Formal Ontology in Information Systems
R. Ferrario and W. Kuhn (Eds.)

doi:10.3233/978-1-61499-660-6-153

Concepts as (Recognition) Abilities[1]

Fausto GIUNCHIGLIA and Mattia FUMAGALLI
Department of Information Engineering and Computer Science (DISI)
University of Trento, Trento, Italy

Abstract. So far, most of the work in Knowledge Representation has modelled concepts as *classes,* i.e., as sets of instances. However, as from the work in *Teleosemantics*, concepts can also be thought as *abilities* of performing certain (biological) *functions*. The shift is from the study of the means by which the world is represented to the study of the reasons and means by which such representations are generated. In this paper, which is grounded in the seminal contribution by the philosopher Ruth Millikan, we focus on *substance concepts*, namely on concepts as *recognition abilities*, and on how this notion can be mapped to that of concepts as classes. The ultimate goal is to provide a unified theory of perception and knowledge representation that, eventually, will allow us to go beyond the limitations and lack of robustness of current Artificial Intelligence (AI) systems. We provide three main contributions: *i)* a model of concepts as abilities, with a focus on recognition abilities, *ii)* an early version of an *Ontology of (Recognition) Abilities* (called *RAO*) and *iii)* the beginning of a methodology for how to use RAO for discovering which classes, among those contained in the state of the art ontologies, correspond to recognition abilities.

Keywords. Concepts, recognition ability, ontology, cognition, foundations.

1. Introduction

*Concept*s are an essential notion for the understanding of human thought. They allow us to give an account of phenomena such as knowledge acquisition and representation, language understanding, inference, and categorization [1]. A mainstream line of research on this topic, called in the philosophical literature *Descriptionism* [2], takes concepts to be classes. According to this view, a concept of something in the world is a *representation* of this “something”, articulated in terms of sets of properties. Descriptionism has had a large influence on the work in Artificial Intelligence (AI) and Knowledge Representation (KR) and has motivated various KR languages. The main focus of this work has been (and still is) on how concepts can be used to organize knowledge via the *classification* of instances into classes as a function of their properties. Although KR formalisms have been used in several applications with many success stories [3, 4], there are still many open issues related, for instance, to the several roles played by concepts in cognition, see, e.g., [5] for a discussion of some of the issues which arise with this approach.

Lately, the field of *Teleosemantics* [6] has proposed an alternative approach. According to this school of thought, concepts implement suitable (biological) *functions*. The shift is from the study of the means by which the world is represented to the study of the means by which such representations are generated. Here the notion of function

[1] This research has received funding from the *Smart Society* Collaborative Project, funded by the European Commission's 7th Framework ICT Programme for Research and Technological Development under the Grant Agreement no. 600854.

is the same as that used in neurobiology when attributing functions to components of the brain (as in "the function of processing visual information"). According to this view, concepts are components (*devices*) of the human brain characterized by sets of *abilities* of performing, under certain conditions, specific functions. Most relevant to us is the work by Ruth Millikan [7]. Millikan's work concentrates on what she calls *substance concept*s, namely, specific types of concepts which can be characterized as *abilities of recognizing* a certain type of items, that she calls *substances*, which are perceived as being part of the real world [8]. *Substance concepts* have the main function of collecting and accumulating knowledge from the world.

The goal of the work described in this paper is to lay the foundations of a unified theory of perception and KR that integrates the results from the two approaches above. The underlying intuition is to think of *all* types of concepts as abilities, to identify the different forms of functions, and corresponding representations, and to study how these functions can be composed as part of an overall *process* enabling cognition. Thus, if substance concepts are recognition abilities, when we concentrate on the classification task, as it is the case in KR, we think of concepts as *classification abilities*, namely as abilities "… *of simplifying the environment, of reducing the load on memory, and of helping us to store and retrieve information efficiently* …" [2, 9]. This paper provides the following three contributions:

a) It provides a model of concepts as recognition abilities by clarifying their role and by defining their main characteristics. This work can be seen as providing a rationalization and formalization of Millikan's work. The main result is a precise characterisation of the similarities but also the (non-trivial) differences between concepts as recognition abilities and concepts as classification abilities.
b) Based on the results above, it provides the definition of an (early version of an) ontology *RAO*, for *Ontology of (Recognition) Abilities*, as the basis for an integrated study of the two types of concepts.
c) It provides the beginning of a methodology for how to use RAO for discovering which concepts as classification abilities, among those contained in the state of the art ontologies, correspond also to recognition abilities.

It is important to notice that, within KR, various approaches have attempted to provide broader notions of concepts and/or to overcome some of the existing limitations. Some examples are: methodologies for making explicit the semantics of the underlying conceptual models inside KR languages [10], the analysis of cognitive and ontological principles that ground knowledge engineering processes [11], the implementation of new conceptual theories, with a sound cognitive foundation, such as conceptual spaces [12, 13], the perceptual symbol system approach [14, 15], the proxy-type theory [16, 17]. In addition to these theories we may find works addressing the problem of empirical classification and of how to build representations from "observations" [18, 19]. The work described in this paper is orthogonal to this work and, as far as we know, it is the first attempt to provide a unified view of concepts as recognition abilities and as classification abilities.

