Towards Adaptive Open-World Modeling

Achim Kuwertz

Vision and Fusion Laboratory
Institute for Anthropomatics
Karlsruhe Institute of Technology (KIT), Germany
achim.kuwertz@kit.edu

Technical Report IES-2011-10

Abstract: In this technical report, extensions to the Object-Oriented World Model (OOWM) are proposed allowing for adaptive open-world modeling for artificial cognitive systems. In cognitive systems, a world model can serve as a central component for integrating, storing and disseminating information from an observed environment. Thus, a world model depicts an abstract, simplified representation of an observed real-world domain. For allowing high-level information processing on a semantic layer, representations of real-world entities, which are observed by sensors, can be semantically enriched by domain models. In general, a domain model contains only a fixed number of a priori defined concepts from a closed world. However, in many real-life applications, the considered environment is not closed. For coping with changing environments, a cognitive system must be equipped with an adaptive world model able to adjust to an observed open environment. Thus, this technical report discusses how the OOWM can be extended to facilitate semantically grounded open-world modeling.

1 Introduction

Information has become one of the most valuable assets in modern life. In fact, information is the building block from which knowledge is built, concerning almost any area of today's life like politics, business or even daily life. Information arises whenever data is interpreted in a meaningful way, be it the data of a radio signal in communications, scientific data observed in experiments or the data retrieved from online resources. With data nowadays being collected in growing quantities and the increasing presence of network enabled devices, information seems to be ubiquitously accessible. Thus, the unavailability of information in many cases no longer seem to pose a problem. Nevertheless, it still can be tedious to find, retrieve and integrate specific information from available sources in order to, e.g., establish a task-related situational overview or answer dedicated questions. In fact, the availability of a flood of information even worsens this problem. In this situation, modern technologies can be applied to guide and aid us in the task of processing and managing available information, supporting us just with relevant information or even information outlined on the right level of abstraction for our needs.

Another area of modern life where innovation and technological progress has led to the possibility of aiding and supporting humans by technical solutions is the field of assisted living and working. In this area, mechatronic systems autonomously manipulate their environment in order to fulfill supportive work or even specific tasks like heath care. In order to succeed at such tasks, autonomous technical systems must be able to sense, i.e., to acquire information about the environment they are operating in in order to determine the current environment state. Perception and adequate interpretation of a variable and possibly unknown environment are important prerequisites for a successful operation of these systems. Thus, autonomous systems also need the ability to acquire, process and integrate relevant information for establishing a task-related situational picture.

1.1 Artificial Cognitive Systems

In conclusion, one of the most relevant problems of today's profound decision taking no longer is the lack of information, but rather the tasks revolving around the separation of relevant from non-relevant information, the integration of retrieved information to form a consistent knowledge, the drawing of conclusions based on such knowledge and the presentation of desired derived information on the right level of abstraction. Hence, there is a need for technical systems being able to (semi-) autonomously fulfill these cognitive tasks. In this report, systems able to cope with such tasks will be subsumed under the notion of *artificial cognitive systems*.

In order to fulfill one of their central tasks, evaluating a sensed operational environment, artificial cognitive systems need to be equipped with several subsystems:

- a *sensing* input subsystem possibly consisting of multiple sensors of heterogeneous types and sensing modalities.

- an *information integration and storing* subsystem that consistently fuses all available sensor data and integrates it with a priori information into a comprehensive knowledge base.

- a *reasoning* subsystem able to process and answer information requests by interpreting acquired knowledge (e.g., extract current environment state).

- an output subsystem providing a channel for *relaying and displaying* the results of information requests on different levels of abstraction.

Furthermore, artificial cognitive systems can possess more functional subsystems like actuators for environment manipulation and locomotion, or subsystems for planning and scheduling such interaction. In addition to these functional aspects, artificial cognitive systems need to be provided with semantic information about their operational domain in order to evaluate their sensor observations in a meaningful manner. Such information can be given as an a priori *domain model*.

Areas of application for artificial cognitive systems include assisted living and working, disaster management, security and safety in public places, remote sensing for environment protection, maritime and land border surveillance as well as reconnaissance. In all these areas, data and information available from heterogeneous sources have to be integrated to support a situational awareness for decision making by establishing an overview of the current situation at hand.

1.2 Object-Oriented World Modeling

The task of integrating different information into a comprehensive semantical representation can be solved by a world modeling system as depicted in Fig. 1.1. The OOWM [GHB08, BKFB12], developed at Vision and Fusion Laboratory in collaboration with Fraunhofer IOSB, is an example of such a system. Its main purpose is to provide a consistent and integrated view on the current state of an observed environment. It thereby acts as an information hub for sensory and cognitive processes, allowing all subsystems to access the information required to perform their tasks. The OOWM possesses the following features ([GHB08, EGB08]):

- *object-oriented* information representation – observed real-world entities are represented as objects with attributes and relations.

- *semantic* background knowledge – observed entities are associated with semantic concept descriptions stored in a domain model.

- *bayesian* information processing – information stored in the world model is treated as inherently uncertain and characterized by probability distributions interpreted as degrees of belief, which allows for integrating new information in a uniform way by applying fusion methods. [HGPLB10].

Figure 1.1: World modeling: a given domain, the world of interest, is observed by heterogeneous sensor systems. Sensor data and extracted information are passed on to a world modeling system in order to be integrated, stored and serve as a basis for evaluating the current state of the world of interest.

- *dynamic* environment description – stored information is not static but can both be dynamically updated by new observations and age in time.
- level-of-detail *abstraction* – information provided to subsystems can be represented on a task-relevant hierarchical level of abstraction [HGPLB10].

