Proceedings of the QQML 2019: 11 th Qualitative and Quantitative Methods in Libraries
International Conference, European University Institute
Florence, Italy, 28-31 May, 2019
http://www.qqml.org

Ontologies and Classification: The Unavoidable Interplay Between Human Reasoning and Machine Reasoning

Mauricio B. Almeida[1], Eduardo Felipe[2] Benildes C. Maculan[3]

[1] Associate Professor at the Graduate Program in Information Science, School of Information Science, Federal University of Minas Gerais, Brazil
[2] Phd. Candidate at the Graduate Program in Information Science, School of Information Science, Federal University of Minas Gerais, Brazil
[3] Adjunct Professor at the Graduate Program in Information Science, School of Information Science, Federal University of Minas Gerais, Brazil

Abstract. Artificial Intelligence (AI) is currently a vigorous subject widely studied and publicized. In recent times, we have seen notable AI advances that lead scientists to put the question if there is classification after arrival of the search engines. The question impacts Library and Information Science (LIS) insofar as in the data environments, there are no shelves, no physical constraints that demand the organization for which classification schemes were created. Indeed, some have defended that no prior scheme can say in advance what a user needs. This advances are made possible by AI algorithms that allow automatic reasoning. In this session, discuss to what extent automatic reasoning can replace human reasoning in classification. First, we introduce some essentials of both human reasoning and of AI. Then, we describe the operation and the reasoning produced by AI approaches – ontologies and machine learning – in addition to compare them with human classifying activity.
Keywords: ontologies; automatic classification; machine learning; artificial intelligence; knowledge organization

1. Introduction

What seems to connect human beings and machines, the latter running AI algorithms trying to replace people in some tasks, is the ability of reasoning or producing inferences. In general, if one infers, she draws a conclusion. From a philosophical point of view, one infers a proposition “p” from a proposition “q” when she asserts “p” on basis of “q”. Also, one infers a proposition “p” from the fact “F” when she asserts “p” one basis of the occurrence of “F”. In addition, one infers a proposition “p” from a proposition “q” when she notes that everyone who asserts “q” is committed to assert to “p” (Sparkes, 1991). For example, Tom asserts that he was in London last Tuesday at 12:00 noon. For example, John infers that Mary could not have been seen lunching in Brighton on Tuesday, but John saw Mary lunching in Brighton on Tuesday. Then, John knows that Mary is not telling the truth.

People also use reasoning to classify, an activity considered fundamental within LIS. Likewise, classification have been also performed with success by current AI algorithms. The connection between these fields – AI and LIS – has already been raised in the literature: “human beings use categories to think and to speak, and also algorithms of AI need some kind of category schema to properly run” (Saracevic, 1996, p.1). However, some inaccuracies incurred by AI machinery have drawn the attention not only of scientists, but also of the lay people. An example is the recent case of the misclassification of the United States Declaration of Independence as a hate speech by a Facebook algorithm (The Guardian, 2018). This type of gaffe show us that there is still much work to be

done to make AI part of our routine. Within LIS, questions have also been raised about the interplay between classifications performed by humans and machines (Hjørland, 2012).
In this session, we critically shed some light about what can be really done by AI. To achieve our goals, we introduce some essentials of both human reasoning and classification (section 2). Then, we clarify some AI essentials and explain the two main current approaches for automatic reasoning – ontologies and machine learning – describing their operation (section 3). Finally, we discuss in to what extent automatic reasoning could replace human reasoning in tasks of classification and related (section 4).

2. Human reasoning – an overview

This section is devoted to explain basic concepts of human reasoning and the main scientific theories that underly it (section 2.1), in addition to introduce activities that can be performed by reasoning, particularly, classification (section 2.2).