This paper is organized as follows. Section 2 defines the notion of substance concept. Sections 3 and 4 analyse the various kinds of substance concepts. Section 5 provides a comparison between concepts as recognition abilities and concepts as classification abilities. Section 6 introduces RAO and its main categories. Section 7 provides an example of how to use RAO for the identification of substance concepts among the concepts which are used in state of the art ontologies. Finally, Section 8 analyses the implications of the results presented in this paper for the development of complex AI systems which integrate recognition and knowledge representation, as a first (small) step towards a unified architecture for cognition.

2. Substances and Substance Concepts

We model how things are in terms of *subjects* able to experience the *world*, where by world we mean anything that is external to the subjects themselves. We call these subjects, *perceptual-cognitive systems (PCSs)* [20] to emphasize our focus on the study of systems where perception and knowledge are integrated.

Time is the horizon over which PCSs and the world "meet". A PCS experiences the world through *encounter*s. An encounter is the event through which (a portion of) the world *manifests* itself to a PCS. We call such part of the world, *substance*, where, quoting Millikan, "… *substances are those things about which you can learn from one encounter something of what to expect on other encounters, where this is no accident but the result of a real connection*" [2]. The uniquely identifying characteristic of substances is their ability to manifest some form of invariance through multiple encounters. This invariance is grounded in what we call the substance *causal factor* [21], meaning by this *an inner characteristic which is associated to a substance and which is the cause of its invariance across encounters.* In turn, this invariance takes the form of a set of *outer characteristics* which occur across encounters and allow for the recognition of a substance. Thus, for instance, cats[2], like all species, are characterized *by a homeostatic mechanism* which, in turn, causes them to possess a certain set of common traits (e.g., in shape or weight) and, often but not always, to look similar.

As from the above quote, substances are subjects of learning, namely, of the generation of new knowledge from perception. This process is enabled by *substance concepts*, where substance concepts are taken to be *recognition abilities*, namely *abilities which allow a PCS to realize that the substance involved in the current encounter is the same substance as from previous encounters*. Substance concepts implement functions that allow to recognize a substance as such and to learn and to cumulate the new knowledge about it through a sequence of encounters. They allow to recognize sameness of content in time and also to *group pieces of information together, as being from the same substance* [2]. Substance concepts are innate abilities, which are at the core of cognition, which match the stimuli coming from substances (what we call *signals*) and which allow humans to generate knowledge from signals. Consider, for instance, the substance concept "cat"[3]: we can observe today that cats drink milk or that scratch when we disturb them and this knowledge will be confirmed in future encounters.

The set of (outer) characteristics that a substance manifests over encounters are matched with a set of *substance property concept*s, or simply, *(substance) properties*, which are associated to its corresponding substance concept. Substance properties play a central role in the recognition of substances. A substance property is an *ability to discriminate a substance characteristic over encounters*. This ability is manifested in sameness of reaction to substance characteristics. There are two types of properties that we call *determinables* and *determinates* [22], where determinables can be thought as slots that collect determinates. Thus, for instance, *colour* is a determinable which is

[2] Throughout the paper we write cat meaning *Felis catus*.

[3] To distinguish between substance concepts and substances we write the former in "quotes". Thus for instance, "cat" is an example of substance concept which corresponds to the substance cat.

used to collect determinates such as *blue*, *red*, or *yellow*. From a biological point of view, determinable properties correspond to the use of neurons located in certain early sensory areas of the brain (e.g., colour is tuned to neurons in *visual* cortex) while, on the other hand, determinate properties (e.g., red) would represent single states produced as a reaction to perception (e.g., a red stimulus) [23].

3. Kinds of Substance Concepts

There are two types of substances and, correspondingly, two types of substance concepts, namely *individuals* and *real kinds*. Individuals are single units, scattered in space, enduring through time. In language, individuals are usually revealed by the use of proper nouns or definite descriptions. Examples of named individuals are Barack Obama, my cat Garfield and the Empire State Building. On the other hand, we usually think of real kinds as clusters of elements, what we usually call the real kind *members*, which are characterized by *a common, empirically observable, connection grounded in some*, most often natural, *law*. Real kinds "*... allow successful inductions to be made from one or a few members to other members of the kind not by accident*" [2]. Examples of real kinds are: stuff, e.g., gold or water, biological species, e.g., cat and *Quercus Alba*, artefacts, e.g., chair and car, and also social roles, e.g., doctor and father. The members of real kinds, what we perceive as a "generic" chair or cat, are substances as well.

A first observation relates to the statement that real kinds, their members and individuals are all substances, a statement which is somewhat counter-intuitive for anybody working in KR. For someone coming from this field, the most obvious way to think of the world is to map real kinds to classes and individuals to instances which, in turn, are members of classes. This mapping is discussed in detail in Section 5 below. Here it is worthwhile noticing that with substance concepts we focus on recognition, modelled as an ability. In this respect, both individuals and real kinds share the property that, during an encounter, they are only partially perceived by PCSs. In the same way as we *always* perceive only one or a few members of a kind, we *always* get only a partial view of an individual (e.g., the back or the front). The best way to understand this commonality is to think of substances, no matter whether they are individuals or real kinds, as *wholes* which are only partially presented, with some of their *parts*, by their manifestations to a PCS. In perception there are neither sets nor instances, there are *only* wholes (substances) that are perceived *only* partially. There is however a key difference between real kinds and individuals which is at the basis of the KR representation of the world in terms of classes and instances. Real kinds have the property, not possessed by individuals, of being in multiple places at the same time, meaning by this the fact that any kind can have, at the same time, multiple occurrences inside one or more (contemporary) encounters. This property, clearly, does not hold for individuals: a PCS will perceive at most one individual as part of the same encounter. Thus, for instance, I can perceive two cats together on top of the wall in front of me, but I can only perceive (at most) one occurrence of Garfield per encounter.

