

Theories of Meaning for the Internet of Things



Michael Färber, Yulia Svetashova, and Andreas Harth

1 Why Traditional Knowledge Representation Is Insufficient

Future information systems, such as virtual assistants, augmented reality systems, and semi-autonomous or autonomous machines (Chan et al. 2009; Hermann et al. 2016), require access to large amounts of world knowledge in combination with sensor data. Consider a smart home scenario involving interconnected light bulbs. Here, a desired rule could be: “switch on the light in the hallway when somebody enters the home and set the light level in the hallway to below 50 lux.” In this scenario, there needs to be a common understanding (i.e., *semantics*) of all the information (concepts and facts) mentioned in this command between the user and the device, such as “light,” “hallway,” “50 lux,” but also of situational aspects, such as “when somebody enters the home.” In a Health 2.0 scenario, connected devices measure parameters concerning a patient’s health. The data need to be transformed (ideally automatically) into symbolically grounded knowledge and combined with the existing knowledge about health, diseases, and treatments (Henson et al. 2012).

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-69823-2_3) contains supplementary material, which is available to authorized users.

M. Färber (✉)

Karlsruhe Institute of Technology (KIT), Institute AIFB, Karlsruhe, Germany

e-mail: michael.farber@kit.edu

Y. Svetashova

Robert Bosch GmbH, Corporate Research and Advance Engineering, Robert-Bosch-Campus 1, 71272 Renningen, Germany

e-mail: svetashova@gmail.com

A. Harth

Friedrich-Alexander-University Erlangen-Nuremberg, Nuremberg, Germany

e-mail: andreas.harth@fau.de

Fraunhofer IIS-SCS, Nordostpark 84, 90411 Nuremberg, Germany

© The Author(s) 2021

L. Bechberger et al. (eds.), *Concepts in Action*, Language, Cognition, and Mind 9, https://doi.org/10.1007/978-3-030-69823-2_3

These examples demonstrate that knowledge representation for Internet of Things scenarios is needed. Specifically, on closer inspection, they indicate that three aspects are particularly essential for the Internet of Things knowledge representation:

1. How are perceptions and actions grounded and represented (in the Internet of Things terminology, sensors and actuators)?
2. How can machines and humans agree on a common understanding when referring to concepts and facts, and how can this common understanding be shared?
3. How can changes in the world be used in knowledge representation?

In the past, research on knowledge representation in computer science has mainly focused on developing and using static ontologies (i.e., as a *formal, explicit specification of a shared conceptualization in a domain of interest* (Studer et al. 1998)) and knowledge graphs (Fensel et al. 2020; Färber et al. 2018). Ontological languages, such as the *Resource Description Framework* (RDF) (Cyganiak et al. 2014), *RDF Schema* (Brickley and Guha 2014), and the *Web Ontology Language* (OWL) (Bechhofer et al. 2004), have been established to model parts of the world. To connect the world knowledge with sensor data, a few ad-hoc solutions have been proposed (e.g., Bonnet et al. 2000; Ganz et al. 2016, and Sect. 3). However, in our minds, all these technology is not capable of sufficiently incorporating the aspects of the Internet of Things as outlined above.

In this chapter, we want to take up the previous considerations on knowledge representation in the context of the Internet of Things; we thereby make use of content from epistemology—particularly, the semantic theories—for our discussion on an optimal knowledge representation, addressing research question 1 “How can we formally describe and model concepts?” outlined in Chap. 1 of this book. We can show that the problem of knowledge representation for the Internet of Things is by no means trivial and that questions about concrete implementations lead to fundamental questions of knowledge representation, such as the symbol ground problem (Harnad 1990) and the intersubjectivity problem (Reich 2010).

The topic of this chapter is highly interdisciplinary. Consequently, it is written for a diversity of user groups:

- Computer scientists, cognitive scientists, and IT practitioners who work on artificial intelligence systems and who are particularly interested in knowledge representation for the Internet of Things (e.g., designing ontologies for the Internet of Things);
- Philosophers who want to make autonomous systems and the Internet of Things accessible for their theories.

The chapter is structured as follows: After a detailed statement of the research problem in Sect. 1.1, we outline in Sects. 1.2 and 1.3 how our research problem is embedded in the scientific landscape of philosophy and computer science, respectively. In Sect. 2, we present a scenario in the Internet of Things context, which is used in the following sections to illustrate the concrete influences of theories of meaning on Internet of Things applications. Section 3 is dedicated to several semantic theo-

ries originating from philosophy and how they can be used to address our research problem. The chapter finishes with a summary in Sect. 4.

1.1 Problem Statement and Methodology

Problem Statement. The Internet of Things (IoT) refers to the idea of the “pervasive presence of a variety of things or objects around us—such as Radio-Frequency Identification (RFID) tags, sensors, actuators, mobile phones, etc.—which, through unique addressing schemes, are able to interact with each other and cooperate with their neighbors to reach common goals” (Atzori et al. 2010). The Internet of Things has emerged as an important research topic and paradigm that can greatly affect a variety of aspects of everyday life. In the private setting, examples are smart homes, assisted living, and e-health. In the business setting, the Internet of Things is used, among other things, for automation and industrial manufacturing, logistics, and intelligent transportation.

We focus on the connection between the Internet of Things and knowledge representation. As such, we consider *intelligent agents*—defined as objects acting rationally (Russell and Norvig 2010) and often perceived as being identical to smart information systems—that

- have *sensors* (e.g., cameras, microphones, sensors for temperature and humidity) that allow them to perceive the environment,
- have *actuators* (e.g., displays, motors, light bulbs) that allow them to act in an environment,
- have an *interface* (e.g., buttons and dials, speech) to communicate with humans and
- can act semi-autonomously or autonomously.

In the future, humans and agents will increasingly co-exist side by side. For instance, humanoid robots with conversational artificial intelligence capabilities might become omnipresent. Moreover, agents will communicate with each other and thereby exchange knowledge to accomplish tasks in an autonomous way. However, obtaining a common understanding of the shared world and having the ability to refer to the same objects during communication is from an epistemological point of view nontrivial and by no means a matter of course. The crucial aspect in this context is the gap between the *represented world* (also called the *model*) and the *actual world* (see Fig. 1). It is related to mind-body dualism and specifically Descartes’ mind-body problem in philosophy (Skirry 2006). Agents have access to the outside world (typically called *perception of the environment*) and are able to trigger changes in the world via actuators (i.e., they can change the outside world). This aspect is also related to the following questions: How can someone obtain the meaning of a text in a language unknown to him or her? How can someone interact with people without

the ability to speak the language of the people (see the Chinese room argument (Cole 2014))?

Methodology. We will outline the possibilities of modeling things for scenarios in the world of the Internet of Thing. Given the Internet of Things, an environment in which agents are situated with other agents, a *theory for knowledge representation on the Internet of Things* needs to

1. include sensors (perception) and actuators (action) in its formalization;
2. include the notion of multiple subjects (machines and humans) that interact with the environment and with each other (intersubjectivity);
3. include ways to describe changes in the represented world that should mirror changes in the real world, and vice versa (dynamics).