2.1. Theories of reasoning

There are theories from Cognitive Sciences and Psychology considered the most acceptable current approaches to understand how people reason. Indeed, these fields have accumulated a large amount of empirical results during the last decades: i) the sentential-rule or ruled-based approach; ii) the mental model approach; iii) the probabilistic approach; and iv) relational complexity approach (Warren, 2017).
The rule-based approach, which traces back the 19th century, supports that human reasoning is based on formal systems. Such formal systems involve the capacity of draw conclusions from sentences using both logical connectives, like "and", "if", "or", "not"; and quantifiers, like "all", "some". A subset of these quantifiers gives rise to predicates for syllogisms, as originally analyzed by Aristotle. For example, if anyone knows that John thinks all parties are boring, one can deduce that John will not be present at the party tonight (Rips, 1983). This example embeds a familiar logic law, namely, the modus ponens, a rule of inference that can be summarized as "P implies Q and P is asserted to be true, therefore Q must be true." However, not all situations are so straightforward and extensive theories have been conceived to explain more complex operations within the rule-based approach (Braine and O'Brien, 1991; Stalnaker (1968, 1984).
An issue emphasized in the literature is the great number of models produced by the rule-based reasoning, for example, there are 64 possible logical premises for a simple syllogism containing two premises and one conclusion. Having so many models in memory, there are chances that some of them are forgotten, neglected or even not created. This could lead to inconsistent conclusions or produce a response bias on valid conclusions (Johnson-Lard and Bara, 1984). Accordingly, the ruled-based approach have been considered a particular case of another approach, the so-called mental model approach, which is considered the most general model. The mental model approach become popular in Psychology, Cognitive Science and Linguistics by defending the principle that, in order to reason, people create models of reality and therefore verify the conclusions taken against these models (Johnson-Laird, 1999).
Another approach, called probabilistic, offers empirical evidences to demonstrate that people make large and systematic errors in performing standard reasoning tasks. These people could be characterized as irrational if one follows the tenets of other reasoning approaches, though the defenders of the probabilistic approach disagree with this interpretation. To reach reasonable

results, the probabilistic approach uses conditional probability rather than logic. In spite of its relative success, the probabilistic approach does not provide any explanation about the operation behind the human reasoning (Oaksford and Chater, 2001).

The relational complexity approach was created to quantify the complexity associated with reasoning about relations. The complexity of a relation is the number of elements associated with that relation, in other words, its cardinality: unary, binary, ternary and so forth. For most people, the relational complexity of cardinality four is the top boundary. om other words, four is the number of relations that people can deal with because of constraints of memory (Halford et al., 2007). Ultimately, the relational complexity approach is considered a kind of mental model approach in which models are conceptualized in a different way (Zielinski et al. 2010).

Finally, it is worth emphasizing that there is no enough evidence available that allows one to establish an unified theory for human reasoning. The variations in the way people reason, the complexity of human brain, the distinct approaches taken, and the different ways to face problems have hindered scientists to find a definitive theory (Stenning and Yule, 1997; Stenning and Van Lambalgen, 2008).

2.2. Classification and Categorization

Nevertheless reasoning is a crucial research subject addressed by several theories of notable fields – Cognitive Science, Neuroscience, Artificial Intelligence, Linguistics, Psychology, to mention a few – it is chiefly involved with another significant investigation: the studies in classification. Indeed, it is impossible to think without categories. In additions, some decisions, as putting something in a category, may have paramount impacts in society. The sentence "when does a human being begin to exist?" (Smith and Brogaard, 2003, p. 1) could be transcribed as when does an organism has features that qualify it to belong to the category of human beings?

To be accurate, one cannot say that "class" is a synonym for "category", even though the terms are sometimes used in an interchangeable way (Jacob, 2004). The issue here is viewing categories either as natural or as constructed by humans. The former case considers natural kinds, many times referred to as Aristotelian entities that exist regardless human thoughts; the latter one considers classes, human artifacts created for specific purposes. (Jansen, 2008). In this sense, to illustrate the difference, one can tell about the "category of trees", which encompasses all trees of the world; and, the "class of trees in the garden of that university", for example, with the aim of designating a gardener to take care of the plants. From now on, we admit these foregoing distinction, even though we try to respect the use of one or another term by different authors.

Research involving categories and classification are a central subject within LIS, in which one can observe a history of at least 200 years of research (Tennis, 2016). Nonetheless, throughout centuries several debatable questions about both classes and categories (as we have already said, each author favoring the use of one or another term) have populated the mind of scientists: how should we classify the world's entities? What is the relationship between classes and the world itself? How scientific classifications should be constructed? Is there an unique top-most category or several ones? How to distinguished one category from another? (Studtmann, 2013).

Philosophy have offered numerous answers, with emphasis to three main schools: i) essentialism, which sorts things according to their essential natures, a method originated in Aristotle; ii) cluster analysis, which divides entities into groups whose members share a cluster of similar features, though none of those features are essential, as advocated, for example, by Wittgenstein; and iii) historical classification, which classifies entities according to causal relations rather than their qualitative features, for example, Darwin's suggestion for classifying organisms (Ereshefsky, 2000).