A second observation relates to the fact that the inductive grounding that allows the recognition of the same real kind across encounters is very much the same as for individuals. The key observation is that a real kind manifests itself through its members. Both in the case of individuals and of real kinds, the PCS is faced, in time, with similar characteristics that allow a substance to be recognized as being the same from a previous encounter. Thus, for instance, the members of the real kind cat, what we usually call cats, like all species, possess a certain set of common properties (e.g., similar shape and weight) and, consequently, often look similar. Analogously, Garfield, like all individuals, looks pretty much the same across encounters. The ability of substances to manifest some form of invariance through multiple encounters is grounded in their causal factor, as defined above. But the nature of this causal factor is very different between real kinds and individuals. In the first case it consists of some causal connection that is shared by all members of a real kind while in the second case it is related to the fact that the same individual usually changes slowly in time.

The third observation relates to the process by which substances get recognized through substance concepts. This observation is also crucial to understand the distinction between individuals and members of a real kind, a distinction that in KR is blurred into the notion of instance. This distinction is, again, deeply rooted in the profound difference which exists between recognition and classification. Consider for instance an encounter with Garfield. What will the PCS recognize: the individual Garfield or the (member of the) real kind cat, what we usually call "a cat"? We have the first case when recognition happens via the individual substance concept, the second case when the real kind substance concept is enabled. The "selection" of the substance concept is related to the substance properties being recognized. This process is not univocal and depends on many factors. The most important seems to be the actual goal of the PCS (is it looking for Garfield because it wants to feed it or is it just trying to avoid hitting a cat running in front of the car?), but it also depends on the context (e.g., it is harder to recognize an individual at night), on which characteristics are manifested and/or grasped (it is harder to recognize an individual from the back) and so on. Notice that the recognition of Garfield will most likely exploit different characteristics from those used in the recognition of a cat. In the first case, the PCS will exploit those characteristics that uniquely identify Garfield among the other cats, while in the second case it will exploit those characteristics that uniquely identify cats among the other animals. These two sets of characteristics overlap only partially.

The fourth and last observation is that the same substance changes over different encounters thus presenting a set of continuously evolving characteristics. Thus, for instance, two encounters with the real kind cat may produce very different manifestations, though looking similar to two other manifestations which in turn look similar to two other manifestations which …, eventually, will look similar. As a paradigmatic example, under what conditions a person is (recognized as being) the same person as 30 years ago? If I meet a person after 30 years, most likely I will not recognize her as being the same individual. Dually, with an individual with no salient distinguishing marks, there is a high probability to fail its recognition over encounters. Think for instance of forks. In this case what usually happens is that only the real kind fork is recognized as there is no interest in distinguishing among the various individuals. We just look for any fork. This of course will not be the case with that specific fork that I was playing with when I was a kid.

4. Kinds of Real Kinds

Real kinds can be further divided into two more specific categories, i.e., *eternal kind*s and *historical kind*s.

Eternal kinds correspond to what is often called *stuff*, e.g., "gold" and "water". The members of eternal kinds share some fundamental characteristics without being historically related to one another. This inner structure remains stable over time without exceptions. Eternal kinds are often expressed through their mass and/or atomic number, are named using uncountable nouns and are said to have "essences" in a very classical sense, i.e., essences that can be discovered through empirical investigation.

Historical kinds are real kinds "*… for which historical location does play a role in explaining likeness*" [2]. Examples of historical kinds are species, artefacts and social roles, e.g., "doctor" and "baker". Historical kinds are named in this way because their members bear a certain common relation which has evolved in time. For instance, consider species. All their members have a connection with some prior member from which they derive their characteristics. Similarly, all artefacts can be seen as being derived from some prior member, i.e., a prototype, a model of a chair. Finally, younger doctors learn how to act from older doctors.

Historical kinds are strongly correlated and, for what we have figured out so far, include as sub-kinds what the psychologist Eleanor Rosch calls *basic level categories* (and *objects*). These are the concepts that children learn first and use to categorize the world. They are the easiest to recognize via sensory (e.g., visual) and motor interaction with substances. Basic level categories can be detected by running experiment(s) like the one described in [24]. As shown in this experiment, in a hierarchy (a classification) of categories, basic level categories maximize the number of characteristics shared by their members and minimize the number of characteristics shared with the members of their sibling categories. A further characterization of basic level categories is that, usually, the members of their superordinate categories share a very small number of characteristics while the members of their subordinate categories, usually, share a large number of characteristics that, however, are shared also by the members of the sibling categories. The consequence is that the members of basic level categories have a much higher probability of successful recognition than the members of their superordinate or subordinate categories. Recognizing a cat, for instance, is much easier than recognizing an animal or a Siamese cat. In other words, basic level categories provide the ideal balance between the similarity of their members and the dissimilarity of the members of their sibling categories. One interesting observation is that, contrarily to what was initially expected by anthropological and linguistic researchers, biological basic level objects are at the level of abstraction of species, namely one level up from the level of abstraction of the basic level objects which are artefacts (e.g., furniture, as from the experiment by Rosch).