1.3 Adaptive World Modeling

The ability of a world model to represent and semantically evaluate the current state of an observed domain strongly depends on the presence of a semantic domain model. In general, domain models are created both manually by human experts and prior to the operations of a cognitive system. Therefore, the semantic expressiveness of a world model is limited by the set of a priori modeled concepts. As a consequence, only a closed world segment can be represented. To integrate unforeseen events and entities, the world model has to be able to dynamically adjust its prior knowledge to the current state of the observed domain. This ability to adjust and extend existing knowledge is at the heart of *adaptive world modeling*.

A main part of adaptive open-world modeling is concerned with extending knowledge by learning new concept definitions. Concept learning for artificial cognitive systems comprises of several different subtasks. These tasks include deciding when and which concept should be learned at all, acquiring a perceptual representation and a symbolic name for a learned concept as well as defining its meaning by learning unobservable semantic attributes. In order to enable concept learning for world models, this technical report proposes to extent the OOWM. For this purpose, a formal non-shallow knowledge representation will be introduced and an interoperable top-level ontology for OOWM knowledge will be defined.

This technical report is structured as follows. Sec. 2 gives a brief overview of the object-oriented world modeling system and its notions. Sec. 3 proposes extensions to the object-oriented world model for allowing adaptive world modeling. Sec. 4 then takes a closer look at semantics, meaning and symbol grounding for different artificial cognitive systems. In Sec. 5, various concept learning approaches relevant for grounding adaptive world modeling are examined. The report concludes with an outlook on planned future work.

2 Notions of Object-Oriented World Modeling

The term *world modeling* refers to generating an abstract representation of a real-world domain, as depicted in 1.1. The resulting model contains simplified, but formal representations of real-world entities. In world modeling different approaches for taking into account time-dependent features can lead to different types of models. On the one hand, static models describe the static, non time-dependent properties of a modeled domain, e.g. facts, structures or statements that hold true regardless of the point of time they are considered. Such models are often called domain models. On the other hand, in contrast to static modeling a dynamic model tries to capture the current state of an observed domain. Based on sensor measurements of time-dependent properties, dynamic world models focus on providing a snapshot of features like inherent states, positions or constellations of entities. In order to represent a domain in a comprehensive way, both state information and static knowledge are necessary. Thus, the object-oriented world model presented in the following employs a domain model and a dynamic model in combination.

2.1 System Overview and Notions

Within object-oriented world modeling, the relevant part of the real world is called the *world of interest* [FHB10]. The world of interest is a domain-specific or task-specific view on a spatio-temporal segment of the real world. Conceptually, the world of interest consists of various different *entities* like objects, persons or abstracta. This view of the world of interest, taking into account not only physical objects but also abstract concepts, constitutes a human-related perspective to arti-

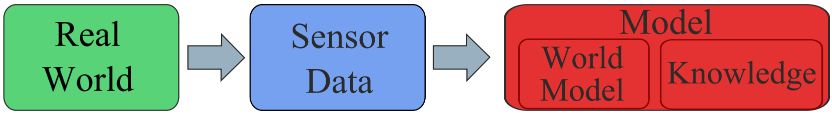

Figure 2.1: High-level overview and notions of object-oriented world modeling.

ficial cognition. This ensures a compatible, consistent view and a mutual comprehension for human-related artificial cognitive systems and their human operators.

For acquiring dynamic information, observable *attributes* of individual entities like their extent, color or position are monitored by sensors. The sensor data, representing a view or projection of the world of interest to sensing modalities, serves as input for the OOWM. The OOWM consists of two modeling subsystems [BKFB12]: the dynamic *World Model* representing the current state of an observed world of interest, and the *Knowledge* subsystem representing domain knowledge. An overview of object-oriented world modeling is depicted in Figure 2.1.

2.1.1 Knowledge

The *Knowledge* subsystem provides a static model of the world of interest. As an important part, the semantics of the considered domain have to be captured within the model. In object-oriented world modeling, the semantics of entities are formally represented by *entity concepts*, which describe the observable and unobservable attributes of an entity and its defining relations. Attributes comprise time-varying state attributes (e.g., a position) and immutable property attributes (e.g., a color). Relations can either describe necessary relationships, which have to hold any time and for all instances of a concept, or possible relationships which, e.g., can hold only for certain instances or distinct periods of time. Entity concepts can be modeled prior to system operations by domain experts. Thus, they serve as a priori information. Yet, OOWM knowledge does not just consist of prior knowledge, as it is designed to support adaptive open-world modeling and can be extended by newly learned concept definitions.

Using a semantic domain model enables an artificial cognitive system to share a meaningful set of symbols with its environment. Based on prior knowledge it is not only possible to classify a real-world entity given observations of its features, but also to derive unobservable features or to interpret symbolic information requests, e.g. being phrased in natural language.

2.1.2 World Model

The *World Model* in the OOWM is responsible for processing and storing the dynamic environment data given by sensor observations. Sensor observations are associated with a *representative* of the real-world entity they originate from and are stored as a container set of attribute values [KBS+10], as exemplary illustrated in Fig. 2.2. Observed attribute values are represented by probability distributions characterizing the measurement value as well as its uncertainty as a Degree-of-Belief (*DoB*). For integrating heterogeneous sensor observations, techniques for data and information fusion are employed, supported by techniques for information management (like an aging mechanism, consistency checks, etc.) [BKFB12]. Using a set-based entity representation allows to represent unknown entities as well, as long as the observed attributes are known. Furthermore, blank objects allow to represent just the presence or rather existence of something. If relations between different objects are observed, the current state of a world of interest is represented as a semantic network in the World Model. In addition, past environment states are represented in a history of recent state information. The World Model thus is time-dependent and dynamic.

