

Millikan + Ranganathan – From Perception to Classification

Fausto Giunchiglia¹, Mayukh Bagchi¹

¹*Department of Information Engineering and Computer Science (DISI), University of Trento, Via Sommarive, 9 I-38123 Povo (TN), Italy.*

Abstract

We assume that *substances* in the world are represented by two types of concepts, namely *substance concepts*, as originally introduced by Ruth Millikan, and *classification concepts*, the former instrumental to (visual) perception, the latter to (language based) classification. Based on this distinction, we introduce a general methodology for building lexico-semantic hierarchies of substance concepts, where nodes are annotated with the media, e.g., videos or photos, from which substance concepts are extracted, and are associated with the corresponding classification concepts. The methodology is based on Ranganathan's original faceted approach, contextualized to the problem of classifying substance concepts. The key novelty is that the hierarchy is built exploiting the *visual* properties of substance concepts, while the *linguistically defined* properties of classification concepts are only used to describe substance concepts. The validity of the approach is exemplified by providing some highlights of an ongoing project whose goal is to build a large scale multimedia multilingual concept hierarchy.

Keywords

Teleosemantics, Concepts as etiological functions, Conceptual hierarchies, Faceted approach, Multimedia lexico-semantic hierarchies

1. Introduction

Concepts are a foundational notion for any theory of mind, no matter whether these theories are more theoretically oriented (as in, e.g., the Philosophy of Mind, Lexical Semantics), or more application oriented, (as in, e.g., Information Systems, Artificial Intelligence (AI), Computational Linguistics). As of today, it is somehow universally accepted that concepts are mental representations of substances, as they occur in the world, but what these representations are, what they are for, and how many of them exist, these are all questions for which there is no uniform view. So far, the mainstream approach, mainly but not only, in the application oriented work, has been to model concepts as *classes*, which in turn are populated by *instances*. This approach was termed *Descriptionism* in [1] to emphasize the fact that, in this work, the main goal is *to describe* what is the case in the world, this description being instrumental to the intended (specific) use of concepts, i.e., that of providing an account of and also implementing phenomena such as knowledge acquisition and representation, reasoning, natural language processing and classification.


CAOS 2021: 5th Workshop on Cognition And Ontologies, held at JWOW 2021: Episode VII The Bolzano Summer of Knowledge, September 11-18, 2021, Bolzano, Italy

✉ fausto.giunchiglia@unitn.it (F. Giunchiglia); mayukh.bagchi@unitn.it (M. Bagchi)

🆔 0000-0002-5903-6150 (F. Giunchiglia); 0000-0002-2946-5018 (M. Bagchi)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

In earlier work [2], the authors propose a complementary view, building upon some recent results in the field of Teleosemantics [3], previously called Biosemantics [4], and in particular of the work of Ruth Millikan [1, 5, 6]. In this work, concepts are seen as (*etiological*) *functions*. Here, the notion of etiological function must be read as meaning ‘*intended for*’ or ‘*devised for*’ with respect to a referent device [7], this notion being adapted from the notion of biological causation. To exemplify, the occipital lobe of our brain is devised for visual perception, where the occipital lobe is the referent device ‘intended to’ perform the function of vision. As from [2], this view allows us to distinguish between two types of concepts, as follows:

1. *substance concepts*, namely mental representations which support the *recognition of what is the case in the world*,¹ and
2. *classification concepts*, namely mental representations which support the *classification of what has been recognized as substance concepts*, classification being the key activity which allows to talk and reason about substance concepts.

Classification concepts capture the Descriptionistic view of concepts as classes. The more novel notion of substance concepts has been successfully tested in the development of a system capable of performing human-like object recognition [8, 9]. Notice that we assume the existence of two types of concepts, where substance concepts and classification concepts perform different functions and are built following independent processes.² The problem which arises is how to make sure that, given a certain substance in the real world, its substance concept and classification concept mental representations both encode the substance itself. The work in [8, 9] provides a general solution to this problem by ensuring that: (i) there is one-to-one correspondence between the two types of concepts; (ii) substance concepts are organized into a hierarchy based on *Genus-Differentia* intensional definitions, as it is usually the case in Lexical Semantics [12, 13], and (iii) the one-to-one mapping is actually correct in the sense that the selected classification concept properly names and describes the substance been perceived.

The goal in this paper is to provide a general methodology for building such an integrated hierarchy of classification and substance concepts. The ultimate goal is to develop a lexico-semantic hierarchy where classification concepts are annotated with media, e.g., photos or videos, representing the corresponding substance concepts. The approach we propose is based on Ranganathan’s faceted methodology for building conceptual hierarchies [14]. The main novelty is that, contrarily to what has been the case so far, see, e.g., all the work in the construction of lexico-semantic hierarchies [13, 15], the hierarchies are built based on the visual properties of substance concepts, rather than by exploiting the linguistically defined properties of classification concepts. The properties of classification concepts are *only* used to describe substances, as perceptually represented by substance concepts and, accordingly, organized in hierarchies. This fact, in turn enables the definition of a very compelling set of rules which allow for the construction of high quality classification hierarchies which at the same time encode the visual properties of substances, as well as their linguistically defined properties, thus being amenable for multi-media classification.

