955) *Einführung in die operative Logik und Mathematik*. Springer Verlag, Berlin
- Lund H (2012) Definitions of forest, deforestation, reforestation and afforestation. Tech. rep., Forest Information Services, Gainesville, VA, URL <http://home.comcast.net/%7Egyde/DEFpaper.htm>
- Lund HG (2006) Guide for classifying lands for greenhouse gas inventories. *Journal of Forestry* 104(4):211–216
- Madin JS, Bowers S, Schildhauer M, Krivov S, Pennington D, Villa F (2007) An ontology for describing and synthesizing ecological observation data. *Ecological Informatics* 2(3):279–296
- Mark DM (1993) Toward a theoretical framework for geographic entity types. In: Frank AU, Campari I (eds) *Spatial Information Theory A Theoretical Basis for GIS*, Lecture Notes in Computer Science, vol 716, Springer Berlin Heidelberg, pp 270–283
- Masolo C, Borgo S, Gangemi A, Guarino N, Oltramari A (2003) *Wonderweb deliverable d18: Ontology library*. Trento, Italy
- Mitchell T (1980) *The need for biases in learning generalizations*. Cbm-tr 5-110, Rutgers University
- Nipkow T, Wenzel M, Paulson LC (2002) *Isabelle/HOL: a proof assistant for higher-order logic*. Springer-Verlag, Berlin, Heidelberg
- Quine W (1951) Two dogmas of empiricism. *The Philosophical Review* 60:20–43
- Quine W (1974) *The roots of reference*. Open Court Publishing, La Salle, Illinois
- Quine W (2001) *Word and object*, 24th edn. MIT Press, Cambridge, Mass.
- Riedemann C, Kuhn W (1999) What are sports grounds? or: Why semantics requires interoperability. In: Vckovski A, Brassel KE, Schek HJ (eds) *Interoperating Geographic Information Systems*, Lecture Notes in Computer Science, vol 1580, Springer Berlin Heidelberg, pp 217–229
- Sausser EL, Billard AG (2005) View sensitive cells as a neural basis for the representation of others in a self-centered frame of reference. In: 3rd Int. Symp. Imitation in Animals & Artifacts, pp 119–127
- Scarantino A (2003) Affordances explained. *Philosophy of Science* 70:949–961
- Scheider S (2012) Grounding geographic information in perceptual operations, *Frontiers in Artificial Intelligence and Applications*, vol 244. IOS Press, Amsterdam
- Scheider S, Kuhn W (2010) Affordance-based categorization of road network data using a grounded theory of channel networks. *Int J Geogr Inf Sci* 24(8):1249 – 1267
- Shirky C (2009) *Ontology is overrated*. URL http://www.shirky.com/writings/ontology_overrated.html
- Steels L (1997) Constructing and sharing perceptual distinctions. In: Someren M, Widmer G (eds) *Machine Learning: ECML-97*, Lecture Notes in Computer Science, vol 1224, Springer Berlin Heidelberg, pp 4–13

- Steels L (2002) Grounding symbols through evolutionary language games. In: Cangelosi A, Parisi D (eds) *Simulating the evolution of language*, Springer-Verlag New York, Inc., New York, NY, USA, pp 211–226
- Steels L (2003) Social language learning. In: Tokoro M, Steels L (eds) *The Future of Learning*, The MIT Press, Amsterdam, pp 133–162
- Steels L (2008) The symbol grounding problem has been solved. So what's next? In: de Vega M (ed) *Symbols and Embodiment: Debates on Meaning and Cognition*, Oxford University Press, Oxford, chap 12
- Stuckenschmidt H, van Harmelen F (2003) *Information Sharing on the Semantic Web*. Springer-Verlag, Heidelberg
- Tomasello M (1999) *The cultural origins of human cognition*. Harvard University Press, Cambridge, MA
- Wittgenstein L (2003) *Philosophische Untersuchungen*. Suhrkamp, Frankfurt a. M.