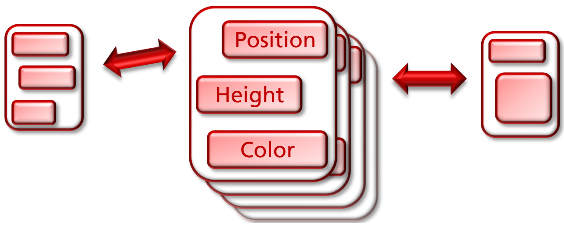

Figure 2.2: The World Model as a semantic network with history and past states.

2.2 Areas of Application

The OOWM is designed to be an information integration and storage facility of general purpose being applicable, as a part of a cognitive system, in various areas of application. Up to now, the OOWM has been successfully applied in several domains including assisted living [KBS+10], indoor security [EGB08], real estate surveillance or maritime border surveillance [FB10]. In many cases though, only a shallow model was employed to represent semantic domain knowledge. In order to allow for a semantic open-world modeling, a formal structure for representing non-shallow domain knowledge must be defined. This extension to object-oriented world modeling is introduced in the following section.

3 Extending the Object-Oriented World Model

In adaptive world modeling knowledge can be modified by acquiring new concept definitions during the time of system operations. As stated in [KHB10], the task of dynamic concept acquisition constitutes an extension to the classical concept learning problem [Mit97]. Classical concept learning is aimed at learning concepts, represented as boolean functions over sets of attributes, from given labeled training data in a supervised manner. As a kind of inductive learning problem, concept learning requires a formal structure in which concepts can be represented, i.e., a concept space. The process of learning new concepts can be regarded as a search within this space for the most suited representation of a considered concept given a set of training data [Mit97]. The structure for concept representation constitutes the inductive bias necessary for learning and generalization.

3.1 A Meta-Model for Entity Concepts

In object-oriented world modeling, up to now no explicit structure for entity concepts exists. To allow for sophisticated concept learning in adaptive world modeling, a meta-model formally defining the space of entity concepts will be introduced in the following. Extending our ideas presented in [BKFB12], a formal structure for the most basic entity concepts in OOWM knowledge will be defined. In summary, entity concepts are governed by intra- and inter-class relationships. Intra-class relationships are modeled by attributes (e.g., the height of an entity), which are instances of attribute concepts (e.g., a height attribute being a non-negative real-valued variable of a certain length unit). Inter-class relationships are modeled by relations (e.g., one entity being part of another), which are instances of relation concepts. Relation concepts again are characterized by attributes. Fig. 3.1 shows an overview of the meta-model for entity concepts in OOWM knowledge described subsequently.

3.1.1 Individual Concepts

The essential entity concepts in object-oriented world modeling are *individual concepts*, which model the common features of entities like things, persons, roles, functions, qualities, etc. Individual concepts are defined by

- a unique name,
- a parent individual concept (given as a reference),

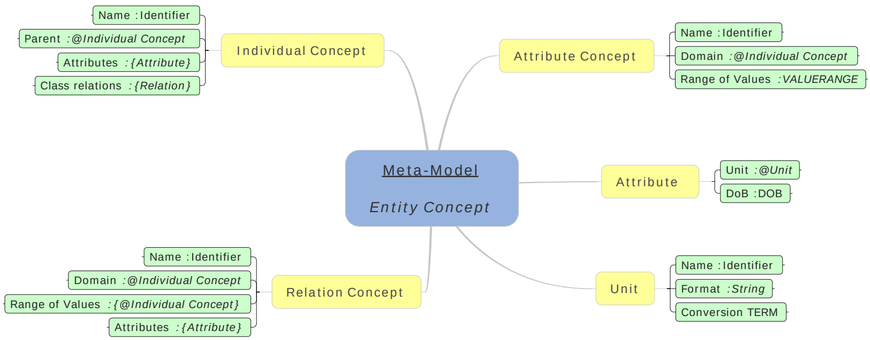

Figure 3.1: Meta-model structure for entity concepts in OOWM knowledge. The displayed concept definitions specify the syntax for individual concepts, relation concepts, attribute concepts, attributes and units.

- a set of attributes, and
- a set of class relations.

Attributes and class relations (i.e., relations necessarily holding for all instances of a concept) serve to semantically define specific features of individual concepts. The reference to a parent concept allows to establish a concept hierarchy.

3.1.2 Attribute Concepts

For characterizing an individual entity, attributes describing the entity features are instantiated from *attribute concepts*. An attribute concept is specified by

- a unique name,
- an attribute domain (given as a reference to an individual concept), and
- a range of values defining admissible attribute values.

By defining a range of values, the type of an attribute concept is determined as well (e.g., nominal vs. cardinal attributes). This allows the use of specialized fusion and learning methods adjusted to attributes for this concept. The reference to an individual concept serves as semantic definition for this attribute.

3.1.3 Attributes and Units

When instantiating an *attribute* from a given attribute concept (e.g., a height), an attribute additionally needs to define

- its unit (given as a reference to a unit concept), and
- a DoB probability distribution (characterizing the expected attribute values).

Not specifying units within attribute concepts, but later within attributes allows for multiply instantiating a single attribute concept with attributes of different units. Following the presented structure, the height of a cup can for example be described as being given in cm and being equally distributed from 9 to 10 (cm). In order to enable this structure, *unit concepts* are defined as consisting of

- a unique name,
- a syntax string, and
- a conversion term (allowing to convert related units, e.g. cm to m).