¹We restrict therefore ourselves to substances which can be perceived, e.g., *objects*.

²It is worth noticing how, according to the work in Neuroscience, the brain actually holds multiple representations of the world (at least one per sense) [10, 11].

The paper is structured as follows. In Section 2 we provide a short synthesis of the teleological view of substance and classification concepts. In Section 3 we introduce Ranganathan's faceted view of concepts. In Section 4 we show how Ranganathan's approach naturally maps to the proposed etiological view of concepts. In Section 5 we introduce a set of canons, adapted from Ranganathan's faceted approach, which prescribe how to build high quality hierarchies. In Section 6 we provide some insights about how the ideas presented in this paper can be exploited in the construction of a multi-media lexico-semantic hierarchy. Finally, Section 7 concludes the paper.

2. A Teleological View of Concepts

We assume that the world is populated by *substances* [1, 2], where, in the parlance of Teleosemantics, 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*” [1]. Notice how we have access to substances only indirectly through *encounters*, with encounters being events through which substances partially exhibit themselves to perception over *space-time*. As from the work in Teleosemantics, substances are recognized as *substance concepts*. Substance concepts are of two types: (i) *individuals*, which are *representations of sets of temporal occurrences of substances* commonly identified by employing proper nouns (such as the Colosseum, my pet fish which I named Oreo), and (ii) *real kinds*, which “*allow successful inductions to be made from one or a few members to other members of the kind not by accident*” [1] (such as gold, fish, professor).

The recognition of substance concepts happens via *sets of encounters*. Notice how there is a *many-to-many mapping* between (encounters with) substances and substance concepts. To take an example, depending upon the context, focus or purpose, the encounter with a substance can be recognized as the real kind *chinook salmon*, or as a *salmon*, or as a *fish*, or as the individual *Nemo*, the fish that I have in the aquarium at home. Dually, the same substance concept *fish* can be recognised from many encounters with substances, all differing for some specific feature, e.g., orientation of the fish, light. The choice of the substance concept being perceived depends not only on what is being perceived but also on the mental attitude of the perceiver. Notice, also that what a substance exactly is, and what it does outside encounters, is completely irrelevant to the process by which we perceive substance concepts.

The recognition of a substance concept is grounded in its *causal factor*, i.e., a set of *inner characteristics* which are responsible for its broad *invariance* across encounters (for instance, homeostasis is a causal factor in living beings), and which manifest themselves as a set of *outer (visual) characteristics* which allow to uniquely perceive the substance concept (e.g., color, shape). We model these outer characteristics as *visual objects*, where visual objects are sequences of similar *visual frames* [8, 9]. Via visual frames it is thus possible to compute the visual (dis)similarity among what is perceived. A concrete example would be, for instance, the incremental generation of knowledge about the substance concept *fish*, as recognized from a set of encounters. In such sequences the *fish* could be described, for instance, by visual objects which depict how it appears from the front, from the side, at night, in dark water and so on.

If substance concepts are the means for the recognition of substances, *classification concepts*

are the means for describing and classifying substance concepts. The fact that classification, the ability of “*reducing the load on memory, and of helping us to store and retrieve information efficiently*” [2], is distinct from recognition is reinforced by their very nature as representations - the former consisting of linguistic terms, with the later consisting of visual objects. Classification concepts encode concepts, in Lexical Semantics jargon, *word senses*, and thereby model the diversity of the world in terms of *classes* (i.e., sets of instances lexicalized as nouns) and corresponding *properties*. These are, thus, abilities functioning towards “*organizing instances into classes as a function of their properties*” [2] in the form of lexical-semantic (classification) hierarchies. Such hierarchies are built following the intensional paradigm of *Genus-Differentia* [12]. Here *Genus* refers to an existing intensional definition characterized by a shared set of properties across distinct objects, for instance, the common linguistically defined properties of a trout fish, e.g., its color, size, weight, movements. *Differentia*, instead, implies a set of novel properties, different from the ones of *Genus*, utilized for discriminating amongst objects with the same genus. One such example would be the set of distinguishing properties for *rainbow trout*, *steelhead trout* and *Dolly Varden trout*, respectively, which are among the many sub-classes of *trout*.

A few observations. First, the definitions of substance concept and of classification concept are very different, where the former is provided in terms of temporal sequences of frames, e.g., 2D or 3D videos, while the latter is provided in terms of linguistic descriptions, e.g., in terms of *glosses* articulating Genus and Differentia [13]. Second, substance concepts, differently from classification concepts, do not distinguish between kinds (i.e., classes) and individuals (i.e., instances); these are just two types of substance concepts which are recognized according to the same recognition process, and which can be distinguished only because of their causal factor. Third, the coherence between what is visually represented by substance concepts and linguistically described by classification concepts is guaranteed by the fact that both types of concepts are ultimately generated from the same substances. Fourth, and most importantly, notice how there is a *many-to-many mapping* between substance concepts and classification concepts. Thus, on one hand, it is hardly ever possible to completely describe in natural language what has been perceived, while, dually, any natural language description will hardly ever uniquely identify a substance concept.