Following Rosch we can further distinguish basic level categories, namely historical kinds, into *biological* and *non-biological basic kinds*. The former are the basic units of biological classification, i.e., biological species, while the latter are defined as the complement of the former and are therefore not well characterized. Examples of non-biological basic kinds, are artefacts like "car" or "chair" (subsumed by superordinate categories like vehicle and furniture, respectively) or social roles like "doctor" and "baker".

5. Recognition Ability and Classification Ability

In KR, the main focus so far has been on classification more than on recognition. As a result, knowledge is modelled in terms of *instances* (e.g., Garfield), *concepts* (e.g., cat), namely sets of instances and *properties* defined as the Cartesian product of two classes, e.g., being of colour yellow, being near something). Concepts are associated to sets of properties and the values of the latter allow to make distinctions among the members of the former. We call below this kind of concepts, *classification concepts*, or simply *classes*, to distinguish them from substance concepts. We also talk of *classification properties* when we need to distinguish them from substance properties.

The mapping between the work and notions defined in this paper and these notions coming from KR can be established based on the following steps:

a) We think of *classification as* the *ability of organizing instances into classes as a function of their properties*. This is the ability that generates and manipulates classes, classification properties and instances as representations of the world.
b) With an overloading of the terms, we talk of substance concepts and substance properties meaning not only the corresponding functions and abilities but also the representations generated by such functions, and dually for classification concepts. This allows to eliminate the difference in approach between us and the "usual" KR approach. For both classification and recognition, we distinguish among devices, abilities, functions and representations only when needed.
c) We acknowledge that recognition and classification are two distinct abilities which generate and manipulate distinct representations of the world, the first being a perception-oriented representation the second being a semantic language-oriented representation of the world. This implies that classes and substance concepts, classification properties and substance properties, instances and individuals are actually distinct representations. It is important to notice that this assumption is coherent with the most recent discoveries in neuroscience which provide evidence that perception and "semantic" oriented representations are actually stored in two different parts of the brain [25].

As a result of these assumptions we are now in the condition of studying the pairwise similarities between the recognition and classification representations of the world. The existence of these similarities is the obvious consequence of the fact that *substance concepts and classification concepts are both representations of substances*. However this mapping is far less obvious than one would expect, the motivation being rooted in the very nature of the functions of recognition and classification. Substance concepts allow to recognize substances over encounters and to acquire knowledge about them, while classification concepts allow to group together substances about which we already have some knowledge. Thus the former are *representations of sets of occurrences of substances*, while the latter are *representations of sets of substances*. With substance concepts we describe substances *over* time, while with classification concepts we describe substances *in* time. Similarly, individuals are *representations of sets of occurrences of substances*, while instances are *representations of (single) substances*. This generates various crucial distinctions.

Let us start from individuals. An individual is a set of occurrences of the same substance and, as such, it can be mapped to the single instance representing that substance. A crucial difference is that individual substance concepts need not have names. Names play no role in the recognition process, while they are crucial in the deployment of classification abilities: you need an identifier to be able to refer to an instance, this is a prerequisite to classification. Furthermore, the mapping individual – instance is not one-to-one. Thus, I can have two or more individuals for the same instance because I did not recognize them as being (sets of) manifestations of the same substance, e.g., myself at the age of five and the age of fifty, or myself dressed as Santa Claus and being recognized as such. The contrary, namely having two or more instances for the same individual, seems to be the case only when there is a need to reason about occurrences of individuals in different moments in time, e.g., because reasoning of the colour of my hair at the age of five and at the age of fifty. Furthermore, I can have an instance which does not correspond to an individual, e.g., the *Minotaur*, because it is a product of the mind with no existence in the real world, or *Homer* who I have never seen in person or described in any form; but I can also have an individual which is not an instance, e.g., a specific part of the mountain I see every day from the window of my office. I look at this view every day, I love it but I do not need to name it because there is no need for me to classify it and reason about it.

Real kinds can be mapped to classes. Classes are sets of instances where real kinds group sets of encounters, one set per instance. Again, as in the case of individuals, and for the same reason, real kinds and their members need not have names. It is interesting to notice that in natural language, when we speak of a member of a real kind, e.g., "cat", we speak of *a cat* meaning a generic cat while we speak of Garfield meaning the specific individual. And the kind of mental image and reasoning that is performed in the two situations is usually different. In most cases we have a one-to-one mapping between a class and a real kind. However, as for individuals, this mapping presents lots of exceptions. Thus we may have a real kind with no corresponding class, similarly to what happens for individuals. We may also have a class with no corresponding real kind. Following Millikan [2] we call these types of classes, *Nominal Kinds* or simply, *Nominals*. Nominals are sets of instances which do not share a causal factor but that, rather, *are grouped together according to a definition provided in terms of the set of properties they share*.