Acquiring the correct underlying foundations—and, in philosophical terms, the correct conditions of possibilities for acquiring and exchanging knowledge—is crucial to enabling the manifold benefits that arise from increased automation and human-computer interaction. As an example, let us take one of the prominent scenarios in the specific context of the Internet of Things—the so-called “onboarding” of devices. Onboarding is the process of connecting a sensor or a more complex Internet of Things device to the Internet and to a platform establishing an initial configuration and enabling services (Balestrini et al. 2017; Gupta and van Oorschot 2019). This process can either be automated or involves broad communities of device owners. In both cases, the problems of device-platform communication and deciding on identifiers (how to address a specific new device) require the acceptance of an adequate theory of meaning in the open context system. Such a system interacts with the changing world and needs to adapt accordingly.

This fact has already been noted by noteworthy philosophers and cognition scientists, such as (Gärdenfors 2000):

When building robots that are capable of linguistic communication, the constructor must decide at an early stage how the robot grasps the meaning of words. A fundamental methodological decision is whether the meanings are determined by the state of the world or whether they are based on the robot’s internal model of the world. (Gärdenfors 2000, p. 152)

Gärdenfors does not describe scenarios involving intelligent agents and does not show how the perception layer of a robotic system fits into his model of geometric spaces, which is the problem we address in this chapter. Specifically, we focus on *perception, multiple subjects, and world changes*.

1.2 Existing Solutions in Philosophy

In philosophy, the study of what knowledge is and how it can be represented (i.e., *epistemology*) and the study of how to acquire knowledge from an environment (i.e., *philosophy of perception*) are highly relevant to addressing the problem of knowledge




Fig. 1 Mediated reference theories distinguish between the world and a model of the world. Direct reference theories, in contrast, do not distinguish between the model and the world (i.e., the model is the world; illustration adapted from Sowa (2005))

representation for the Internet of Things. From these research areas, we can highlight the following aspects.

Theories of Meaning. Defining the meaning (particularly in the context of language also referred to as *semantics*) has always been an integral part of philosophy. In the 20th century, philosophy shifted its focus to language and the role of language in understanding. Particularly noteworthy is the groundbreaking work of Gottlob Frege (1848–1925), which can be seen as the basis for many achievements in the area of artificial intelligence. Frege’s ideas come together in a mediated reference theory (see Fig. 1).

Frege challenged the belief that the meaning of a sentence directly depends on the meaning of its parts. The meaning of a sentence is its truth value and the meaning of its constituent expressions is their reference in the extra-linguistic reality. First, he explored the role of the proper names (which have direct reference) and concepts (which gain meaning only when their direct referent is specified). He then studied identity statements (in the form of $a = a$ or $a = b$) and came to the conclusion that direct reference theories do not adequately capture the meaning of identity statements. In particular, he pointed to the fact that the statements “Hesperus is the same planet as Hesperus” and “Hesperus is the same planet as Phosphorus” do not mean the same thing, even though the terms “Hesperus” and “Phosphorus” refer to the same extra-linguistic entity, the planet Venus. Thus, he came to an important distinction: the reference (*Bedeutung*) of a sentence is its truth value and the sense (*Sinn*) is the thought which it expresses. The questions that originated from Frege’s arguments gave rise to many theories of meaning in logic and computer science and contoured the definition of meaning we accept in this chapter. Overall, Frege as a philosopher provided categories that other scientists questioned and developed.

We define meaning pragmatically as follows:

- *The meaning of symbols.* Following the idea that meaning is referential (see Gärdenfors 2000, pp. 151), the meaning of symbols (including names) are the objects to which they refer. For real-world entities, one can point directly towards the objects or mention their proper names. For abstract concepts (classes) and properties, we take the Cartesian-Kantian two-world assumption as a basis and assume that, besides the real world we can see (*res extensa* or *phaenomenon*), there exists the world of ideas or thoughts (*res cogitans* or *noumenon*), and that classes and properties exist in this intellectual world.
- *The meaning of sentences.* Symbols can be arranged together to express statements (used synonymously to facts; sentences are the written counterpart). Statements bring us to a new level of meaning, the level of truthfulness. Each sentence can be *true* or *false*.

Theories of Truth. Given that statements can be true or false, questions of how statements stand in relation to the world and how statements can be tested concerning their truthfulness arise. Among the most commonly used theories of truth are as follow.

- *The correspondence theory of truth:* This theory can be considered the most basic theory of truth. True sentences capture the current state of affairs—objectively, without an observer. This theory is very much based on the actual world. The theory does not consider linking new knowledge to existing knowledge of a subject and differentiating between the varying knowledge levels of different people and how these people agree on the same meanings.
- *The coherence theory of truth:* This theory is coherence-centric and takes into consideration how new knowledge is incorporated into existing knowledge. Statements are considered to be true if they are consistent with the statements (i.e., knowledge) obtained so far.
- *The consensus theory of truth (pragmatic theory of truth):* This theory takes the different views of observers into consideration and is designed to align subjective views to other views. Statements are considered true if people (i.e., observers) agree on them.

It becomes immediately clear that these theories of truth do not exclude each other but rather have different foci. We argue that a comprehensive theory needs to take all of the theories' aspects into account. Particularly noteworthy is the fact that the theories focus mainly on knowledge and truth at a given point in time (see the construction of ontologies in computer science). Dynamic aspects, and thus the modeling of events, are insufficiently covered by these foundational theories.

1.3 Existing Solutions in Computer Science and Logic

In computer science and cognitive sciences, specifically the fields of knowledge representation and logic, the problem of how to represent knowledge about the world for Internet of Things scenarios has been addressed to some degree.

Theories of Meaning. In the past, computer scientists and logicians defined the meaning of objects in their knowledge representation models (e.g., ontologies) and methods for describing the world largely without an explicit connection to reality and perception. In particular, *model theory* (Tarski 1944) is the established way of defining the meaning of logic-based knowledge representation languages, such as the semantic web languages RDF, RDFS, and OWL.

Moreover, in the area of knowledge representation, it became popular to use *ontologies* (Staab and Studer 2010) and *knowledge graphs* (Fensel et al. 2020) as world models. Freely available open knowledge graphs form the Linked Open Data (LOD) cloud, which is used in various applications nowadays (Färber et al. 2018). However, since logic and model theory are very formal disciplines, there was no need to link knowledge representation to perception. Works on *ontology evaluation* and *ontology evolution* consider the process of creating and evaluating ontologies (in the sense used in computer science, i.e., as a formal model of a small domain of interest) as finding the lowest common denominator for modeling parts of the world. However, researchers mainly discuss common and best practices a team of developers can use to create an ontology. Early attempts at defining an ontology which incorporate temporal dynamics were made by Grenon and Smith (2004) and Heflin and Hendler (2000).

Overall, existing methods for modeling the world and defining meaning have the following drawbacks: (1) They disregard any explicit connection to reality. (2) They are omniscient and try to capture an (imposed) objective view of the world. (3) They are only able to express static knowledge but not changes in the world to a sufficient degree. In Sect. 1.2, we have carved out similar drawbacks regarding existing theories of meaning and truth.

If symbols are only identifiers, how can our minds create a link to an object in the real world (or in our conceptual worlds of ideas or thoughts)? How can we make sure that other subjects/minds have the same meaning; that is, link to the same object (e.g., when we only mention the object’s identifier, such as <http://dbpedia.org/resource/Karlsruhe> or <http://wikidata.org/entity/Q1040>)? Is the meaning directly connected (grounded) to non-symbols? This problem is known as the *symbol grounding problem*: “How can you ever get off the symbol/symbol merry-go-round? How is symbol meaning to be grounded in something other than just more meaningless symbols?” (Harnad 1990). In the Semantic Web and Linked Data context, URIs are used as symbols for objects. The symbol-grounding problem is not often considered (Cregan 2007) or even solved. In particular, the aspects of perception, multiple subjects, and changes in the world—the focus in our chapter—for knowledge representation are not covered sufficiently. In the Internet of Things domain, we find only a

few works in this respect, such as the article by Hermann et al. (2017), who present grounded language learning in a 3D environment.