In spite of the school chosen, a good start point to reflect how challenging is to deal with categories is to remind that something can be a member of more than one category. For example, a "flower" can be thought, and therefore categorized, as a "plant", a "sort of decoration", an "item for sale", and so forth. The assignment of a variety of categories to only one entity is made possible by the use of hierarchies. A hierarchy allows one to consider that "flower", for example, falls under a category of "physical object", along with "car" and "stone". In addition, flower can be found within categories like "plant", "living being", and so forth (Gorman and Sanford, 2004).

Hierarchies, however, do not solve all issues since one can relate categories in several other ways. Again, using a "flower" as example, it is a "plant" and it is also an "item to sale. On the other hand, a "car" can be "something to sale", but it is not a "plant". Not "everything to sale" is a "plant". Accordingly, categories will overlap and then pose obstacles to a hierarchical organization of things in the world. A way of avoiding such overlapping is to adopt the simple aforementioned distinction: a flower is, of its own nature, a kind of plant, a natural entity; on the other hand, an "item for sale" is not natural, it is a class, which exists because someone took the decision of selling it in a shop. Thus, if one considers just natural categories, there will be no overlap and the hierarchical order can be preserved (Gorman, 2004).

However, as we know, human beings can create divisions, and also categories on the world. A "pedestrian crossing" marked on the ground in the parking lot is an example of this kind of division; a "zip code" is another one. In fact, people create divisions through human activities of demarcation. Then, the issue in this case is to discuss if either the categories created by people will be mapped to the world or if they will impose divisions on the world. To illustrate the former, in the design of an architectural blueprint, reality should match the blueprint through the construction of something; likewise, to illustrate the latter, a corporation should match the statute of the corporation that assigned obligations to its members and boundaries to its actions (Bittner and Smith, 2008).

Reality can be viewed in different ways and each of these views involves a particular categorial scheme, or a perspective, from the side of the observer. For example, engineers view "cars" differently of clients that buy the very car. This lead us to other issues, for example, if there are more than one scheme, one for each specific purpose, how we could co-related all of them? On the one hand, we can solve problems from a practical point of view in creating a scheme for each need; on the other hand, in creating several schemes, we become owners of several distinguished arrangements without a good parameter to determine which is the better one. Thus, how there is no reliable way to decide what is the better or the most basic, it is up to us to decide which one to use in each situation.

It is worth reminding that all these foregoing issues and situations are ontological by nature. They do not cover several other usual situations regarding, for example, morality, values, rights, to mention a few, but provide an introduction to the complexity of dealing with categories.

3. Machine reasoning

Whilst in Computer Science, particularly in AI, the word inference is usually understood as logical inference, there is another of type of inference suitable to solve problems anywhere the classic logic does not fit well. It is the domain of uncertainty or incomplete knowledge, in which inferences are generally referred to as probabilistic inferences. Both foregoing types of inference declares to be able to perform machine classification processes. In this section, we evaluate the meaning of the term "classification" in the scope of two approaches representative of logical and probabilistic inferences, respectively, ontologies and machine learning (ML).

3.1. Logical Inferences with Ontologies

To access the realm of IA, a complex field that combine Computer Science, Cognitive Science, Linguistics, to mention a few, one needs to be informed about some essentials and basic concepts. In this section, we first provide such essentials to then explain the automatic classification that can be provided by ontologies.

In the scope of IA, there are representation languages in general composed by sets of logical sentences, usually called knowledge bases, which are able to represent some knowledge domain for computational purposes. Here, the word "sentence" has a technical sense: it is an assertive, a sentence that states facts about the world. Such sentences have a syntax and a semantic: while the former specifies all well-structured sentences, for example, "x +y = 4" is well-structured and "x 2 y + =" is not; the latter defines the truth of the sentences within a model, for example, the semantics of the arithmetic specifies that a sentence "x + y = 4" is true in a model in which x = 2 and y = 2, but it is false in a model in which x = 1 and y = 1. So, models express the notion of truth by enclosing a set of sentences in which all sentences are true. A logical inference here is performed by algorithms for model checking, listing the possible models and verifying if a given sentence is true in all models that compose the knowledge base (Russell and Norvig, 2007).

In this context, logical inference involves logical entailments, for example, if sentence B is true in all models in which sentence A is true, we say that "A entails B". Logical entailments can be applied to derive conclusions. A classical example is the case in which from the sentence "Socrates is a man" and "All men are mortal", it can be inferred that "Socrates is mortal". In addition to logical entailments, other notions are important for the understanding of inferences from the point of view of modern AI, namely: consistency, soundness, completeness, equivalence, validity, satisfiability (Genesereth and Nilsson, 1988).