The syntax string of a unit defines the format of values given in this unit. Representing the data format of values then allows to define a conversion term, which can be used for automatic conversion of given values to a standard unit.

3.1.4 Relation Concepts

An individual concept is not only characterized by its attributes (defining intra-class relationships), but also by relationships to other entity concepts. As mentioned earlier, there are two types of relations within world modeling: class relations holding for all instances of an entity concept and thus being part of its definition, and instance relations that can hold just for individual instances (e.g., being situated near to something). In general, a *relation concept* consists of

- a unique name,
- a domain for which this relation can be applied (given as a reference to an individual concept),
- a range of values for this relation (given as set of references to individual concepts), and
- a set of attributes.

When specifying an individual concept as domain or as part of the range of values for a relation concept, relation instances can be applied to instances of the given individual concept as well as all its children in the concept taxonomy.

3.2 Top-Level Knowledge Ontology

Semantic concept learning, besides a syntactic metastructure for OOWM knowledge, also requires an initial semantic structure for organizing its entity concepts. When knowledge is represented by ontologies, abstract high-level ontologies called upper ontologies or top-level ontologies can be employed to introduce a semantic structure. A *top-level ontology* thereby defines high-level entity concepts, i.e., entity concepts that occur independent of and across many application domains, in an abstract and extendable way. Examples of top-level ontologies include, e.g., DOLCE [GGM+02] or WordNet [MBF+90]. An important purpose of top-level ontologies is to enable semantic interoperability between different application-specific domain ontologies. Thus, a top-level ontology structure seems to be good anchor for a semantic and interoperable extension of OOWM knowledge as required for adaptive open-world modeling.

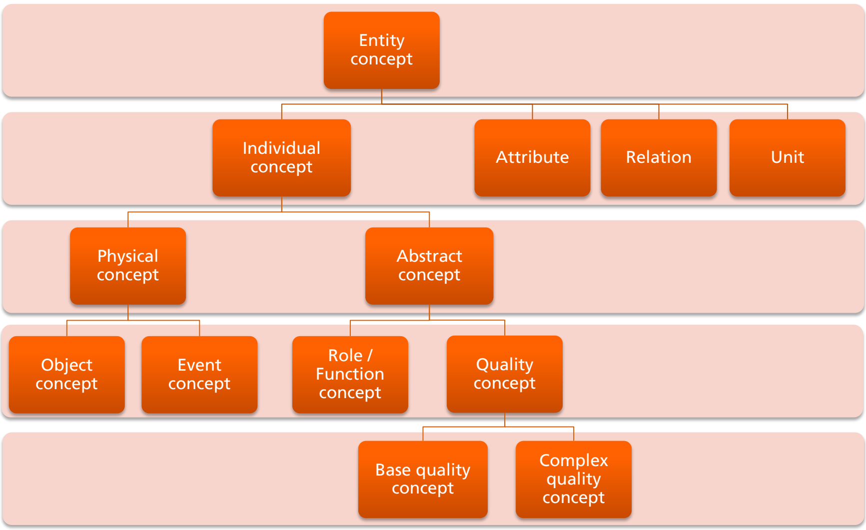

Figure 3.2: Top-level ontology for structuring entity concepts in knowledge.

As a semantic structure for OOWM knowledge, the approach depicted in Fig. 3.2 was chosen. This approach is in some parts based on the principles used for designing the DOLCE ontology [GGM+02]. As can be seen in Fig 3.2, entity concepts are semantically subdivided into four *basic concept categories*, namely individual concepts, attribute concepts, relation concepts and unit concepts, reflecting the structure used within the meta-model. At the moment, only a conception for further partitioning individual concepts has been elaborated. Individual concepts are semantically subdivided into concepts for physically existing entities and concepts for abstract entities. *Physical concepts* are further subdivided into object concepts for physical entities existing in space and event concepts for physical entities existing in time. Object concepts thereby include concepts for persons and things. *Abstract concepts* are divided into role-function concepts on the one hand, describing the roles a person can embody as well as the functions a thing can fulfill. On the other hand, quality concepts describe the features that real-world entities can possess. In future work, the structure of the remaining basic concepts as well as more level-of-detail for the presented concepts has to be elaborated.

3.3 Towards Adaptive Object-Oriented World Modeling

Figure 3.3 gives an overview of system components and their interactions relevant for adaptive object-oriented world modeling. As in the present OOWM, information within knowledge is used to classify real-world entities based on sensor observations of their attribute values [BKFB12]. When classified, additional information, for example on unobservable attributes, can be derived from knowledge. During operations, the system is likely to encounter real-world entities that have not been a priori modeled in knowledge and thus cannot be uniquely mapped to an entity concept. Being represented as set of observed attribute values, such entities can then be passed to a concept learning subcomponent responsible for adaptively extending knowledge. Besides being driven by sensor observations, concept learning can also be necessary if the world model is queried for instances of concepts not yet represented in knowledge.

The concept learning component is bound by the meta-model structure in respect of how information is learned and stored in knowledge and by the top-level ontology in respect of what information can be learned. The meta-model describes the syntax of concept definitions and thus formalizes what a system itself “knows” about the information stored in knowledge. This means that all features described by the metal-model can be interpreted in a meaningful way by the system. The top-level ontology provides a semantic foundation for concept learning by abstractly defining the basic partition of concept space as well as defining the semantics of measured and learned attributes. For managing the information modeled in meta-model and top-level ontology or represented in knowledge, a model editor component seems to be a valuable supplement. This editor in addition could serve as a tool to display and manually verify newly learned concept definitions.