With nominals we have two possible situations. In the first case we have one nominal kind that can be mapped to two or more real kinds. This situation arises with concepts like "stone" or "animal". Thus, "animal" can be thought as the union of the real kinds "cat", "dog", "lion", and so on. With stones, for instance we may focus on their shape, weight, composition and so forth. These types of nominal kinds are concepts that are high in the abstraction hierarchy and are very useful to classify and organize real kinds. But this conventional, theory driven, characterization of stones has nothing to do with the rich, recognition driven substance concepts. When I say "animal" which image of which animal should come to my mind? A cat, a crocodile, or …? In the second case we have a real kind that can be mapped to two or more classes. This situation arises any time we distinguish among the members of a real kind by assigning them some specific property. Thus, for instance, "cat" can be thought as the union of "white cat", "black cat", and so on.

If we concentrate on the mapping between substance properties and classification properties we have pretty much the same situation as for instances and real kinds. Even if some mapping exists, it is not one-to-one and it presents various exceptions. Thus, for instance, as from [23], the (substance) properties that we use in recognition are often quite different, and much more complex from the classification properties we use to describe substances. There is also a further interesting twist. While most classification properties map to substance properties, this turns out not to be the case for *social roles*. Social roles are real kinds that in KR are modelled as properties. Social roles are real kinds as a consequence of the fact that their members, e.g., doctors, manifest similar behaviours. No matter the concrete individual who is playing the role of doctor, given a certain situation, her activities will be similar to those of other doctors, as they would be needed to carry out their duty.

6. An Ontology of (Recognition) Abilities

Figure 1 introduces *RAO*, for *Ontology of (Recognition) Abilities,* a very early version of an ontology which organizes the notions introduced above. RAO must be understood as follows. As from the root, we concentrate on substance concepts taking them to be abilities as well as the representations generated by these abilities. Nominal kinds are added (in dashed lines) for completeness meaning by this that they can be thought as extreme cases of substance concepts with no recognition ability.

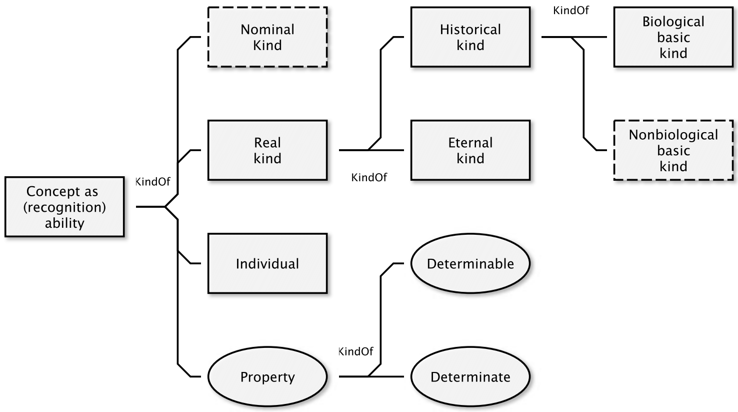

Figure 1. The RAO ontology.

The structure of RAO is motivated by causal factors, as defined in the rest of this section. Thus, with individuals we have the following:

a) *A natural conservation law.* Individuals have the ability to preserve their properties from day to day [21]. Take for instance a person. If she has brown eyes, is tall, is a good tennis player and knowledgeable about informatics, it is likely that she will have these same traits also tomorrow. A similar argument applies to the other kinds of individuals, e.g., to artefacts.

It is important to observe that the causal factor of individuals works very much in the same way as the physics law of the conservation of energy.

The members of real kinds share an empirically observable connection grounded in some law. Real kinds are taken to be the union of historical kinds and eternal kinds. The connections characterizing historical kinds are provided by the possession of one or more of the following four causal factors:

b) *Being the result of a copying activity*. In this case, historical kinds share determinate properties because of some form of previous "copying" activity. We say that a substance B is a "copy" of a substance A, or that B is modelled on A, or that B is a reproduction of A [22]. B can be a true copy of A, as in the case of genes and viruses. B can also be an indirect copy of A resulting from a wider reproduction process. Thus for instance the heart of a person is an indirect reproduction of the heart of her parents while an artefact is another form of indirect reproduction from some abstract model.

c) *A function.* In this case, historical kinds are associated with a function which defines their purpose. This property is possessed in particular by artefacts, and its concrete appearance is often influenced by the cultural context [21]. Chairs for instance are defined by the function of allowing people to sit on them, and Japanese chairs are very different from European chairs. Social roles, e.g. mother, are examples of human functions.
d) *A similar training.* In this case, historical kinds are living beings, e.g., persons, who have characteristics or skills that are transferred across generations through training. Example kinds are socially constructed substances such as roles, e.g., doctors and bakers.
e) *A homeostatic mechanism.* For instance, the members of biological species can be seen as "*… homeostatic systems […] amazingly well-buffered to resist change and maintain stability in the face of disturbing influences*" [2]. The key observation here is that, despite having many different properties, the members of a species remain stable and relatively similar in time (e.g., adult weight, internal temperature and so forth). This is because species evolve as a result of continuous adaptation and, at the same time, of the necessity for the various genes in a gene pool to be compatible with one another.