Theories of Truth. Theories of truth are traditionally proposed in philosophy. When we apply the theories of truth as introduced in Sect. 1.2 to the established and widely used semantic web technologies, such as RDF and OWL, and to knowledge representation ideas like knowledge graphs and linked open data, we can observe the following: (1) The RDF data model (Hayes and Patel-Schneider 2014) might fit to the *correspondence theory of truth* and to the *consensus theory* in the context of the Internet of Things. (2) Linked data can be regarded as an implementation of the *consensus theory* in the sense that data publishers and data consumers need to agree on common terms to use the linked data in a reasonable way. However, applications in the Internet of Things require more, since the (linked) data are subjected to changes over time and dependent on the perception (see, e.g., the sensor data from devices).

In recent years, approaches based on neural networks have been presented to represent entities and relations in knowledge graphs—as an implementation of a knowledge representation—in the form of vectors in a low-dimensional vector space (called *embeddings* Mikolov et al. 2013; Wang et al. 2017). Apart from the context of the entities and relations in the knowledge graph, external data sources have also been used to build these implicit knowledge representations. For instance, data from several modalities (text, images, speech, etc.) can be combined to form a unified, comprehensive representation in a low-dimensional space (Bruni et al. 2014). In the Internet of Things context, the representations are created based on sensor data, and thus, perceptions. We can argue that the formal method and technology of obtaining the sensory data (e.g., images, text, etc. of an object) and of transforming it into a common vector space (e.g., via machine learning techniques) has a direct influence on the meaning of objects or even constitutes the meaning itself.


2 Motivating Scenario

In this section, we describe a smart home scenario, which will be used in the upcoming sections as an example of an Internet of Things scenario. It will show how the theories of meaning considered by us affect the way of modeling knowledge.

Consider the home of Alice (see Fig. 2) with four rooms: the hallway, the living room, the bathroom, and the bedroom. Each of the rooms is equipped with a light bulb that can be controlled via a network interface. Each room also has a window with controllable window blinds. Moreover, each room has a sensor to measure the light levels. The door has a sensor that detects when it is opened. A virtual assistant called Bob provides a user interface to the smart home via speech interaction. The more data and knowledge about the smart home is coupled with the virtual assistant, the more generic and flexible the virtual assistant needs to be.

Considering this scenario, we can point out several issues with respect to knowledge representation. The first issue concerns *naming*. Both Alice and Bob have to

Fig. 2 A smart home with various sensors (i.e., luminance sensors and door sensors) and actuators (i.e., network-controlled lamps and window blinds)



agree on the meaning of “the living room,” so that Alice can ask Bob, “Is the light on in the living room?” Similarly, to affect a change in the world, Alice and Bob have to agree on names as references to objects, so that Alice is able to tell Bob to “switch off the light in the living room.” A more elaborate command could be to “switch on the light in the hallway when somebody enters the home and set the light level in the hallway below 50 lux.”

We assume that a shared understanding between the virtual assistant and the human user has to be configured when setting up the smart home (the so-called “onboarding problem”). The problem also arises when a new human user wants to interact with the smart home (e.g., when Carol visits Alice and wants to turn on the lights).

We can think of various other Internet of Things scenarios in which theories of meaning (also called *semantic theories*) become important for modeling the scenarios. For instance, in a Health 2.0 scenario as outlined by Henson et al. (2012), the sensor data gathered by Internet of Things devices need to be collected and transformed into symbolic information. This transformation allows the system to interpret the information and combine it with other existing, symbolically grounded knowledge (e.g., about diseases). Questions concerning the representation of perception, the inter-subjective agreement of concepts and facts, and the representation of dynamically changing knowledge arise.

3 Applying Theories of Meaning to the Internet of Things

Several theories of meaning have been proposed to link the real world with actual knowledge about it. In this section, we review the following semantic theories:

1. Model-theoretic semantics;
2. Possible world semantics;
3. Situation semantics;
4. Cognitive and distributional semantics.

These semantics have been chosen due to their popularity and “baselines” in previous work (Gärdenfors 2000, pp. 151). The first formalism is sometimes referred to as “extensional semantics” and the second formalism is referred to as “intensional semantics.” Furthermore, some authors, such as Gärdenfors (2000), refer to “extensional semantics” instead of model theory and “intensional semantics” instead of “modal logic.” Given the various and sometimes incompatible uses of “extensional” and “intensional” in the literature (Janas and Schwind 1979; Helbig and Glockner 2007; Lanotte and Merro 2018; Franconi et al. 2013), we use the terms “model-theoretic semantics” and “possible world semantics” for clarity.

In the following sections, we cover each theory of meaning in detail and apply it to the Internet of Things. Within each section, we first give a definition of the theory and outline its characteristics. Subsequently, we describe how the theory can be applied to model Internet of Things scenarios. We thereby focus primarily on the perception, intersubjectivity, and dynamics, because modeling these aspects is particularly crucial in the context of the Internet of Things (see Sect. 1).

3.1 *Model-Theoretic Semantics for the Internet of Things*

3.1.1 Definition and Current Use

Model-theoretic semantics can be encoded in various ways. In the following, we assume that the knowledge in embodied systems (e.g., a smart home) is described using sentences in first-order (predicate) logic. The meaning (i.e., the truth value) of the sentences is given via mapping to a world represented using set theory.

Extensional semantics is considered one of the realistic theories on semantics (Gärdenfors 2000). Expressions (names) are mapped to objects in the world (see the *theory of correspondence*). Predicates are then applied to a set of objects or relations between objects. Generally, using such a map, sentences can be assigned *true/false* values (see *truth conditions*). The “extension” of the sentence “Lassie is famous” is the logical value “true,” since Lassie is famous. There is no anchoring of the language in a body (i.e., the meaning of words is modeled independently of individual subjects). This is known as the human capability of abstraction. All sentences being true constitute the world.

First-order predicate logic provides the foundation for formalizing current Semantic Web languages, such as RDF, RDFS, and OWL.

3.1.2 Application to the Internet of Things

While the languages with a formalization in model theory are mature and widely used, they do not cover the dimensions required in scenarios around the Internet of Things as outlined in the following:

Perception

The set-theoretic structure representing the world does not have any connection to the external world. Whether or not the term “Lassie” refers to Lassie the dog in the external world does not have any bearing on the truth value of the sentence. However, such a connection is needed to take perception (e.g., sensor data) in Internet of Things scenarios into account for modeling the world.

Intersubjectivity

The theory does not address the problem of reaching agreement on the meaning of terms across different agents. For instance, in the case of the semantic web languages RDF, RDFS, and OWL, there exists no defined mechanism that ensures different agents have the same notion of terms and sentences. Finding a shared understanding is left to the agents.