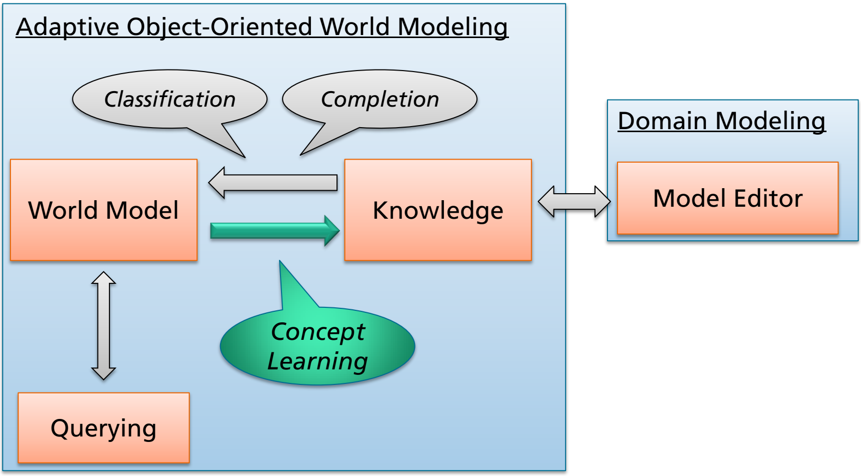

Figure 3.3: Overview of adaptive object-oriented world modeling. In addition to the interaction of world model and knowledge for classifying observed entities and completing the information represented for this entities, dynamic information from world model is used as input for extending knowledge by concept learning.

3.4 Concept Learning Example

The conceptions of adaptive world modeling shall be illustrated by a brief high-level example. Suppose a number of entity concepts is given, formally describing tangible real-world objects as defined in the meta-model. In fact, a more specific structure defining a kind of prototype concept for tangible objects could be imagined (for example, a common parent node in concept taxonomy). The attributes of this prototype concept (e.g., its extent, shape, mass, position, color, etc.) then have to be semantically defined within an extended top-level ontology, for example by defining related attribute concepts and synonymic symbolic names, and within an extended meta-model by defining the data type properties for each attribute.

Now, suppose a yet unmodeled entity is observed and several of its attributes get measured. First, each sensor measurement has to be associated with an attribute concept. This for example is possible, if all sensors formally describe their sensing modalities a priori as attribute concepts. Then, one can try to learn the distribution of values for each attribute of the unmodeled concept. Knowing the data type of attributes, different methods for learning this distribution can be employed, for example to learn nominally, ordinally and cardinally scaled attributes and represent them as Gaussian mixture models. Having an initial estimate of the distribution of each attribute, one then can try to separate attributes relevant for concept definition from non-relevant attributes, e.g. by evaluating the entropy of each attribute distribution: The lesser an entropy, the more relevant an attribute is since it then seems to be a good indicator for separating concept instances from non-instances based on observed values.

If not enough measurement data is given for an observable attribute, or when trying to learn unobservable attributes, more data can be acquired by specific request e.g. to a human operator or by searching information on the World Wide Web. Acquiring information in such ways on the one hand requires a symbolic name or description for the concept to be learned. On the other hand, semantic information on relevant attributes is necessary to find data which can serve as training data for learning attribute distributions. Which attributes have to be learned could thereby, for example, be determined by a high-level service in order to performing its tasks.

4 Meaning and Semantics in World Modeling

An important aspect of world modeling for artificial cognitive systems is to represent an observed environment in a meaningful way. A cognitive system shall not only store its observations according to world modeling data structures, but should also be able to understand its environment and perhaps interact with it in a useful manner. Ideally, an artificial cognitive system should be able to autonomously explore its environment and "get a feeling" for the entities it encounters. Getting a feeling for entities means to establish in internal representation, a model, of the perceptual and action-related features which real-world objects do possess. Furthermore, a cognitive system should be able to learn abstract concepts from textual descriptions, e.g. explanations given by humans. Moreover, again ideally, a cognitive system should end up with the same idea of its environment as a human would have. Thus, artificial cognitive systems require human-related world modeling subsystems, that is a world model at least capable of establishing an environment representation being compatible to a human's conception.

4.1 Semantics in World Modeling

The term meaning can be interpreted in a lot of different ways according to in which scientific discipline or even philosophical movement it is used. One way to interpret meaning is to say that it subsumes all information related to finding a *referent*, i.e., a real-world entity, for a given symbol, e.g., a word. Within world modeling, this kind of meaning is relevant for identifying entities in the real world which are instances of a concept given by a symbolic description, for example by natural language. Thus, following the referent interpretation, the meaning of a concept is basically given by its observable attributes and its defining relations deducible from this attributes.

Yet, this interpretation disregards the remaining attributes which, for example, describe unobservable features. When looking for a cup, we not necessarily use the feature of a cup being a vessel to carry liquids for discriminating real-world objects. Nevertheless, being a vessel for liquids is part of the semantics of a concept for cups. Thus, another way to interpret the meaning of a concept is to subsume all its features, attributes and relations, especially features that describe functions or roles. In some way, these features can be seen as higher-lever semantics.

For artificial cognitive systems, meaning can be defined within the knowledge of a world modeling subsystem. Concepts contained in knowledge characterize the mapping from real-world entities to representatives in the world model in a homomorphic way [Syd82]. The main purpose of this mapping for cognitive systems is to share a meaningful set of symbols with their environment: on the one hand this mapping allows for classifying and identifying observed entities based on their attribute values, thus fulfilling the referent interpretation of meaning. On the other hand concepts allow to interoperably exchange meaningful information, e.g., when receiving requests from humans. Based on the mapping of representatives to concepts, human requests and sensor data can be interrelated and, furthermore, represented entities can be enhanced with unobservable semantic features.