The members of biological basic kinds possess factors (b) and (e) while the members of historical kinds that are held together according to (b) or (c) or (d), or their combination, can be grouped in the catch-all category of non-biological kinds.

Finally, eternal kinds can be characterized by the following causal factors:

f) *An inner structure or a single underlying cause.* Eternal kinds have a sort of real essence that can be discovered by empirical investigation.

Notice that the above list of properties, and therefore RAO, is neither claimed to be final nor complete. It is a first characterization that organizes the state of the art and which can be further extended. Among others, an open research issue is whether non-biological kinds can be divided into more fine-grained categories according to some specific applications of (b), (c) or (d). For instance, a possible subordinate category of non-biological kinds could be the "artefact kind", whose members are alike due to a special case of (b), which should be applied in a non-biological sense and (c), applied as in the chair example above.

7. From Classification Abilities to Recognition Abilities – A Case Study

Starting from RAO it is possible to devise the beginning of a methodology for identifying which classification concepts from existing ontologies correspond to substance concepts. In this section we show how this can be done by mapping RAO to the light version of DOLCE[4], i.e., DOLCE-Lite, and to the PIZZA domain ontology[5]. DOLCE-Lite provides a large repertoire of very abstract concepts while the PIZZA ontology classifies more concrete concepts. The resulting mapping is depicted in Figure 2 below.

[4] http://www.loa.istc.cnr.it/old/DOLCE.html
[5] http://protege.stanford.edu/ontologies/pizza/pizza.owl

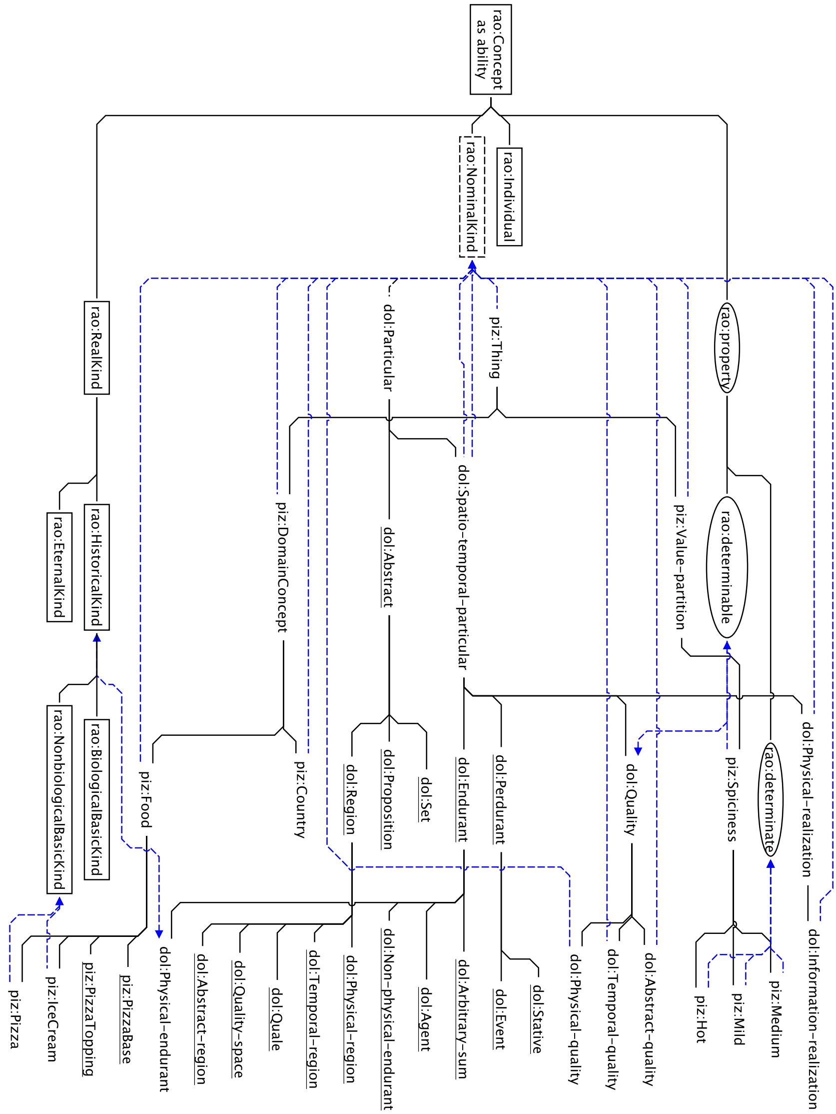

Figure 2. Mapping RAO to DOLCE and the PIZZA Ontology.

Due to lack of space we have reported our results only for the first three levels of the input ontologies. The RAO concepts are identified with boxes (kinds/individual) and ellipses (properties). `dol:`, `piz:`, `rao:`, stand for concepts coming from DOLCE, PIZZA and RAO, respectively. The underlined terms denote classes which are not mapped. Dashed edges denote the mapping, where two arrows mean equivalence and one arrow means subsumption.