3. A Knowledge Organization View of Concepts

The teleological view of concepts gives us a functional categorization of concepts. The next issue is how to define a methodology which allows us to build hierarchies which are coherent with the teleological view. This is where Ranganathan’s work takes a crucial role [14]. Following his approach, we organize concepts according to the *analytico-synthetic* paradigm. Such a paradigm, for any domain of discourse, derives its power from its two component procedures- *analysis* wherein ideas are broadly recognized and decomposed into elemental facets which subsequently undergo *synthesis*, involving the semantic composition of appropriate facets to form concepts. See [16] for a Knowledge Representation (KR) formalization of the analytico-synthetic paradigm. Within this approach, of specific interest to this work is the stratified mechanism proposed by Ranganathan [14] which conjoins the perceptual recognition of concepts with their lexical-

semantic expression and organization, grounded in the analytico-synthetic paradigm. The main novelty relates to the fact that this “*separation facilitates the understanding and exploitation of each plane*” (quote from [14]). We have the following phases (that Ranganathan calls also *planes* or *stages*):

1. *Pre-Idea Stage*, which is focused on the perceptual generation of concepts;
2. *Idea Plane*, which is focused on the organization of concepts in a classification hierarchy, based on their perceptual (e.g., visual) properties;
3. *Verbal Plane*, which is focused on the lexical-semantic rendering of the classification hierarchy (i.e., on linguistically naming the concepts); and,
4. *Notational Plane*, which is focused on formally rendering the classification hierarchy in language-agnostic terms, employing a unique numerical identifier for each concept in the hierarchy.

The *Pre-Idea Stage* is the phase which focuses on the *cognitive grounding* of concepts via the process of their perception, recognition and subsequent mental agglomeration. Accordingly, we take perception to be the “*reference of a percept to its entity-correlate outside the mind*” [14], and define two kinds of percepts facilitating incremental recognition - *Pure Percepts* and *Compound Percepts*. Pure Percepts are, quoting from [14], “*a meaningful impression produced by any entity through a single primary sense and deposited in the memory*”, and Compound Percepts are “*the impression, deposited in the memory, as a result of the association of two or more pure percepts formed simultaneously or in quick succession*” [14]. To illustrate using an example, the (machine) visual acuity of a visual recognition system recognizes the impression produced by a fish eating a shrimp, where the impression produced of eating a shrimp is the *pure percept*, and the object fish corresponds to the *entity-correlate*. The same system, in a successive set of encounters, recognizes the impression produced by the same fish, but this time eating flake food (thus forming a different *pure percept*). The system associates the two impressions together to form the compound percept fish in its memory, which is what we refer to as the “*formation, deposited in memory, as a result of the association of percepts - pure as well as compound - already deposited in memory*” [14]. The process of incremental assimilation of such “*newly received percepts and newly formed concepts with the concepts already present in the memory*”, is what we call, from [14], *apperception*, and the agglomerated memory which is characteristically in continuous evolution across encounters is referred to as *apperception mass*. As from [8, 9], in our terminology, a pure percept is the set of visual objects perceived during an encounter, a compound percept is the result of multiple encounters with the same substance, and apperception mass is the cumulative *memory* of what has been perceived so far, i.e., an *object*.

The *Idea Plane*, being “*a paramount plane which is both a map and foundation*” [17], is built over the apperception mass through *perceptual organization* of the perceived concepts, which Ranganathan terms *Ideas* [14]. Such perceptual organization is pragmatically effectuated by constructing *perceptual subsumption hierarchies*, where these hierarchies correspond to the *visual subsumption hierarchies* of visual concepts defined in [9]. The design of such hierarchies is not based on intuition but informed by a “*a panoply of canons and postulates for designing and evaluating classification systems*” [17]. To illustrate with an example, when a visual recognition

Table 1

Mapping between Knowledge Organization and Teleological View of Concepts

Knowledge Organization View	Teleological View
Pre-Idea Stage	Substance Concept Generation
Idea Plane	Substance Concept Hierarchy
Verbal Plane	Substance Concept Hierarchy in words
Notational Plane	Substance and Classification Concept Hierarchy

system will encounter a successive stream of visual frames composed of different aquatic animal-objects, it will be able to organize them (i.e., the visual concepts induced by images) into a visual subsumption hierarchy by forming genus and differentia in terms of their visual properties, and most importantly, guided by a set of established principles for rendering them ontologically thorough.

The *Verbal Plane* employs “*an articulate language as medium for communication*” [14] of the concepts which are still in the form of a ‘perceptually-grounded’ concept hierarchy. The crux of this phase is to seamlessly annotate such concepts (for instance, visual concepts in the form of objects in images) by employing semantically equivalent linguistic labels (mostly, nouns) from any number of natural languages or domain-specific vocabularies,³ including also namespaces, and thus, in effect, assigning language label(s) to each such concept. As from [8, 9], in our terminology, the *Verbal Plane* is the visual subsumption hierarchy transformed into a lexico-semantic classification concept hierarchy by labeling all the substance concepts with linguistic labels, articulating their Genus and Differentia, with respect to the other substance concepts. Notice that, as from [15], there will be a different hierarchy for each distinct natural language (e.g., English or Italian or Hindi) and that these hierarchies do not necessarily have the same shape, because of multilingual *Lexical Gaps*.