4.2 Symbol Grounding

After having seen that meaning in world modeling is useful for identifying and classifying real-world entities as well as processing symbolic requests, the next step is to take a look at how meaning can be acquired by cognitive systems. This is closely related to the symbol grounding problem in cognition as stated in [Har90]. The question at hand in symbol grounding is how a "semantic interpretation" of symbols, e.g., attribute and relation concepts in the case of world modeling, can "be made intrinsic to the system rather than just" existing as "meanings in our heads" [Har90]. Recognizing and addressing this problem can be seen as a first step to bridging the semantic gap. If one, for example, formally models a given set of facts within an ontology, one can reach semantic interoperability in a sense that different humans can interpret the symbols contained in the ontology in a consistent way. Yet, one does not necessarily enable a system employing this ontology to semantically understand its symbols. This difference can be made clear by imagining an ontology in which the symbolic concept names have been exchanged for numbered variables. Since this is the perspective a system has onto an ontology, the question arises if this ontology still does define meaning?

To solve this problem, different solutions are imaginable. For example, it is possible to try to formalize a set of facts or a whole domain in such a rigid way (e.g., by using a formal language) that, from a human perspective, by exchanging symbolic names no information would be lost. But this approach seems to be rather complex and time-consuming. Another solution is to ground the occurring symbols as proposed in [Har90]. Basically, a bottom-up grounding of semantic concepts within non-symbolic representations is proposed, where concepts are being hierarchically constructed from grounded symbols. The basic non-symbolic representations are given by sensor observations of entity features. How this approach, envisioned as a psychological model of human cognition, could be transferred to world modeling in artificial cognitive systems shall be examined in the following section.

4.3 Symbol Grounding for Artificial Cognitive Systems

The notion of artificial cognitive systems presented in this report subsumes autonomous systems and information systems. Prior to considering a symbol ground for these systems, their similarities and differences shall be pointed out.

4.3.1 Types of Cognitive Systems

A *cognitive autonomous system* is a system which is able to autonomously, i.e., without human interaction, make decisions, based on observations of its environment on the one hand and internally stored knowledge on the other hand, and then, as implied by these decisions, accordingly interact with its environment. Thus, an autonomous system can be characterized by its subsystems fulfilling the tasks of

- environment perception,
- data and information processing,
- decision making (goal-directed), and
- environment interaction (e.g., manipulation, mobility, etc.).

The two tasks of actively perceiving and interacting with its environment are fundamental for autonomy: on the one hand, a autonomous system has to be able to interpret signals from its environment in order to capture the current state of the environment and react to state changes. On the other hand, autonomous systems are constructed to serve a specific purpose and therefore need ways to interact with their observed environment. To relate perceived environment observations to interactions, the cognitive tasks of information processing and decision making are necessary. Examples of autonomous systems situated in different environments range from humanoid robots (perceiving and manipulating their environments in a human-like sense) over autonomous service robots (e.g., for performing some specialized tasks like room cleaning) over feedback control systems to software agents (perceiving and interacting for example within a simulated or abstracted environment as avatars, like trading agents for stock exchange). Some types of autonomous systems are designed to perform more complex tasks in an open environment, like interacting with humans. These systems have to be equipped with sophisticated cognitive processing abilities, able to understand a part of their observed environment in a way compatible to human comprehension.

Apart from systems able to interact with their environment in a physical (relevant or simulated) way, there are artificial cognitive systems primarily concerned with processing, managing and displaying high-level information. These *cognitive information systems* like autonomous systems consist of subsystems responsible for the cognitive tasks of

- information processing and
- decision making.

Yet unlike autonomous systems, cognitive information systems do not directly interact with an observed environment but communicate their processing results, like conclusion or alerts, to a human operator. Also these systems do not directly observe an environment but are presented with preprocessed data and information as input. Examples of such systems include information management systems, decision support system and situation awareness support systems.

4.3.2 Grounding Attribute Concepts

Since in world modeling sensor observations serve as input for attribute values, concept grounding has to start with attribute concepts. Attribute concepts in world modeling can be divided into sensomotoric and semantic attributes. *Sensomotoric* attributes, being either perceptual attributes or effector attributes, are especially relevant for cognitive autonomous systems as these systems directly observe and interact with their environment. Perceptual attributes (e.g., the height, color or speed of an entity) are grounded within the sensor components providing value information for these attributes to the world model. Effector attributes (like weight, shape, etc.) are grounded within the actuators used to manipulated real-world entities based on the represented values for these attributes [Ken98a]. As can be seen, there is no sharp boundary between perceptual and effector attributes, and which attribute is grounded by which sensomotoric component of an autonomous system in general depends upon the task currently to be executed. By relying on this sensomotoric grounding, cognitive autonomous systems are in principle able to autonomously acquire symbolic descriptions of concepts for observable entities grounded in sensor projections of the real-world. They are further able to ground action-related attributes of physical concepts in their interaction and manipulation of the respective real-world objects.

Furthermore, there are attributes in knowledge that can neither be associated with a sensor or an actuator component. Examples are attributes concerning roles and functions of entities and information like if and how an entity can be moved or where it is usually resting. These attributes constitute *semantic* attributes as they either result from acquired knowledge or long-term experience. Semantic attributes are unobservable and include physical properties, behavior and capabilities of real-world entities. They are the kind of attributes usually dealt with in cognitive information systems. As they originally may have been, at least partially, provided by a sensor system, they get preprocessed prior to being input to the information system. Thus these attributes cannot be grounded in sensor projections. Yet these semantic attributes also play an important role in world modeling for autonomous systems. Here they describe higher-level information like objection functions, e.g. a cup serving as a drinking vessel. In order to ground semantic attributes, symbolic descriptions can be used, e.g. given as text. In that way, semantic attributes can either be related to attributes and entity concepts grounded in sensor projections, as necessary for autonomous systems, or they can be defined in the context of semantic networks, as might be sufficient for cognitive information system that do not directly interact with their environment.