Focusing on DOLCE-Lite, 36% of its classes are mapped to RAO. Two classes are mapped with an equivalence relation. The first is `dol:PhysicalEndurant`, which is mapped with `rao:RealKind`, the second is `dol:Quality`, which, according to our current understanding, maps to `rao:determinable` category. 28% of the

DOLCE classes (78% of all the mapped classes) are mapped to `rao:NominalKind` via a subsumption relation. As an example of the kind of reasoning which motivates the mapping with nominal kinds, consider the `dol:SpatioTemporalParticular` class, (Figure 2), which is labelled as nominal. The reason for this choice is that the class "spatio-temporal particular", as from DOLCE-Lite, is a "*dummy class for optimizing some property universes*". Here we may find concepts such as endurants (i.e., `dol: Endurant`), perdurants (i.e., `dol:Perdurant`) and physical realizations (i.e., `dol:PhysicalRealization`), which cannot be grounded in a causal factor.

85% of the PIZZA ontology classes are successfully mapped, resulting in eleven subsumption relations. For instance, `piz:hot` and `piz:mild` are subsumed by `rao:determinate`. These concepts map well to our definition of determinate. They are distinguishable because of a set of common features (in this case, features related to spiciness, i.e., `piz:spiciness`, which perfectly map to `rao:determinable`). Similarly, classes like `piz:Pizza` and `piz:IceCream` are subsumed by `rao:NonBiologicalBasicKind`. In fact, if, as from the experiment by Rosch, `piz:Food` is a nominal, pizzas are artefacts grounded in a "copying" causal factor. 38,5% of "pizza" ontology classes (45,5% of all the mapped classes) are mapped to nominal kinds via subsumption. These classes are clear examples of conventionally defined, theory driven, groupings (e.g., `piz:Value-partition`).

This exercise provides the highlights of a general methodology for identifying substance concepts. At the same time, it also provides evidence of the fact that there is a need to further refine RAO: 64% of the concepts from DOLCE-Lite and 15% of the concepts from the PIZZA ontology are unmapped. Let us consider some examples. For instance, following Millikan's suggestion that *events* are substance concepts, the classification concept `dol:Perdurant`, should be mapped to a new recognition ability. As another example, the concept `dol:Endurant` and its more specific concepts should be linked to RAO through several more classes which are more specific than the ones we have provided so far. For instance, `dol:AmountOfMatter`, is a `dol:PhysicalEndurant` which can be successfully mapped to what Millikan would call "stuff", i.e., a kind of `rao:EternalKind`. A further issue is whether quality spaces, as defined in DOLCE, can be used as kind of "determination dimensions" and whether they can be employed to guide the linking between determinables and determinates. This analysis is essential to explore the relation between the concept of "determinate" and the concept of "quale" thus providing a contribution to the modelling of properties [19]. All these examples define a path of research that will allow us to provide a clear mapping of how, in practice, we could deploy concepts which implement a recognition function, a classification function, or both in an integrated manner.

8. Implications on AI Systems

The differences between substance concepts and classification concepts highlighted above provide interesting insights on how to build integrated AI perception and reasoning systems. The general question is which concepts should be selected for artificial recognition (e.g., vision, sensory) systems [27] and how we should treat them in relation to the classification concepts which are represented inside KR systems, for instance as elements of ontologies. How to create a mapping between these two kinds of concepts is a very well-known open problem, i.e. the *semantic gap* problem, which has been solved only in very particular situations [28].

From the point of view of recognition, substance concepts are the ones where most efforts should be concentrated, as they are the concepts that, thanks to the causal factor in which they are grounded, have a more immediate mapping with their appearance. Thus for instance, the recognition of cats and dogs from a set of pictures will need to handle less diversity than the recognition of animals. Dually, the function of classification is very valuable and worthwhile with nominals, not only because it allows to organize substance concepts (thus delegating them the recognition function) but also because the definitions of nominals are very stable, with essentially no exceptions. These definitions can in fact be provided as sets of properties with no need to map with the complexity and infinite variety of the real world. Animals, for instance, are best defined and thought of as cats or dogs or …, without trying to recognize them in terms of the sensory input.

A further interesting situation arises when there is a need to recognize some specific sub-kinds of real kinds, for instance when we need to distinguish cats by their color. This type of categorization turns out to be useful in substance recognition from sensors as it allows to apply to recognition the compositionality of meaning which is intrinsic in knowledge representation. The work described in [29] is a rather successful experiment in this direction. We are aware of the consequences of what discussed by Millikan in her critique to Fodor [8], namely that the compositionality of properties in KR may not correspond to the compositionality of properties in recognition. Thus, even if I know how to recognize a person and how to recognize a hat I may fail to recognize that very same person wearing that very same hat. However, our early attempts to apply compositionality to recognition hold a lot of promise, the main reason being that they exponentially decrease the cost of training of the learning components [29].

The big challenge is how to manage and reason about those substances for which we have both classes and substance concepts. The problem is that the static definition of classes does not fit well with the variance of appearance of their substances, and therefore, with the corresponding variability of substance concepts. A long discussion on this issue leading, among other things, to the distinction between *conception* and concept can be found in [2]. As a small example, it is essentially impossible to provide a definition of what the real kind “cat” is, as what makes a cat “a cat” is its causal factor while its apparent characteristics change in time. A discussion of this issue is out of the scope of this paper. Our general approach, which will be the topic of a follow-up paper, is that the (substance) properties which are chosen when defining, e.g., the class “cat”, should not be fixed a priori but, rather, should be *adapted* at run time as a function of the goal which the definition must serve, for instance the alignment with what is being recognized by a vision system.