Dynamics

Traditional first-order predicate logic was developed to describe properties of things. That is, one can name things (“Lassie”) and assign properties to them (“is famous”). The focus of such representations is to deduce new declarative sentences based on the given sentences. Some applications use first-order logic to represent events (e.g., “Lassie rescues the girl from drowning”), where the event (“rescuing”) is treated as a property. While such representations might be suitable for some derivations, they do not cover the dynamics behind events sufficiently for scenarios in the Internet of Things.

Benefits and Limitations for the Internet of Things. The focus of model theory is to provide a notion of truth of sentences that allows for the specification of logical consequence. Logical consequence can help one check for satisfiability of sentences with regards to the world. It provides means to integrate data from multiple sources. However, model theory does not consider many aspects relevant in the Internet of Things, such as the connection of symbols and sentences to the real world or the question of how multiple agents can agree on the meaning of symbols. Furthermore, model theory lacks means to adequately formalize change, since the sentences are classically interpreted over a static model of the world.

3.2 Possible World Semantics for the Internet of Things

3.2.1 Definition and Current Use

The origins of possible world semantics can be traced back to Carnap (1947), Kripke (1959), and Montague (1974). Without loss of generality, we assume for the remaining part of the chapter that the possible world semantics are implemented via modal logic. More on the idea of possible worlds as the conceptual underpinning

of the modal logics can be found in Hughes et al. (1996) and Menzel (2017). In the following, we review the modal logic and its applicability for the Internet of Things scenarios.

With modal logic, expressions are mapped to a set of possible worlds, instead of a single world. Otherwise, the setting is the same as for the extensional semantics theory: sentences can have “truth conditions”, and each proposition (sentence) has worlds in which it holds true.

To model these possible worlds, modal logic adds two new unary operators: \Box (“necessary”) and \Diamond (“possibly”) to the set of Boolean connectors (negation, disjunction, conjunction and implication). The proposition is possible, if a world may exist in which this proposition is true. The proposition is necessary, if it has to be true in all worlds.

Dependent on the application context, modal operators can have different intuitive interpretations. For example, if one wants to represent temporal knowledge, $\Box_{future} P$ may mean that proposition P is *always* true *in the future* and that $\Diamond_{future} P$ means P is *sometimes* true *in the future*. These different ways to interpret modal connectives give rise to various types of modal logics: tense, epistemic, deontic, dynamic, geometric, and others (see more in Goldblatt (2006)). Thus, they represent facts that are “necessarily/possibly” true, true “today/in the future”, “believed/known” to be true, true “before/after an action”, and true “locally/everywhere.”

3.2.2 Application to the Internet of Things

We see many possibilities to use modal logic to capture the semantics in Internet of Things scenarios. As an example, Fig. 3 shows a system that interprets the voice input “Turn on the light” and acts differently depending on the location of the user. We can also consider such parameters as time of day and define different scenarios with temporal logics.

Modal logic as a kind of formal logic extends predicate logic by allowing it to express possibilities. Modal logic has mainly been used in formal sciences, such as logic (e.g., “ontology of possibilities”). However, it has not been applied extensively in computer science and, specifically, in Internet of Things contexts. We can observe that modal logic as an implementation of possible world semantics is better suited to the Internet of Things than model-theoretic semantics. However, modal logic is not perfectly suited for modeling knowledge of Internet of Things agents. This can be demonstrated by evaluating perception, intersubjectivity, and dynamics.

Perception

Similar to first-order predicate logic with a model-theoretic formalization, modal logic does not have any connection to the external world.

Intersubjectivity

Modal logic and its semantics are still based on a realistic idea (i.e., coordinating extra-linguistic entities to linguistics expressions). However, subjects’ interpretations

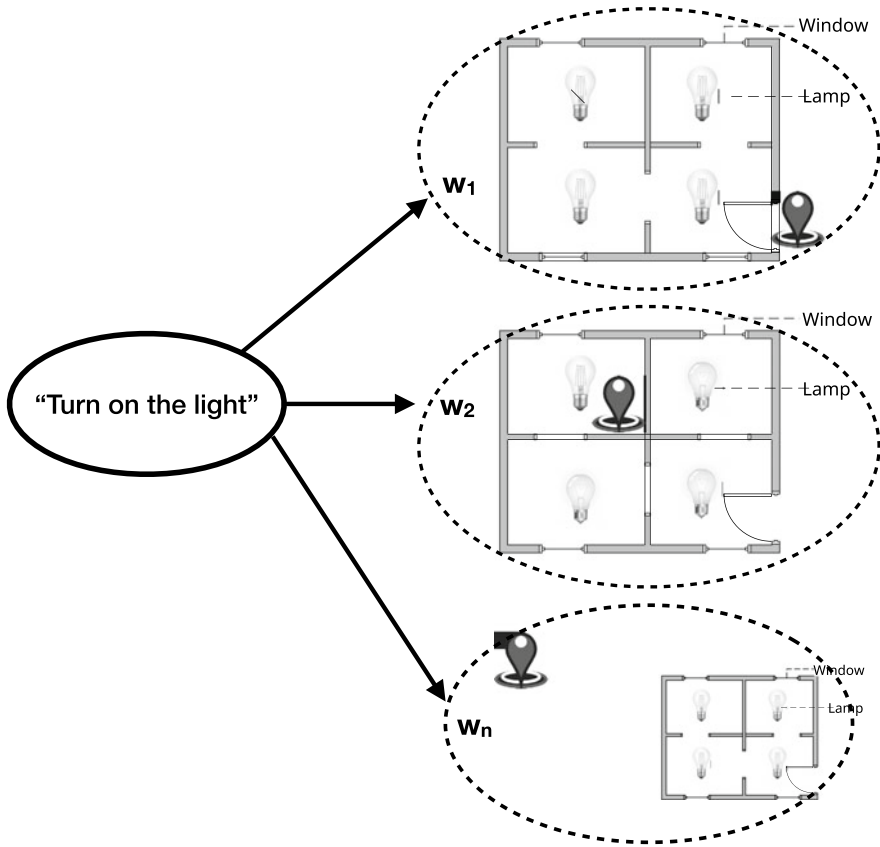


Fig. 3 Possible worlds in the smart home scenario

of the world can be represented as distinct worlds. In this way, modal logic allows us to model multiple worlds and to represent the knowledge of several agents (i.e., subjects).

Dynamics With the ability to add temporal operators, modal logic allows us to keep track of states of resources over time. Furthermore, with the ability to keep track of state over time, one can detect events (i.e., state changes) and thus represent knowledge evolving over time.

Benefits and Limitations for the Internet of Things. The focus on logical consequence of sentences is one of the properties that possible world semantics shares with model-theoretic semantics. Neither has an explicit connection to the real world. Modal logic as an implementation of possible world semantics stands out from implementations of model-theoretic semantics by taking the aspects of intersubjectivity

and dynamics into account. Nevertheless, the possible world semantics only provide means to describe a changing world with sentences and to reason over such sentences, but not to actually affect changes in the world.

3.3 Situation Semantics for the Internet of Things

3.3.1 Definition and Current Use

The theory of situation semantics, another kind of realistic semantics, was developed by Jon Barwise and John Perry in their seminal book *Situations and Attitudes* (1983). In contrast to its predecessor *possible worlds semantics*, it postulates the principal of *partiality* of information available about the world. Limited parts of the world that are “clearly recognized [...] in common sense and human language” and “can be comprehended as a whole in [their] own right” (Barwise and Perry 1980) are called *situations*. Situations stand in contrast to processes and activities. According to (Galton 2008):

I believe that open processes and closed processes are very different kinds of things. The fact that we use the word ‘process’ for both of them perhaps lends some support to Sowa’s use of this word as the most inclusive term, corresponding to what others have called situations or eventualities.