5 Types of Concept Learning

In order to enable adaptive open-world modeling for cognitive systems, techniques have to be employed for extending the existing knowledge of a world model with new entity concepts. This acquisition of concepts is often denoted as concept learning. The term concept learning is widely used in different scientific disciplines. In cognitive sciences it for example describes human concept learning and its computational implementations. In computer sciences, it describes artificial computational concept learning as part of or extension to machine learning and AI. In computational concept learning, there are again several different terms used to describe methods for learning new concepts, e.g.

- dynamic category building (e.g. [BB09]),
- concept acquisition (e.g. [Ken98b]),
- autonomous concept generation (e.g. [Ken98a]), or
- category learning (e.g. [DLST09], [TA08]),

which all focus on different aspects of concept learning. Furthermore, there is classical concept learning as e.g. stated in [Mit97]. Besides using various terms to denote concept learning, there are several subtypes of concept learning concerned with learning different kinds of models from observed data. These models differ in respect of what attributes are to be learned for representing an entity, and thus operate on different levels of semantics. The subtypes of concept learning include

- visual concept learning,
- symbolic concept learning,
- object learning, or
- ontology learning.

For adaptive open-world modeling, several of these concept learning subtypes have to be combined to learn about perceptual as well as symbolic and semantic features of observed entities. In this section, a short overview of some of these concept learning subtypes shall be given.

5.1 Visual Concept Detection and Learning

Visual concept learning or visual categorization has been a topic of research for decades [DLST09]. In visual concept learning, one tries to learn, detect, recognize and classify objects in video and image data by relying object shapes or even classifying whole situations depicted in imagery. The goal is to recognize already learned concepts in new imagery data and continuously update their representations. In this kind of learning, concept models emphasize syntax, not semantics, but the methods are for example applicable for automating image exploitation.

5.1.1 Cross-Modal Visual Concept Learning Based on Collateral Data

In [SLS05b] video data is tagged with concepts contained or depicted in single frames - for the purpose of semantic information retrieval and indexing. Concepts, hereby understood as keyword tags modeling image features, are detected within video data by classifiers trained in a supervised manner. For training, the approach proposed in [SLS05b] does not require tedious manual tagging of extensive training samples. Instead, an automatic labeling algorithm is employed, based on recognizing speech from correlated audio data, e.g. contained on the sound track of a video. This multimodal approach thus allows for learning visual concepts by abstracting their defining features from relevant image samples indicated by collateral audio data. Furthermore, the learned concepts are coincidently labeled with semantically interpretable names. In [SLS05a], the authors pick up on this fact and try to retrieve visual concepts solemnly based on accompanying speech information. The speech can be analyzed for keywords being related to a concept in terms of keyword expansion based on ontologies like WordNet.

5.1.2 Web-Based Visual Concept Learning

The World Wide Web (WWW) can be seen as another data source of collateral information for image data. This fact is for example exploited in [UBB10], where weakly labeled web videos are used in an unsupervised training for video indexing. A weakly label video thereby denotes a web video, for example retrieved from a video platform like YouTube, which has been manually tagged by a human user. Classifiers resulting from training with these videos shall then be able to detect objects, locations and even activities in videos.

In a similar fashion, as [UBB10] states, it is possible to use Google Search to find image examples for text-based concept descriptions. Google Search thereby uses the correlation between the textual descriptions of or located near by pictures in websites. Based on Google labels, there exists methods to create more robust labeling, e.g., by clustering the results and using the most promising cluster for concept detecting. Also, a manually chosen subset of Google Search results or a subset provided by text and meta-data analysis could be used as training data for training more robust concept classifiers.

5.1.3 Statistical Data-Driven Visual Concept Learning

Another approach to visual concept learning is taken by [DLST09]. In this approach, a special focus is laid on concept representation in order to achieve a high generalization performance for the learned models. Visual concept models thereby are composed of different shape primitives modeling the contours of objects. They are built in a *hierarchical* data-driven manner by an unsupervised *statistical* bottom-up learning algorithm. Combining statistical learning and a hierarchical representation allows for selecting only relevant shape primitives, based on given image data, on a lower level, which then serve as building parts for objects contours on a higher level. On the highest level, a few hand-picked training examples are used for learning the specific category models, which then can be assigned a label (e.g., ”cups”). The presented approach reaches a promising generalization performance for correctly classifying object shapes in previously unseen images.

5.2 Ontology Learning

Since visual concept learning focuses on the perceptual attributes of entity concepts, ontology learning seems to be a suitable supplement for adding semantics to perceptual concepts. Therefore, in future work methods for ontology learning, perhaps based on acquiring symbolic knowledge from online resources like the WWW or existing ontologies, have to examined in more detail.

6 Conclusion

This technical report proposes a conception to adaptive open-world modeling for artificial cognitive systems. Several important aspects concerning semantic modeling and concept learning have been discussed, including the structure of concept representation, the semantic grounding of knowledge and the different existing approaches to concept learning. All these aspects thereby play an important role when designing an adaptive system for open-world modeling.

Future work should further examine many of these areas and elaborate them in more detail. In total, adaptive world modeling in general and concept learning for world models in particular can be regarded as an information management problem. Future work to this problem includes aspects of:

- Defining a process or procedure able to determine when concept learning should be initiated or, respectively, what concepts should be learned in order to be of use for a world modeling system.