Consistency describes a set of sentences in which no sentence contradicts any other. Consistency, however, is not a process, it is a condition. When the algorithm comes up with only entailed sentences, we say that it preserves the truth and possess soundness. Another relevant property is completeness, which consists in the capacity of the algorithm in deriving any entailed sentences. Also, a sentence is logical equivalent to another sentence when both sentences are true for the same set of models. A sentence have validity only if it is true in all models. Finally, satisfiability describes a sentence that it is true in some model, or at least it is satisfied by some model of the knowledge base (Levesque, 1989).

Contemporary applications, within Semantic Web, contain devices called inference engines, capable of automatic inferences in the scope of ontologies as artifacts. In this context, inference engines – also called automatic reasoners – are pieces of software designed to organize and reason over categories. The most popular current inference engines work on Description Logics (Baaeder et al., 1992), a family of logics that provides representation languages to create definitions for both categories and relationships between these categories.

One of the most widely common usages for inference engines is classification. In this context, classification is the process of checking whether an instance belongs to a class. In addition to classification, inference engines can also perform additional tasks, for example: subsumption, which checks if a class is a subset of other classes by comparing their definitions; and consistency, which checks whether the membership criteria are logically satisfiable. A currently popular web-based language, composed of description logic profiles, is the Web Ontology Language (OWL), a fragment of the First Order Logic (FOL) that serves as the representation language in the Semantic Web (W3C, 2019).

An example can help to understand the operation of an inference engine (Sattler, Stevens and Lord, 2014). Consider a set of OWL files containing "O", an ontology that encloses several axioms encompassing classes. Here, a logical sentence considered to be true is called axiom. In this context, if one calls a inference engine to work, it carries out three tasks: a consistency test, a satisfiability test, and a subsumption test.

First, the inference engine executes a consistency test checking whether there exists a model of O, i.e., whether there exists a structure of classes and relations in which all axioms in O are true. For example, an ontology composed by five axioms (Table 1) would fail in this test since "Bob" cannot be an instance of two disjoint classes, namely, "human" and "sponge". Classes are disjoint if they can have no members, and therefore no subclasses, in common. If only one axiom is violated, the O ontology fails in the consistency test.

Table 1: some axioms in O ontology

Logical version	Natural language version
1- *Class*: Human	1- There is a class called Human
2- *Class*: Sponge	2- There is a class called Sponge
3- *Individual*: Bob types: Sponge	3- Bob is an instance of a class Sponge
4- *Individual*: Bob types: Human	4- Bob is an instance of a class Sponge
5- *DisjointClasses*: Human, Sponge	5- Instance of Human class cannot be an instance of Sponge class

Second, the inference engine executes a satisfiability test checking whether there exists a model of O containing an instance "x" of a class "A". In other words, it checks whether exists a structure of classes and relations that satisfies all the axioms in O, in addition whether the class A has an instance named x.

Table 2: axioms and relations in O ontology

Logical version	Natural language version
1- EukaryoticCell *SubClassOf*: *Cell and* (*hasPart some* Nucleus)	1- An Eukaryotic Cell is a kind of cell that have a nucleus as part
2- *RedBloodCell SubClassOf*: EukaryoticCell *and* (*hasPart only* (*not* Nucleus))	2- A Red Blood Cell is a kind of Eukaryotic Cell that do not have a nucleus as part
3- Blood S*ubClassOf*: (*hasPart some* RedBloodCell)	3- Blood is a kind of entity that has a Red Blood Cell as part

In this case, a "red blood cell" has a "nucleus" because it is an "eukaryotic cell" (Table 2, line 1); and it has no "nucleus" since it is described as without a "nucleus" (Table 2, line 2). Thus, a model of O cannot have an instance of red blood cell, since such instance would have both "nucleus" and "no nucleus". Such O is then said to contain unsatisfiable classes, namely "RedBloodCell" and "Blood".

Thirdly, for two class named "A" and "B" that belongs to O, the inference engine tests if "A is subsumed by B". In other words, the test checks whether, in each model of O, each instance of A is also an instance of B. There more than way of testing subsumption. An example is to build a model with an instance of A and "notB" and, if it fails, then A and notB cannot have an instance in any model of O. Therefore, A is subsumed by B.