Devlin (2006), who formalized the basic notions of situation semantics and extended it to situation theory, emphasizes that information is always given “*about* some situation.” It is constructed from discrete information units, called *infons*. An infon (σ) is a relational structure of shape, $\langle\langle R, a_1, \dots, a_n, 0/1 \rangle\rangle$, where R is an n -place relation, a_1, \dots, a_n are objects appropriate for the argument roles i_1, \dots, i_n , and $0/1$ are the *polarity* values indicating whether or not the objects a_1, \dots, a_n stand in the relation R .

Objects in the argument roles of an infon include individuals, properties, relations, space-time locations, situations, and parameters. *Parameters* in situation semantics act as variables (i.e., they reference arbitrary objects of a given type). To set parameters to concrete real-world entities, Barwise and Perry (1983) introduce an assignment mechanism called an *anchor*.

Unlike model-theoretic or possible worlds semantics, situation theory claims that an infon—roughly corresponding to a fact or statement—can be true (or false) only in the context of a particular situation. This relationship is written as $s \models \sigma$ (read as “ s supports σ ”), meaning that the fact represented by infon σ holds true in situation s .

Figure 4 shows an illustrative example of situation semantics for the Internet of Things scenario *smart home*. In this figure, we can see a limited part (s) of the world where we can distinguish several classes of objects: `WindowBlind`, `Room`, and `LightBulb`. Potentially, instances of these classes can be involved in many situations. One of them (i.e., `TriggerBlindsUp`) is that when it is dark in the room and already light outside in the morning, the window blinds are automatically raised

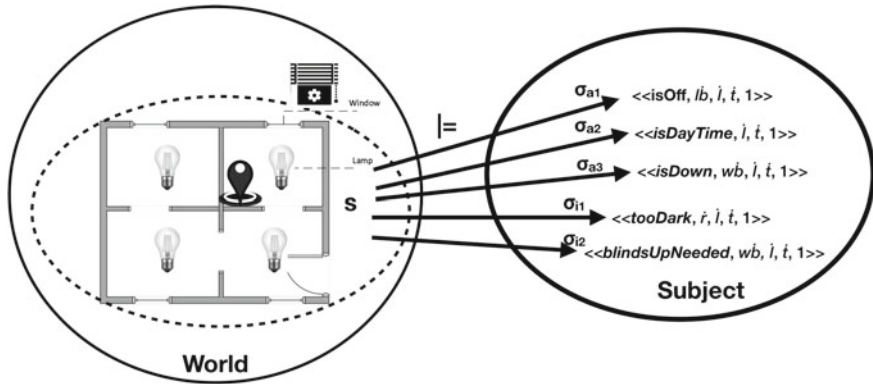


Fig. 4 TriggerBlindsUp situation in the smart home scenario

by the control system. We represent the relevant relations (*isDayTime*, *tooDark*, etc.) with the following infons where parameters \dot{l} and \dot{i} reference arbitrary spatial and temporal locations:

- (σ_{a1}) $\langle\langle isOff, \dot{lb}, \dot{l}, \dot{i}, 1 \rangle\rangle$, where parameter \dot{lb} anchors objects of type `LightBulb`;
- (σ_{a2}) $\langle\langle isDayTime, \dot{l}, \dot{i}, 1 \rangle\rangle$;
- (σ_{a3}) $\langle\langle isDown, wb, \dot{l}, \dot{i}, 1 \rangle\rangle$, with \dot{wb} anchoring `WindowBlind` instances;
- (σ_{i1}) $\langle\langle tooDark, \dot{r}, \dot{l}, \dot{i}, 1 \rangle\rangle$, with \dot{r} anchoring `Room` instances;
- (σ_{i2}) $\langle\langle blindsUpNeeded, wb, \dot{l}, \dot{i}, 1 \rangle\rangle$, with \dot{r} anchoring `WindowBlind` instances.

By using conjunction, disjunction, and anchoring, we can combine infons into more complex structures (i.e., *compound infons*). For situation `TriggerBlindsUp`, the infons form the compound infon: $s \models \sigma_{a1} \wedge \sigma_{a2} \wedge \sigma_{a3} \wedge \sigma_{i1} \wedge \sigma_{i2}$. The system that relies on this formalism can check whether these infons support the situation `TriggerBlindsUp` and use actuators to trigger the change in the real world.

Situation semantics distinguishes three types of situations: *utterance* situation (i.e., the immediate context of utterance, including a speaker and a hearer), *focal* situation (i.e., the part of the world referred to by the utterance), and *resource* situation (i.e., the situation used to support or to reason about focal or utterance situations (Devlin 2006)).

Meaning is acquired by linking utterances expressed in language to objects in the real world. This link, called the “speaker’s connection” (Barwise and Perry 1983), determines the unique role of a subject in this theory: It is the agent who establishes such a link, and meaning is thus made relative to a specific agent. Figure 4 illustrates this possibly changing perspective. The subject perceives the room as dark: $\langle\langle tooDark, \dot{r}, \dot{l}, \dot{i}, 1 \rangle\rangle$; one can imagine another subject for whom the polarity of the infon σ_{i1} would be 0.

In the area of the Internet of Things, certain information systems employ situation semantics as the core of their modeling of user behavior and sensor observations,

as well as the basis of context- and situation-awareness (see Heckmann et al. 2005; Kokar et al. 2009; Stocker et al. 2014, 2016). In the following, we will refer to these systems to show how situation semantics addresses the problems of perception, intersubjectivity, and representing dynamics.

3.3.2 Application to the Internet of Things

In the process of measurement, sensors transform signals of physical properties into numbers, thus generating numerical data. These data are challenging to store and manage and require near-instant access. The interpretation of the raw values requires modeling, finding patterns, and deriving abstractions. Abstractions reveal the properties of the observed real-world entities, show their dynamics, and place them into relations with their surroundings.

Perception

Sensor networks cannot perceive (“observe”) situations directly; instead, as shown in Fig. 5, several components are needed to derive decisions and to take actions (see Kokar et al. 2009; Stocker et al. 2014). The process can be described as follows: The system takes sensor data as input, which then undergo the semantic enrichment process. Semantically annotated data is then transformed via a rule-based inference, digital signal processing, or machine learning algorithms into higher-level abstractions. These abstractions can be considered situations, which in turn can trigger actions and enable intelligent services. Both Stocker et al. (2014) and Kokar et al. (2009) exemplify how sensor input is transformed into a set of infons (called *observed* or *asserted* in Kokar et al. (2009)) and how new *inferred* infons are derived from them.

Situation semantics, therefore, works as a compliment to the algorithms that can directly process data generated in the perception layer. It is the way to organize sensory input in a task or goal-oriented environment. In addition, Stocker et al. (2014) argue that the persistence of situational knowledge in many cases is a desirable alternative to the persistence of sensor data and the key enabler of useful perceptual data in real time. Henson et al. (2012) describe an approach for deriving abstractions—essentially similar to situations—from sensory observations.