- Defining a framework able to perform inductive probabilistic concept learning for acquiring grounded generalized representations of observed entities.

- Defining a framework responsible for learning semantic attributes from symbolic descriptions for supplementing perceptually acquired entity concepts.

- Defining a procedure for managing ontological concept taxonomies and integrating newly learned entity concepts into such a taxonomy.

Bibliography

[BB09] Joscha Bach and Joscha Bach. The MicroPsi architecture. In *Principles of Synthetic Intelligence PSI: An Architecture of Motivated Cognition*, volume 1, pages 233–265. Oxford University Press, Oxford, England, 1 edition, November 2009.

[BKFB12] Andrey Belkin, Achim Kuwertz, Yvonne Fischer, and Jürgen Beyerer. World Modeling for Autonomous Systems. In *Information System (to appear)*, number 1. InTech, 2012.

[DLST09] Sven J. Dickinson, Ales Leonardis, Bernt Schiele, and Michael J. Tarr. Learning Hierarchical Compositional Representations of Object Structure. In *Object Categorization*. Cambridge University Press, 2009.

[EGB08] Thomas Emter, Ioana Gheta, and Jürgen Beyerer. Object Oriented Environment Model for Video Surveillance Systems. pages 315–320, 2008.

[FB10] Yvonne Fischer and Alexander Bauer. Object-Oriented Sensor Data Fusion for Wide Maritime Surveillance. In *Proceedings of 2nd NURC International WaterSide Security Conference*, Marina di Carrara, Italien, November 2010.

[FHB10] Yvonne Fischer, Marco Huber, and Jürgen Beyerer. World modeling for advanced surveillance systems. Technical Report IES-2010-02, Vision and Fusion Laboratory, Institute for Anthropomatics, KIT, 2010.

[GGM+02] Aldo Gangemi, Nicola Guarino, Claudio Masolo, Alessandro Oltramari, and Luc Schneider. Sweetening ontologies with dolce. In *Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management. Ontologies and the Semantic Web*, EKAW '02, pages 166–181, London, UK, 2002. Springer-Verlag.

[GHB08] Ioana Gheta, Michael Heizmann, and Jürgen Beyerer. Object Oriented Environment Model for Autonomous Systems. In Henrik Boström, Ronnie Johansson, and Joeri van Laere, editors, *Proceedings of the second Skövde Workshop on Information Fusion Topics*, pages 9–12. Skövde Studies in Informatics, November 2008.

[Har90] Stevan Harnad. The Symbol Grounding Problem. *Physica D: Nonlinear Phenomena*, 42(42):335–346, June 1990.

[HGPLB10] Michael Heizmann, Ioana Gheta, Fernando Puente Leon, and Jürgen Beyerer. Information fusion for environment exploration. In *Reports on industrial information technology*, volume 12, pages 147–166. KIT Scientific Publishing, Karlsruhe, 2010.

[KBS+10] Benjamin Kühn, Andrey Belkin, Alexej Swerdlow, Timo Machmer, Jürgen Beyerer, Jürgen Beyerer, and Kristian Kroschel. Knowledge-Driven Opto-Acoustic scene analysis based on an object- oriented world modeling approach for humanoid robots. In *Robotics (ISR), 2010 41st International Symposium on and 2010 6th German Conference on Robotics (ROBOTIK)*, pages 1–8. VDE, June 2010.

[Ken98a] Catriona M. Kennedy. A Conceptual Foundation for Autonomous Learning in Unforeseen Situations. In *Intelligent Control (ISIC), 1998. Held jointly with IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA), Intelligent Systems and Semiotics (ISAS), Proceedings*, pages 483–488. IEEE, September 1998.

[Ken98b] Catriona M. Kennedy. Anomaly driven concept acquisition. 1998.

[KHB10] Achim Kuwertz, Marco Huber, and Jürgen Beyerer. On Adaptive Open-World Modeling Based on Information Fusion and Inductive Inference. Technical Report IES-2010-16, Vision and Fusion Laboratory, Institute for Anthropomatics, KIT, 2010.

[MBF+90] George A Miller, Richard Beckwith, Christiane Fellbaum, Derek Gross, and Katherine Miller. WordNet: an on-line lexical database. *INTERNATIONAL JOURNAL OF LEXICOGRAPHY*, 3:235—244, 1990.

[Mit97] Tom M. Mitchell. *Machine Learning*. McGraw-Hill Education (ISE Editions), 1st edition, October 1997.

[SLS05a] Xiaodan Song, Ching-Yung Lin, and Ming-Ting Sun. Autonomous Learning of Visual Concept Models. In *IEEE International Symposium on Circuits and Systems, 2005. ISCAS 2005*, pages 4598– 4601 Vol. 5. IEEE, May 2005.

[SLS05b] Xiaodan Song, Ching-Yung Lin, and Ming-Ting Sun. Speech-Based Visual Concept Learning Using Wordnet. In *IEEE International Conference on Multimedia and Expo, 2005. ICME 2005*, pages 1138–1141. IEEE, July 2005.

[Syd82] P. H. Sydenham. *Handbook of Measurement Science: Theoretical fundamentals*, volume 1. Wiley, 1982.

[TA08] S. Todorovic and N. Ahuja. Unsupervised Category Modeling, Recognition, and Segmentation in Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(12), December 2008.

[UBB10] Adrian Ulges, Damian Borth, and Thomas Breuel. *Visual Concept Learning from Weakly Labeled Web Videos*, chapter 8, Video Search and Mining. Springer, 2010.