A consequence of the process of the linguistic annotation of the *Verbal Plane*, linguistic phenomena such as *homonyms* and *synonyms* get created, a fact which “*causes aberration in communication*” [14] and should be mitigated. This motivates the fourth and final plane, the *Notational Plane*, which prescribes that language labels should be, quoting from [14], “*replaced by symbols pregnant with precise meaning*” thus formally encoding the “*uniqueness of the idea ... and the total absence of homonyms and synonyms*”. As from [15], and articulated in detail in Section 6, the *Notational Plane* is the true space of alinguistic concepts, uniquely identified, and organized into a classification hierarchy. The key observation is that, in the case of a multilingual hierarchy, as it is the case in the work in [15], because of lexical gaps, this hierarchy is a superset of the hierarchy associated to each and any single natural language.

4. Knowledge Organization View vs Teleological View

The correspondence between Ranganathan’s four-phased logical view of concepts and the teleological view of concepts, as detailed above, is represented in Table 1. The order from top to down explicitly indicates how, progressively, what is being perceived is transformed into a

³That is, an *Object Language* as it is called in [14].

hierarchy of classification concepts. The key overall observation is the central role of substance concepts, and therefore of perception and visual properties, in particular during the first two phases, i.e., the pre-idea stage and the idea plane, where all the decisions about the organization of concepts are taken. This is fully coherent with the work described in [9] where the hierarchy is built using substance concepts and where classification concepts, more precisely, the wording which describes them, are used only to linguistically label substance concepts.

The *first* observation is that Ranganathan's approach imposes a hierarchical organization of substance concepts, which, logically facilitates their mapping to classification concepts. The *second* observation is that the notational plane inherits the substance concept hierarchy as built in the idea plane, where, quoting from [17], the idea plane "*genetically determines the quality of the ultimate product*", i.e., the notational plane. Ranganathan characterizes, quoting from [14], "*the relation between the idea plane and the notational plane*" as being "*the one between a master and a servant*", which is aligned with our own characterization of idea plane as *the determiner of the taxonomic backbone of the (final) classification concept hierarchy*. The *third* crucial observation is that the distinction between substance concepts and classification concepts is logically realized by applying, quoting from [14], "*the Wall-Picture Principle ...*" where "*... Idea first, word next*". The intuition is that, just like a mural cannot be executed in the absence of a wall, there is no existence of (linguistically rendered and subsequently numerically identified) classification concepts without recognition of substance concepts in the first place. *Fourth*, the very stratification of the process of building concept hierarchies aids "*to solve independently, in the first instance, the problems arising in each plane*" [14], thus rendering each phase *characteristically autonomous yet functionally linked*. *Fifth*, it is worth noticing that the four phased mapping above is conceptually governed by the *Law of Local Variation*, which is the principle that "*there should be provision ... for strictly local use, results alternative to those for general use*" [14]. This principle is crucial as it accommodates the fact that the mapping between substances, substance concepts and classification concepts can vary depending on, e.g., the purpose or focus.

Last but not least, notice how the process highlighted in Table 1 enforces a one-to-one mapping between substance concepts and classification concepts, as mentioned in Section 2: the proper natural language label and description will be selected based on the current (partial) view of the object under consideration. So, for instance the *same* substance will be named, e.g., a *person*, a *woman*, *Mary*, depending on the visual details which are perceived. In other words, the many-to-many mappings existing between substances, substance concepts and classification concepts mentioned at the end of Section 2 is properly encoded as a set of one-to-one mappings built by assigning labels not in terms of substances as such but, rather, in terms of the relevant substance concepts. It is worth noticing that this approach provides a solution to a long standing unsolved problem that computer vision systems have, the so called *Semantic Gap problem*, which was already identified in 2010 [18] as (quote) "*... the lack of coincidence between the information that one can extract from the visual data and the interpretation that the same data have for a user in a given situation.*". In this quote we take substance concepts to encode 'the information that one can extract from the visual data' and classifications concepts to encode 'the interpretation that the same data have for a user in a given situation'. Because of how they have been constructed, all hierarchies of media constructed so far, including ImageNet [19], suffer from the Semantic Gap problem (see also Section 6).

5. A Canonical Framework for Concept Hierarchies

The adoption of Ranganathan's methodology enables us to exploit its normative principles, called *canons*, which norm how to dynamically perform knowledge classification [14]. The stress is on, quoting from [20], a "*well designed classificatory language ... capable of individualising microscopic thought-units*", thus facilitating "*the representation of a multi-dimensional continuum in a single dimension*". The pre-idea stage is not governed by canons as it is pre-eminently a *causal phase* in generating new concepts from objects via recognition. We analyse below the other three phases where, as to be expected, the canons for the idea plane are by far the most important.

5.1. Canons for Idea Plane

The canons of the idea plane are organized in four specialized sets, to be applied sequentially one after the other. They are: (i) (canons about) characteristics; (ii) (canons about) *succession of characteristics*; (iii) (canons about) *arrays* and (iv) (canons about) *chains*. Let us analyze them in detail.