Fig. 5 Generic components of a system consuming sensor data

Intersubjectivity

In situation semantics, any relation between a real-world situation and its representation in a formal framework is relative to a specific subject. An agent recognizes or, in the terminology of Barwise and Perry (1983), “individuates” situations. Assigning values to certain parameters in the argument roles of an infon is always done by a particular subject. Situation semantics has an inherent mechanism to encode the subject’s perspective, as well as to represent and to coordinate views of multiple subjects. The Internet of Things is often treated as a decentralized distributed system (Singh and Chopra 2017) where different agents generate situational knowledge individually. In this context, formalizing situation semantics can ease inter-agent communication and data integration (see the discussion in Stocker et al. (2014)).

Dynamics

Having *situation* as its central concept, situation semantics considers static representation of situations (as objects and their relations) and their dynamic aspect. According to Barwise and Perry (1983), “Events and episodes are situations in time ... changes are sequences of situations.” As a consequence, this theory has a built-in mechanism for representing temporal and spatial dynamics; namely, it introduces special types of objects that can fill argument roles of an infon (i.e., TIM, the type of a temporal location, and LOC, the type of a spatial location (Devlin 2006)). Thus, it is possible to represent whether a relation holds between the objects at a particular time in a particular location.

Stocker et al. (2014) use situation semantics to model observed situations in a road traffic scenario. By analyzing the road-pavement vibration data from three accelerometer sensors, they were able to detect vehicles in the proximity of sensing devices (*near*-relation) and their types (*light* or *heavy*). Observations, classified by the signal processing algorithms and modeled as sets of infons of shape, $\langle \langle \text{near}, \text{Vehicle}_x, l_x, t_x, 1 \rangle \rangle$, enabled the inference of the velocity and the driving side of a vehicle via a custom set of rules. This example shows that this kind of representation is suitable for time-oriented data. Time-oriented data is a characteristic of most of the data generated in the Internet of Things (see more in Serpanos and Wolf (2017)).

Benefits and Limitations for the Internet of Things. Barwise and Perry were not the first to include situations as first-class citizens into a knowledge representation theory (see, e.g., situation calculus McCarthy 1963; McCarthy and Hayes 1969). Nevertheless, compared to its predecessors, situation semantics presents a richer formalism capable of representing higher level abstractions over raw sensor data, multiple viewpoints, and temporal-spatial dynamics.

Infons with their argument role structures can be reused across related situation types (e.g., how easy it will be to project the set of infons of the `TriggerBlindUp` situation to `TriggerBlindDown`). In many Internet of Things scenarios, the storage of raw data is not optimal due to the quantity and limitations of existing storage

solutions. Having a system as described in this section will allow us to store more meaningful and actionable pieces of information (situations) for certain signals to act upon in real time.

3.4 Cognitive and Distributional Semantics for the Internet of Things

3.4.1 Definition and Current Use

Cognitive semantics needs to be considered with respect to the general notion of *cognition*: instead of a subject perceiving the world with his senses and with language as the subject's ability to talk about the world, the focus is shifted to the mental representation of the world (i.e., to the subject's cognitive structures). Moreover, language becomes part of the cognitive structure. As such, concepts are elements in the subject's cognitive structure and without a direct reference to a reality. Thus, the meaning of concepts, etc., does not go beyond language, but is nothing else than using the language itself (see Ludwig Wittgenstein's theory of language) and therefore, the cognitive structures. These cognitive structures are subject to constant adaptation due to the interaction with the world. For instance, new concepts are learned and new findings are obtained. The world becomes *viable*. Overall, cognitive semantics is categorized as a non-realistic theory of semantics due to the exclusion of reality.

Focusing on the subject's cognitive structures, the question becomes what these cognitive structures look like and how they are created. Motivated by the biology of the human brain as the basis for any human's cognitive ability (Gärdenfors 2000, p. 257), neural networks and their mechanisms are typically considered the basis for cognition. Inputs, outputs, and internal representations of neural networks are modeled mathematically as geometrical (vector) spaces. Vector spaces are therefore used to represent things in the world, such as entities, concepts, and relations. Thus, knowledge is represented as *distributional representations* (e.g., embeddings) on a sub-symbolic level. *Meaning* is formalized as and reduced to a *distance function*. Similar objects tend to be spatially closer to each other in the vector space induced by the used neural network. Semantics is considered to be distributional (leading to the term *distributional semantics*), geometrical, and statistical.

Cognitive semantics and distributional semantics is not a new phenomenon: In 1954, Harris (1954) proposed that meaning is a function of distribution (see the famous quote: "a word is characterized by the company it keeps" (Harris 1954)). Contemporary philosophers and cognitive scientists use geometrical spaces to explain cognition and how concepts are formed by subjects. Gärdenfors (2000), for instance, considered the geometry of cognitive representations. In this cognitive space, points denote objects, while regions denote concepts (see Fig. 6 and book Chap. 2 for more information about Gärdenfors' cognitive framework).




Fig. 6 Low-dimensional vector space representation in the smart home scenario with instances represented as points, concepts represented as areas, and predicates (relations) represented as vectors

Artificial neural networks have been used to simulate neural networks, and thereby, cognition. With the revival of research in artificial neural networks in recent years, research has been performed on how representations for terms, concepts, and predicates can be learned automatically (see, among other things, the approaches *TransE* and *TransH* (Wang et al. 2014)). The idea is to use the weights to the hidden layers of neural networks as representation (called *embeddings*). Guha (2015) proposed a model theory based on embeddings and adapted the Tarski model theory to embeddings.

In recent years, knowledge graph entities and relations (i.e., explicit knowledge representation formats) have also been embedded, showing that not only expressions can be represented in a distributed fashion, but also concepts and entities, as well as classes and relations. This allows us to model human cognition in a more natural way, because embeddings are learned for specific symbols.

3.4.2 Application to the Internet of Things

We assume that cognitive items, such as concepts, are represented in a sub-symbolic fashion, specifically, distributional semantics. Concepts are thus represented in a geometrical space. We use neural-network-based embedding methods as concrete implementation for distributional semantics. Figure 6 shows an example of representing items for the smart home scenario. Distributional semantics is amenable to modeling perception, intersubjectivity, and dynamics in the following respect:

Perception

Distributional semantics differs (with respect to perception) from other semantic theories in several ways:

- The meanings of concepts and facts are represented in a distributed fashion, not as singular units or symbols. In the smart home scenario, specific light bulbs and rooms are encoded as embedding vectors (see Fig. 6).
- The representations and, thus, the meanings of concepts are not static, but can be subject to constant change. To overcome this issue, time-dependent embeddings can be learned (Nguyen et al. 2018). In the smart home scenario, agents can learn

the embeddings of the different light bulbs and the embedding space of light bulbs per se based on the sensor data used as input for a neural network.

- A symbol grounding is possible as long as some form of input data (i.e., sensor data) is provided (and does not change abruptly).
- There are indications that human cognition aggregates the perceptions of different modalities of one unit (e.g., concept or concrete entity). For instance, the image of a dog and the sound of a dog are immediately perceived as belonging to the same unit. The same phenomenon can be observed when a multilingual person switches between languages whilst referring to the same concepts. As is the case with embedding methods from machine learning, such a fusion of sensor data from different modalities is possible. In the Internet of Things context, an embedding vector can be learned jointly based on different modalities.

Overall, perception is reduced to learning embeddings.