9. Conclusion

In this paper we have provided a characterization of concepts as (recognition) abilities and how they can be mapped to the notion of concepts studied so far in KR, what we call classification concepts. We exploit this characterization to develop RAO, a first version of an ontology of concepts as recognition abilities and we show how it can be used to characterize which classification concepts are only nominals or also substance concepts. Our future work is to further refine RAO and its use in the identification of classification concepts which are also substance concepts. The final goal is to exploit these ideas in the implementation of an integrated system that deploys both a recognition and a knowledge representation function.

References

[1] E. Lalumera (2010). Concepts are a functional kind. *Behavioral and Brain Sciences*, *33*(2-3), 217-218.
[2] R. Millikan (2000). *On clear and confused ideas: An essay about substance concepts*. Cambridge University Press.
[3] F. Baader (2003). *The description logic handbook: Theory, implementation and applications*. Cambridge university press.
[4] R. Brachman, H. Levesque (1985). *Readings in knowledge representation*. Morgan Kaufmann.
[5] D. Vernon, G. Metta, G. Sandini (2007). A survey of artificial cognitive systems: Implications for the autonomous development of mental capabilities in computational agents. *IEEE Transactions on Evolutionary Computation*, *11*(2), 151.
[6] G. Macdonald, D. Papineau (Eds.) (2006). *Teleosemantics*. Clarendon Press.
[7] R. Millikan (2004). *Varieties of meaning: the 2002 Jean Nicod lectures*. MIT press.
[8] R. Millikan (1998). A more plausible kind of "recognitional concept". In *Philosophical Issues*, 35-41.
[9] E. Markman (1989). *Categorization and Naming in Children*, Cambridge, MIT Press.
[10] G. Guizzardi (2005). *Ontological foundations for structural conceptual models*. CTIT.
[11] N. Guarino (2009). The ontological level: Revisiting 30 years of knowledge representation. In *Conceptual modeling: Foundations and applications*, 52-67. Springer Berlin Heidelberg.
[12] B. Adams, M. Raubal (2009). Conceptual space markup language (CSML): Towards the cognitive semantic web. In *2009 IEEE International Conference on Semantic Computing*, 253-260.
[13] P. Gärdenfors (2014). *The geometry of meaning: Semantics based on conceptual spaces*. MIT Press.
[14] L. Barsalou, M. Spivey, K. McRae, M. Joanisse (2012). The human conceptual system. *The Cambridge handbook of psycholinguistics*, 239-258.
[15] G. Pezzulo, L. Barsalou, A. Cangelosi, M Fischer, K. McRae, M. Spivey (2013). Computational grounded cognition: a new alliance between grounded cognition and computational modeling. *Frontiers in psychology*, *3*.
[16] A. Lieto, D. Radicioni, V. Rho (2015). A common-sense conceptual categorization system integrating heterogeneous proxytypes and the dual process of reasoning. In *In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), Buenos Aires, AAAI Press*, 875-881.
[17] J. Prinz (2004). *Furnishing the mind: Concepts and their perceptual basis*. MIT press.
[18] E. Bottazzi, R. Ferrario, C. Masolo. The mysterious appearance of objects, in M. Donnelly, G. Guizzardi (Eds.) (2012). *FOIS 2012*, *Formal Ontology in Information Systems*, Frontiers in Artificial Intelligence and Applications, vol. 239, IOS Press, 59-72.
[19] W. Kuhn (2009). A functional ontology of observation and measurement. In *GeoSpatial Semantics*, 26-43. Springer Berlin Heidelberg.
[20] R. Millikan (2005). *Language: A biological model*. Oxford University Press on Demand.
[21] R. Millikan (Forthcoming). Thoughts of real kinds. *The Oxford handbook of philosophy of psychology. Oxford University Press, Oxford.*
[22] R. Millikan (1984). *Language, thought, and other biological categories: New foundations for realism*. MIT press.
[23] B. MacLennan (1998). Finding order in our world: The primacy of the concrete in neural representations and the role of invariance in substance reidentification. *Behavioral and Brain Sciences*, *21*(01).
[24] E. Rosch, C. Mervis, W. Gray, D. Johnson, P. Boyes-Braem (1976). Basic objects in natural categories. *Cognitive psychology*, *8*(3), 382-439.
[25] A. Martin, L. Chao (2001). Semantic memory and the brain: structure and processes. *Current opinion in neurobiology*, 11(2), 194-201.
[26] N. Guarino, G. Guizzardi (2015). "We Need to Discuss the Relationship": Revisiting Relationships as Modeling Constructs. In *Advanced Information Systems Engineering*, 279-294. Springer International Publishing.
[27] D. Forsyth, J. Ponce (2003). A modern approach. Computer Vision: A Modern Approach, 88-101.
[28] A. Smeulders, M. Worring, S. Santini, A. Gupta, R. Jain (2000). Content-based image retrieval at the end of the early years. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 1349-1380.
[29] J. Stöttinger, J. Uijlings, A. Pandey, N. Sebe, F. Giunchiglia (2012). (Unseen) event recognition via semantic compositionality. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, 3061-3068.