Characteristics (by which we mean, outer characteristics, in our terminology, substance properties) form the basis of classification of substance concepts, and the objective is to select such characteristics as will be helpful for our purpose. Let us consider the four which are most relevant. The *canons of differentiation* and *relevance* are conjoined in their purpose, in the sense that the former ensures that a characteristic employed for classifying substance concepts should, quoting from [14], "*differentiate some of its entities - that is, it should give rise at least to two classes*", whereas the later corroborates that such a characteristic "*should be relevant to the purpose of the classification*" itself [14]. For example, while the impossibility of unambiguous classification of salmon and trout on the basis of (visual) recognition of gills in them is ensured by the canon of differentiation, relevance informs that fin spots are an appropriate (visual) differentiating characteristic if the purpose is to classify fishes as per geographical habitat. Further, the *canon of ascertainability* enforces that a classifying characteristic "*should be definite and ascertainable*" [14], in perceptual terms. To take an example, the presence or absence of *pyloric caeca*, which is a part of internal anatomy in many fishes, cannot be construed as an ascertainable characteristic for visually classifying different fishes. Finally, the *canon of permanence* states that, quoting from [14], "*a characteristic used as the basis for classification ... should continue to be unchanged, so long as there is no change in the purpose of classification*", a direct exemplar of which is the fact that *colour* cannot be used as a (perceptual) classificatory characteristic for those *fishes which camouflage*.

The next step is the *succession of characteristics*, namely the order by which characteristics should be applied. It is important to notice that this ordering is crucial as, in case of shared properties, different orderings generate different hierarchies. As an example, the *canon of relevant succession* posits, quoting from [14], "*the succession of the characteristics ... should be relevant to the purpose of the classification*". To illustrate, let us take the case of a visual recognition application for recognizing different fishes. The first logical (visual) characteristic to differentiate, for instance, between salmon and trout will be the *tail shape*, with respect to which the former has a concave tail whereas the trout's tail is convex shaped. Further, the

presence or absence of *round parr marks* can be used by the application as the second (visual) characteristic to differentiate between different varieties of trout, such as rainbow trout and steelhead trout.

The progressive application of the canons for characteristics and succession of characteristics leads to the formation of *arrays*, which are groups of classes, or categories, bearing coordinate status (i.e., categories which are children of the same node), at all the levels of the subsumption hierarchy. Such formation of arrays are guided by the *canon of exhaustiveness*. Exhaustiveness mandates that classes belonging to an array, quoting from [14], “*should be totally exhaustive of their respective common immediate universes*”, and further, “*any new entity added to the original universe ... should be assigned to any of the existing classes or to a newly formed class*”. This is crucial for visual recognition applications where, for example, all the known varieties of salmon should be made coordinate subclasses of the class salmon with the possibility that a newly discovered variety of salmon can be assigned to any of the existing classes or be classified as a new one based on the recognition of a new set of *visual properties*.

The last step is the formation of *chains*, namely what in graph theory are called *paths*. Here the two canons which are pivotal for developing *taxonomically clean* chains are the canon of *increasing/decreasing extension* and the canon of *modulation*. The canon of increasing/decreasing extension is centered around the correlative notions of extension which “*measure the number of entities or of the range comprised in the class*” [14], and intension which define the properties that can be predicated of a class. Based on these notions, decreasing extension states that while traversing down a chain, quoting from [14], “*from its first link to its last, the extension of the classes ... should decrease and the intension should increase at each step*”. Increasing extension, on the other hand, conveys the exact opposite in case of upward traversal in a chain. The second and last canon that we consider is the canon of *modulation* which states that such a chain should comprise one class “*of each and every order that lies between the orders of the first link and the last link of the chain*”, or in other words, the assertion that a chain shouldn’t have any *missing link*. A direct consequence of this canon on the ability of recognition (especially for *human-like visual recognition*) is exemplified by the established fact [21] that there are certain basic categories that are probabilistically most optimal to be perceptually recognized and can never be missed out (for example, for fishes, we cannot skip the class fish and directly jump from aquatic vertebrate to salmon).

Notice that while some of the canons mentioned above are more or less always followed in the state of the art (linguistically constructed) hierarchies, others are not, thus resulting in classification of low quality. Some examples in the first class are the canon of differentiation and the canon for increasing/ decreasing extension which holds by construction in all hierarchies built using Genus-Differentia. Examples of the second class are: the canon of permanence, the canon of relevant succession, (sometimes) the canon of exhaustiveness and the canon of modulation.

5.2. Canons for Verbal Plane

The next step is to linguistically label substance concepts with language labels (nouns). The *canon of context* prescribes, from [14], that “*the denotation of a term in a scheme for classification*

*should be determined in light of the different classes (Upper links) ... belonging to the same primary chain as the class". It is unified in its purpose with the canon of enumeration which stipulates such a denotation to be also determined "through the subclasses ... enumerated in the various chains having the class ... denoted by the term in question as their common link" [14]. The two canons above logically mediate the many-to-many mapping between substance concepts and classification concepts. To take an example, the contextual recognition of the substance concept *fish* as an aquatic vertebrate, a pet or a food depends on its superordinate classes, whereas its precise extensional meaning, in other words its sense disambiguation in classification concept terms, is defined by the subclasses it enumerates in the context of the linguistic hierarchy. Finally, the *canon of reticence* states, from [14], "the term used to denote a class ... should be the one current among specializing in the subject field", or in other words, it prescribes the usage of an appropriate domain language (such as namespaces) for unambiguous annotation of substance concepts, for instance, images. The main goal here is to avoid the use of synonyms.*

Here it is to be noted that all these canons are most often followed in the state of the art hierarchies, the first holding by construction in Genus-Differentia hierarchies, the second holding any time the get-specific principle is applied [22].