Intersubjectivity

Talking about and reaching an agreement on expressions between several agents can be traced back to using the same learned representations (i.e., embedding vectors) and the same conceptual structures (i.e., the distributional space). Even if different initialization values for the embedding spaces are given, subjects can use the same learning function to learn the same concepts. In the smart home scenario, the agents might differ in the exact points of the single light bulbs and rooms, since they rely on their own embedding learning and usage. However, they can agree on the same instances and concepts if the embeddings share the same characteristics (e.g., having nearly the same distances to other embeddings in the vector space). Overall, learning representations and meaning are reduced to learning and applying the same mathematical functions and models.

Dynamics

Describing changes in the world, such as events, is not sufficiently possible in the cognitive theory of semantics. If embeddings as distributed representations are learned or adapted online (i.e., in a permanent fashion, not only once at the beginning), then changes in the world may change the embeddings. However, the change itself is not represented. In the smart home scenario, an event might be light bulb number 4 switching on. The concepts involved in this event, such as light bulb #4, the positions of the light bulbs, and room #1 remain the same.

Benefits and Limitations for the Internet of Things. A characteristic of cognitive/distributional semantics is that information, such as concepts and facts, is not represented in the form of symbols, but in a sub-symbolic fashion as points and spaces in a vector space. This allows a more continuous distance function and an agreement on concepts and facts in the world as a continuous process. Talking about and reaching an agreement on expressions between several subjects can be traced back to using the same learned representations (i.e., embedding vectors) and the same conceptual structures (i.e., the distributional space). Thus, distributional semantics is heavily based on mathematics, which benefits the modeling of data in the Internet

of Things setting. However, describing changes in the world, such as events, is not sufficiently possible in the cognitive theory of semantics.

4 Conclusion

In this chapter, we have considered the theoretical foundations for representing knowledge in the Internet of Things context. Based on the peculiarities of the Internet of Things, we have outlined three dimensions that must be examined with respect to theories of meaning:

1. *Perception*: How can a theory of meaning incorporate “direct access” to the world (e.g., via sensors)?
2. *Intersubjectivity*: How can the world view of several subjects (i.e., agents in the Internet of Things) be modeled coherently?
3. *Dynamics*: How can the change of knowledge be modeled sufficiently? Which aspects of time can be represented?

We considered the following theories of meaning:

1. The model theory (extensional semantics)
2. Modal logic (intensional semantics)
3. Situation semantics
4. Cognitive/distributional semantics.

The single theories have the following advantages and disadvantages (see also Table 1):

1. *Model-theoretic semantics* is the simplest model in our series of considered semantic theories. This semantic theory can be used to formulate sentences and their truth values. However, it does not provide us with techniques or formalisms for modeling reality to the highest degree (i.e., with its unstable and experiential nature).

Table 1 Overview of how the challenges of perception, intersubjectivity, and dynamics are met by the various theories of semantics

	Model-theoret. semantics	Possible world semantics	Situation semantics	Conceptual/distrib. semantics
Perception			✓ (indirectly, via sensor data processed by machine learning)	✓
Intersubjectivity		✓	✓	✓ (communicating the dimensions)
Dynamics		✓	✓	

2. *Possible world semantics* also does not provide us with an explicit connection to the real world that we could use to adequately model knowledge in the Internet of Things domain. However, it allows implementations for modeling temporal modalities and agent's beliefs. As such, this semantic theory first lays theoretical foundations concerning dynamics and intersubjectivity. However, the foundational questions of how agents reach an agreement and how this fact about the agreement can be represented are not covered.
3. *Situation semantics* can be considered another layer in the pyramid of knowledge representation formalisms, as raw sensor data, multiple viewpoints, and temporal-spatial dynamics can be represented to some degree. However, we believe that this formalism is also not the optimal semantic theory for Internet of Things scenarios, as it leaves too many questions unanswered, particularly concerning perception and intersubjectivity.
4. *Cognitive and distributional semantics* can be judged in a manner similar to situation semantics when it comes to Internet of Things applications. Compared to the previous semantic theories, cognitive and distributional semantics are rather empirical (i.e., data-driven theories). The introduction of different levels of cognition and the fact that symbolic knowledge representation can be connected to sub-symbolic knowledge representation is appealing, particularly when it comes to data gathered by sensors. Intersubjectivity can be reduced to empirical training using data and mathematical functions (encoded in the form of [neural] networks). We see the main lack of this semantic theory in the (elegant) modeling of knowledge change over time.

Overall, we came to the conclusion that each of the semantic theories helps in modeling specific aspects, while not sufficiently covering all three aspects simultaneously. For the future, working on the advancements of situational semantics and distributional semantics and combining them towards a united semantic theory can be very fruitful for developing future intelligent information systems.

References

- Atzori, L., Iera, A., & Morabito, G. (2010). The internet of things: A survey. *Computer Networks*, 54(15), 2787–2805.
- Balestrini, M., Seiz, G., Peña, L. L., & Camprodon, G. (2017). Onboarding communities to the IoT. In *International Conference on Internet Science*, INSCI (Vol. 17, pp. 19–27). Springer.
- Barwise, J., & Perry, J. (1980). The situation underground. In J. Barwise, I. Sag, & Stanford Cognitive Science Group (Eds.), *Stanford Working Papers in Semantics* (Vol. 1, pp. 1–55). Stanford University Press.
- Barwise, J., & Perry, J. (1983). *Situations and attitudes*. Cambridge (Mass.): MIT Press.
- Bechhofer, S., van Harmelen, F., Hendler, J., Horrocks, I., McGuinness, D. L., Patel-Schneider, P. F., & Stein, L. A. (2004). OWL web ontology language reference. Retrieved May 03, 2020, from <https://www.w3.org/TR/owl-ref/>.
- Bonnet, P., Gehrke, J., & Seshadri, P. (2000). Querying the physical world. *IEEE Personal Communications*, 7(5), 10–15.