5.3. Canons for Notational Plane

The canons for the notational plane are aimed at translating the linguistic hierarchy of the verbal plane into a fully formal hierarchy of alinguistic classification concepts, wherein each concept (more specifically, each sense) is associated to a unique numerical identifier. The *canon of synonym* specifies that, quoting from [14], "each isolate idea should be represented by one and only one isolate number", which, in our context, ensures that each classification concept is represented by one and only one identifier. On the other hand, the *canon of homonym* implies that "each isolate number should represent one and only one isolate idea" [14]. Thus, these two canons, in effect, impose a *necessary and sufficient* condition between concepts and their respective identifiers. Further, the *canon of hospitality in arrays and chains* [14], for us, states that a new concept can be appropriately positioned and uniquely identified anywhere in the hierarchy. These canons cumulatively ascribes to the notational plane the quality of *perpetuation*, "the devices necessary and sufficient to represent uniquely and unequivocally—that is, to individualize—every new formation thrown forth ... from time to time" [23], which attests to its continuous evolution. Thus, a true *classification concept hierarchy* emerges in the Notational Plane, with the unique identifiers performing Word Sense Disambiguation (WSD) and also rendering the space, synonym, homonym and polysemy free, at the same time.

These canons, while not holding in general, are satisfied by all WordNet-like hierarchies [13].

6. From Media to Classification Concepts via Substance Concepts

The use of media, e.g., videos or photos, is quite pervasive, in particular, but not only, in the Web. This phenomenon extends also to hierarchies, for instance, in the case of eCommerce,

where the user is able to seamlessly navigate a catalog where each item is annotated by, usually, an image. The key observation is that, in these situations, the main description is provided in natural language, while photos have a complementary role of integrating visually the main content provided linguistically. ImageNet [19] is a very important point in case. ImageNet is a very large image database which is extensively used for the training of Deep Neural Networks. It has been built by taking WordNet, its English Version from Princeton [13], and by populating it with millions of photos collected from the Web. As also described in [19], the construction of ImageNet has been done in a way to preserve a high level of quality. However, for how it has been constructed, *viz.* by populating a linguistic hierarchy with photos, there is no guarantee that the photos provide the information that would be needed to build the visual subsumption hierarchy implicitly assumed by WordNet. In other words, while by construction, ImageNet is a hierarchy of classification concepts, there is no evidence that the photos encode also substance concepts. While this will most likely have no implications in the recognition of single substances, some difficulties may arise in case one is interested in learning concept hierarchies, and also any time the problem of the Semantic Gap raises difficulties (see discussion in Section 3). Furthermore notice how these limitations apply to all approaches where linguistic hierarchies are used to classify objects, see, e.g., [24, 25] or, more generically, to support computer vision, see, e.g., [26].

The goal of the project introduced in this section, whose preliminary name is *MultiMedia UKC*, is to build a resource very similar to ImageNet but with the key difference of being built following the methodology defined in this paper. The starting point is the Universal Knowledge Core (UKC) [15],⁴ a multilingual lexical resource now containing more than one thousand languages and more than one hundred thousand classification concepts. Besides its size, both in terms of concepts and languages, which is a strong incentive towards its use, the UKC seems very well suited for our goals as its organization matches quite naturally Ranganathan's four-phased methodology. Starting from the UKC, we envision the following construction of the MultiMedia UKC:

- *Pre-Idea Stage*: This phase is used to construct (substance) concepts by extracting *visual objects* from media. The media will be selected to depict the terms in the UKC. The extraction of visual objects will be done applying the techniques from [8, 9];
- *Idea Plane*: This phase is used to *construct* a visual subsumption hierarchy via visual objects. This is done by applying the methodology described in this paper;
- *Verbal Plane*: This phase is used to annotate substance concepts with *words*, which, in turn, are annotated with *synsets* [13], i.e., the set of their synonyms. Synsets are further annotated with their definition (i.e., their *gloss*) defined in terms of Genus and Differentia. There is one verbal plane per language. This is achieved by aligning the hierarchy constructed in the idea plane with the UKC hierarchy;
- *Notational plane*: This phase is used to generate a set of language independent classification concepts, as unique, alinguistic identifiers. This is done reusing the concepts which already exist in the conceptual layer of the UKC.

The construction of a multimedia UKC is quite ambitious and it is bound to raise complications in all the first three steps. The notational plane should come for free by reusing the current

⁴An online version of the system can be reached at <http://ukc.datascientia.eu/>.

UKC identifiers. In the pre-idea stage the complexity comes from the issue of how to extract features from media, a problem which can get rather complicated in non-ideal situations, e.g., in presence of noise. The intuition is to solve this problem with the help of the human supervision. In the Idea plane the complication comes from the need of selecting how to apply the canons introduced in Section 5, a task which requires the involvement of an expert. The complication in the verbal plane comes from the need to match the hierarchy built in in the previous step with the hierarchy which already exists in the UKC. In practice we expect the classification process of the second step to be strongly driven by what is already available in the UKC. And for sure there will be a lot of iterations between these two steps.

To illustrate the kind of reasoning that we will have to perform while constructing the Multi-Media UKC, we introduce two small examples, which relate to two key tenets of categorization from the work on basic categories by Eleanor Rosch [21, 27]. The starting point is her empirical observation that in taxonomies (quote from [21]),

“there is one level of abstraction at which the most basic category cuts are made”, where she defines basic categories as “those which carry the most information ... and are, thus, the most differentiated from one another”.

Coherently with the methodology proposed here, she observes that in perceiving and categorizing objects (quote from [27]),

“objects may be first seen or recognized as members of their basic category, and that only with the aid of additional processing can they be identified as members of their superordinate or subordinate category.”

In complement to the above tenet, her second tenet establishes that (quote from [27])

“basic objects appeared to be the most abstract categories for which an image could be reasonably representative of the class as a whole”.

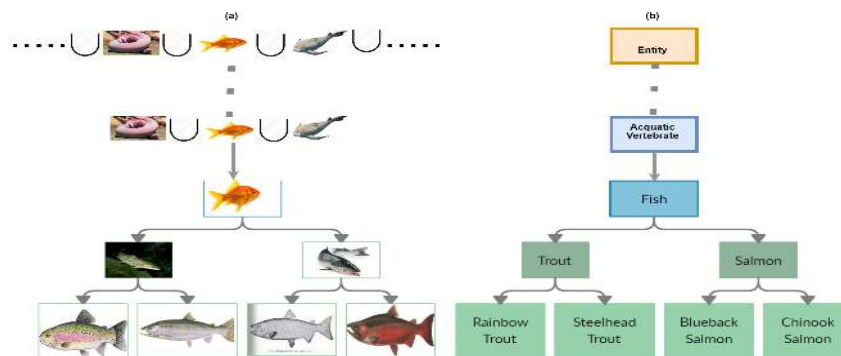


Figure 1: Biological Taxonomies of Fish (picturized from Rosch *et al.* [21])

Let us concentrate on the biological taxonomy depicted in Figure 1, as from [21], where the second classification is a natural language description of the first. Following Rosch’s first tenet, *fish* is a basic category. In fact, fishes share the maximum number of visual properties amongst themselves and are also most differentiable amongst other sub-categories of *aquatic vertebrate*

(e.g., *placoderm* and *agnathan*, not mentioned in the linguistic hierarchy on the right of Figure 1, but depicted in the left figure by the images which annotate the nodes above the node associated to fish).

The first example considers the situation when we move upwards from the basic category *fish*, therefore increasing extension. Here the role of the canon of increasing/decreasing extension provides the logical means of organizing the superordinate categories into a taxonomically clean chain. Let us consider for instance the concept of *aquatic vertebrate*. In Figure 1 the big union symbol sign means that the extension of this concept should be taken as the disjoint union of the elements depicted. As the picture suggests, this concept cannot be visually recognized by any representative image and can only be perceived by considering the images of all the different basic categories, e.g., *fish*, *placoderm* and *agnathan*. In other words more abstract substance concepts should be constructed by joining the substance concepts of the child concepts. But this comes for free, and the distinction is only linguistic and *not visual* as, also for the lower categories, e.g., *fish*, substance concepts are the union of distinct and different visual objects. This is a general phenomenon which we believe it will apply any time we will move towards the root of the hierarchy and that, it seems, has been largely overlooked so far in the mainstream computer vision literature. Another example is the concept of *vehicle*, which can only be visualized as the union of the extension of its subordinate concepts, some of which are basic categories, e.g., *car*, *bike*, *ship*, *train*. Dually, the subordinate categories of *fish*, e.g., *blueback salmon* *chinook salmon*, on the other hand, can be visually recognized by incrementally recognizing the finer visual properties of the basic category *fish* (decreasing extension) over successive encounters. However how far down it is possible to go in the recognition of the subordinate categories before falling into Rosch's second tenet, is an open question for which, at the moment we have no answer. We expect that it will vary a lot from one basic category to another.

The canon of modulation, on the other hand, logically facilitates the identification of co-extensiveness of the categories which are superordinate and subordinate to the basic categories by ensuring the impossibility of missing links. As a matter of fact, this factor definitely confirms the primacy of basic categories in the process of perception. For example, a recognition of *rainbow trout* as a subordinate category of *fish* thus skipping *trout* as the category between *fish* and *rainbow trout* fails in its very purpose of incremental visual classification as it will bring up *rainbow trout* to the same level of visual co-extensiveness as *salmon*. Further, if the system visually recognizes *trout* to be the basic category and *aquatic vertebrate* to be its immediate superordinate category, the system fails in its adherence to human-like vision which, as has been extensively established in [21, 27], recognizes *fish* as a biological basic category. Thus, to restate, the impossibility of missing links (and hence, the canon of modulation) ensures the primacy of basic categories. This at the moment is only an intuition. We believe that the construction of the MultiMedia UKC will allow to (dis)confirm this intuition quantitatively.