- Brickley, D., & Guha, R. (2014). RDF Schema 1.1. Recommendation, W3C. <http://www.w3.org/TR/2014/REC-rdf-schema-20140225/>. Latest version available at <http://www.w3.org/TR/rdf-schema/>.
- Bruni, E., Tran, N., & Baroni, M. (2014). Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, 49, 1–47.
- Carnap, R. (1947). *Meaning and necessity: A study in semantics and modal logic*. University of Chicago Press.
- Chan, M., Campo, E., Estève, D., & Fourniols, J.-Y. (2009). Smart homes-current features and future perspectives. *Maturitas*, 64(2), 90–97.
- Cole, D. (2014). The Chinese room argument. Stanford Encyclopedia of Philosophy. Retrieved September 02, 2019, from <https://plato.stanford.edu/entries/chinese-room/>.
- Cregan, A. (2007). Symbol grounding for the semantic web. In *Proceedings of the 4th European Semantic Web Conference, ESWC'07* (pp. 429–442).
- Cygniak, R., Wood, D., & Lanthaler, M. (Eds.). (2014). RDF 1.1 concepts and abstract syntax. Retrieved May 10, 2020, from <http://www.w3.org/TR/2014/REC-rdf11-concepts-20140225/>.
- Devlin, K. (2006). Situation theory and situation semantics. In *Handbook of the history of logic* (Vol. 7, pp. 601–664). Elsevier.
- Färber, M., Bartscherer, F., Menne, C., & Rettinger, A. (2018). Linked data quality of DBpedia, Freebase, OpenCyc, Wikidata, and YAGO. *Semantic Web*, 9(1), 77–129.
- Fensel, D., Simsek, U., Angele, K., Huaman, E., Kärle, E., Panasiuk, O., et al. (2020). *Knowledge graphs: Methodology, tools and selected use cases*. Springer.
- Franconi, E., Gutiérrez, C., Mosca, A., Pirrò, G., & Rosati, R. (2013). The logic of extensional RDFS. In *Proceedings of the 12th International Semantic Web Conference, ISWC'13* (pp. 101–116).
- Galton, A. (2008). Experience and history: Processes and their relation to events. *Journal of Logic and Computation*, 18, 323–340.
- Ganz, F., Barnaghi, P. M., & Carez, F. (2016). Automated semantic knowledge acquisition from sensor data. *IEEE Systems Journal*, 10(3), 1214–1225.
- Gärdenfors, P. (2000). *Conceptual spaces: The geometry of thought*. MIT Press.
- Goldblatt, R. (2006). Mathematical modal logic: A view of its evolution. In *Handbook of the history of logic* (Vol. 7, pp. 1–98). Elsevier.
- Grenon, P., & Smith, B. (2004). SNAP and SPAN: Towards dynamic spatial ontology. *Spatial Cognition & Computation*, 4(1), 69–104.
- Guha, R. V. (2015). Towards a model theory for distributed representations. In *Proceedings of the 2015 AAAI Spring Symposia* (pp. 22–26).
- Gupta, H., & van Oorschot, P. C. (2019). Onboarding and software update architecture for IoT devices. In *Proceedings of the 17th International Conference on Privacy, Security and Trust, PST'19* (pp. 1–11).
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42, 335–346.
- Harris, Z. S. (1954). Distributional Structure. *Word*, 10(2–3), 146–162.
- Hayes, P., & Patel-Schneider, P. (2014). RDF 1.1 semantics. Recommendation, W3C. <http://www.w3.org/TR/2014/REC-rdf11-mt-20140225/>. Latest version available at <http://www.w3.org/TR/rdf11-mt/>.
- Heckmann, D., Schwartz, T., Brandherm, B., Schmitz, M., & von Wilamowitz-Moellendorff, M. (2005). GUMO: The general user model ontology. In *Proceedings of the 10th International Conference on User Modeling, UM'05* (pp. 428–432).
- Heflin, J., & Hendler, J. A. (2000). Dynamic ontologies on the web. In *Proceedings of the 17th International Conference on Artificial Intelligence and 12th Conference on Innovative Applications of Artificial Intelligence* (pp. 443–449).
- Helbig, H., & Glockner, I. (2007). The role of intensional and extensional interpretation in semantic representations—The intensional and preextensional layers in MultiNet. In *Proceedings of the 4th International Workshop on Natural Language Processing and Cognitive Science, NLPSCS'07* (pp. 92–105).

- Henson, C. A., Sheth, A. P., & Thirunarayan, K. (2012). Semantic perception: Converting sensory observations to abstractions. *IEEE Internet Computing*, 16(2), 26–34.
- Hermann, K. M., Hill, F., Green, S., Wang, F., Faulkner, R., Soyer, H., Szepesvari, D., Czarnecki, W. M., Jaderberg, M., Teplyashin, D., Wainwright, M., Apps, C., Hassabis, D., & Blunsom, P. (2017). Grounded Language Learning in a Simulated 3D World. CoRR. [arXiv:1706.06551](https://arxiv.org/abs/1706.06551).
- Hermann, M., Pentek, T., & Otto, B. (2016). Design principles for Industrie 4.0 scenarios. In *Proceedings of the 49th Hawaii International Conference on System Sciences*, HICSS'16 (pp. 3928–3937).
- Hughes, G. E., Cresswell, M. J., & Cresswell, M. M. (1996). *A new introduction to modal logic*. Psychology Press.
- Janas, J. M., & Schwind, C. B. (1979). Extensional semantic networks: Their representation, application, and generation. In N. V. Findler (Ed.), *Associative networks* (pp. 267–302). Academic Press.
- Kokar, M. M., Matheus, C. J., & Baclawski, K. (2009). Ontology-based situation awareness. *Information Fusion*, 10(1), 83–98.
- Kripke, S. A. (1959). A completeness theorem in modal logic. *The Journal of Symbolic Logic*, 24(1), 1–14.
- Lanotte, R., & Merro, M. (2018). A semantic theory of the internet of things. *Information and Computation*, 259(1), 72–101.
- McCarthy, J. (1963). *Situations, actions, and causal laws*. Technical report, Stanford University, Department of Computer Science.
- McCarthy, J., & Hayes, P. J. (1969). Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer, & D. Michie (Eds.), *Machine intelligence* (Vol. 4, pp. 463–502). Edinburgh University Press.
- Menzel, C. (2017). Possible worlds. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University, Winter 2017 edition.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Proceedings of the 27th Annual Conference on Neural Information Processing Systems*, NIPS'13 (pp. 3111–3119).
- Montague, R. (1974). *Formal philosophy: Selected papers of Richard Montague*. Yale University Press.
- Nguyen, G. H., Lee, J. B., Rossi, R. A., Ahmed, N. K., Koh, E., & Kim, S. (2018). Continuous-time dynamic network embeddings. In *Proceedings of the 2018 World Wide Web Conference*, WWW'18 (pp. 969–976).
- Reich, W. (2010). Three problems of intersubjectivity-and one solution. *Sociological Theory*, 28(1), 40–63.
- Russell, S. J., & Norvig, P. (2010). *Artificial intelligence: A modern approach* (3rd ed.). Pearson Education.
- Serpanos, D., & Wolf, M. (2017). *Internet-of-things (IoT) systems: Architectures, algorithms, methodologies*. Springer.
- Singh, M. P. & Chopra, A. K. (2017). The internet of things and multiagent systems: Decentralized intelligence in distributed computing. In *Proceedings of the 37th IEEE International Conference on Distributed Computing Systems*, ICDCS'17 (pp. 1738–1747).
- Skirry, J. (2006). The mind-body problem. Retrieved May 05, 2020, from <https://www.iep.utm.edu/descmind/#H4>.
- Sowa, J. (2005). Theories, models, reasoning, language, and truth. Retrieved August 31, 2019, from <http://www.jfsowa.com/logic/theories.htm>.
- Staab, S., & Studer, R. (2010). *Handbook on ontologies*. Springer Science & Business Media.
- Stocker, M., Nikander, J., Huitu, H., Jalli, M., & Koistinen, M. (2016). Representing situational knowledge for disease outbreaks in agriculture. *Agrárinformatika/Journal of Agricultural Informatics*, 7(2), 29–39.

- Stocker, M., Rönkkö, M., & Kolehmainen, M. (2014). Situational knowledge representation for traffic observed by a pavement vibration sensor network. *IEEE Transactions on Intelligent Transportation Systems*, 15(4), 1441–1450.
- Studer, R., Benjamins, V. R., & Fensel, D. (1998). Knowledge engineering: Principles and methods. *Data & Knowledge Engineering*, 25(1–2), 161–197.
- Tarski, A. (1944). The semantic conception of truth: And the foundations of semantics. *Philosophy and Phenomenological Research*, 4(3), 341–376.
- Wang, Q., Mao, Z., Wang, B., & Guo, L. (2017). Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering*, 29(12), 2724–2743.
- Wang, Z., Zhang, J., Feng, J., & Chen, Z. (2014). Knowledge graph embedding by translating on hyperplanes. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, AAAI'14* (pp. 1112–1119).

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), 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 license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license 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.