7. Conclusion

Usually, when implementing object recognition systems, and also when building lexico-semantic hierarchies, media, e.g., videos or photos, are organized and classified based on their linguistic description. This is correct as the purpose of language is exactly that of describing what is

known. However, based on some recent results in Computer Vision, this paper suggests that, when in the process of recognizing objects and classifying them in conceptual hierarchies, visual properties are much more relevant. It also suggests to use Ranganathan's faceted approach which is articulated exactly in terms of how visual properties should be progressively refined up to the generation of a not ambiguous linguistic description. The work described here are just the first steps towards our ultimate goal, namely the construction of a large scale multilingual multimedia lexical resource.

Acknowledgements

The research conducted by Fausto Giunchiglia and Mayukh Bagchi has received funding from the "DELPhi - DiscovEring Life Patterns" project funded by the MIUR Progetti di Ricerca di Rilevante Interesse Nazionale (PRIN) 2017 – DD n. 1062 del 31.05.2019.

References

- [1] R. G. Millikan, *On clear and confused ideas: An essay about substance concepts*, Cambridge University Press, 2000.
- [2] F. Giunchiglia, M. Fumagalli, Concepts as (recognition) abilities, in: FOIS, 2016, pp. 153–166.
- [3] D. Papineau, G. Macdonald, *Teleosemantics*, Oxford University Press, 2006.
- [4] R. G. Millikan, Biosemantics, *The journal of philosophy* 86 (1989) 281–297.
- [5] R. G. Millikan, R. G. Millikan, *Varieties of meaning: the 2002 Jean Nicod lectures*, MIT press, 2004.
- [6] R. G. Millikan, *Language: A biological model*, Oxford University Press on Demand, 2005.
- [7] K. Neander, The teleological notion of 'function', *Australasian Journal of Philosophy* 69 (1991) 454–468.
- [8] L. Erculiani, F. Giunchiglia, A. Passerini, Continual egocentric object recognition, arXiv preprint arXiv: 1912.05029 (2019).
- [9] F. Giunchiglia, L. Erculiani, A. Passerini, Towards visual semantics, arXiv preprint arXiv:2104.12379 (2021).
- [10] A. Martin, L. L. Chao, Semantic memory and the brain: structure and processes, *Current opinion in neurobiology* 11 (2001) 194–201.
- [11] E. R. Kandel, J. H. Schwartz, T. M. Jessell, S. Siegelbaum, A. J. Hudspeth, S. Mack, *Principles of neural science*, volume 4, McGraw-hill New York, 2000.
- [12] E. A. Hacker, W. T. Parry, *Aristotelian Logic*, Albany: State University of New York Press, 1991.
- [13] G. A. Miller, R. Beckwith, C. Fellbaum, D. Gross, K. J. Miller, Introduction to wordnet: An on-line lexical database, *International journal of lexicography* 3 (1990) 235–244.
- [14] S. R. Ranganathan, *Prolegomena to Library Classification*, Asia Publishing House (New York), 1967.
- [15] F. Giunchiglia, K. Batsuren, G. Bella, Understanding and exploiting language diversity., in: *IJCAI*, 2017, pp. 4009–4017.

- [16] F. Giunchiglia, B. Dutta, V. Maltese, From knowledge organization to knowledge representation, *KNOWLEDGE ORGANIZATION* 41 (2014) 44–56.
- [17] M. P. Satija, Colon classification, *KNOWLEDGE ORGANIZATION* 44 (2017) 291–307.
- [18] A. W. Smeulders, M. Worring, S. Santini, A. Gupta, R. Jain, Content-based image retrieval at the end of the early years, *IEEE Transactions on pattern analysis and machine intelligence* 22 (2000) 1349–1380.
- [19] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: *2009 IEEE conference on computer vision and pattern recognition*, Ieee, 2009, pp. 248–255.
- [20] S. R. Ranganathan, *Philosophy of library classification*, Sarada Ranganathan Endowment for Library Science (Bangalore, India), 1989.
- [21] E. Rosch, C. B. Mervis, W. D. Gray, D. M. Johnson, P. Boyes-Braem, Basic objects in natural categories, *Cognitive psychology* 8 (1976) 382–439.
- [22] F. Giunchiglia, I. Zaihrayeu, U. Kharkevich, Formalizing the get-specific document classification algorithm, in: *International Conference on Theory and Practice of Digital Libraries*, Springer, 2007, pp. 26–37.
- [23] S. R. Ranganathan, Self-perpetuating scheme of classification, *Journal of Documentation* (1949).
- [24] M. Marszalek, C. Schmid, Semantic hierarchies for visual object recognition, in: *2007 IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, 2007, pp. 1–7.
- [25] B. C. Russell, A. Torralba, K. P. Murphy, W. T. Freeman, Labelme: a database and web-based tool for image annotation, *International journal of computer vision* 77 (2008) 157–173.
- [26] D. Porello, M. Cristani, R. Ferrario, Integrating ontologies and computer vision for classification of objects in images, in: *Proceedings of the Workshop on Neural-Cognitive Integration in German Conference on Artificial Intelligence*, 2013, pp. 1–15.
- [27] E. Rosch, Principles of categorization, *Concepts: core readings* 189 (1999) 312–322.