A real example from the biomedical field can illustrate the subsumption test, also called classification (Figure 1). A knowledge base has the class "omeprazole", which falls under “chemical substance”. There is another relation establishing that “omeprazole” inhibits “hydrogen potassium ATPase”, which is a kind of “proton pump”. Accordingly, the knowledge base can automatically infer that “omeprazole” is also a “proton-pump inhibitor” (red arrow in Figure 1).

Figure 1 – logical inference: omeprazole is a proto-pump inhibitor

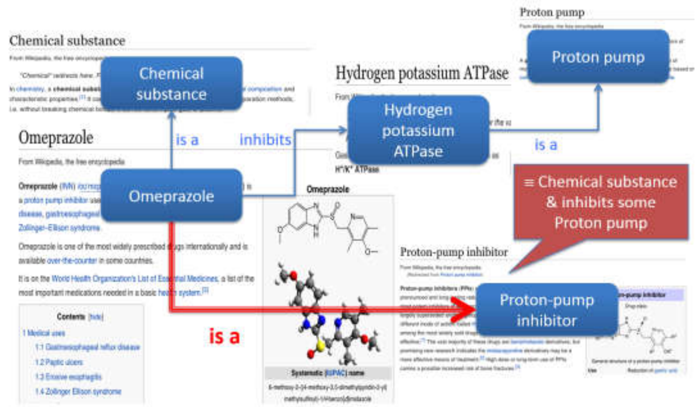

Source: Brochhausen (2018)

3.2. Probabilistic Inferences with Machine Learning

Historically, there are two more important kinds of learning within AI. The first can be illustrated by computers playing chess: they store the moves that resulted in the check-mate and attempts to use them again. The second type of learning is generalization: computers attempt to apply previous learning to situations that are analogous, but not identical. Indeed, current ML consists of algorithms that handle data to make decisions according to probabilities. It works in a similar way to the human brain, relying on probabilities to decide whether it knows something. For example, if a person looks at a Persian cat they rely on their previous experience (data) to identify the other cat as a cat and specifically, as a Persian cat (McLay, 2018).

In spite of its current success, ML it is a field with roots in 1950s. Computers of the past decades, nonetheless, were too slow and lacked memory to deal with large sets of data needed by ML application. Just recently, ML have proven its usefulness. While traditional algorithms need established written rules to govern their use, a ML algorithm is able to learn from a data set, and then both discovering patterns and expanding its capacity of prediction (Apté, Damerau and Weiss, 1994).

The term “predicting”, in general, concerns to the ability of asserting the future, but within ML it means the capacity to assign a value to a variable, not always containing a temporal aspect. If one have an outcome measurement – quantitative or not – that she wish to predict based on a set of features, a ML algorithm can be applied according to the follow steps (Kelleher, Mac and D’Arcy, 2015):

- There is a set of data called data training, which already contains desirable outcomes and features measurements for a set of known objects;
- Through the data training set, the algorithm can build a prediction model enabling the prediction of an outcome for new unknown objects;

Accordingly, the process of building models capable of predictions based on known data, in general, historical series of similar data, is called data predictive analysis (Kelleher, Mac and D’Arcy, 2015). Within such process, three are two main approaches:

- Supervised learning: this modality is called “supervised” because there is an outcome variable to guide the learning process;

- Unsupervised learning: in this modality, there are no measurements of the outcome; the aim is to describe how the data are clustered, as well as to identify patterns;

Within supervised approach, two well-known applications are: i) classification, which classifies the type of a variable, in addition to comes up with an quantitative outcome; and ii) regression, in which the outcome measurement is qualitative. Examples of current problems successful dealt with by supervised approaches are (Hastie, Tibshirani and Friedman, 2008):

- Spams´ prediction of: there is a training data consisting of thousands of emails already classified by people as spam, in addition to the relative frequencies of the most commonly occurring words in email messages. This is a classification problem with the goal of classifying if an email is spam;
- Cancer identification: there are series of correlations between the level of prostate specific antigen (PSA) and other clinical measures in a group of men who were about to suffer prostatectomy. The goal was to predict the log of PSA (lpsa) – a logarithmic scale against time obtained from prostate tumor marker values – along with ten other measurements for variables that impact in the treatment. This is a regression problem because the outcome is qualitative;
- Handwritten identification: a sample of data was extracted from the handwritten ZIP codes on envelopes from postal mail for purposes of handwritten digit recognition. The task is to predict, from a matrix of pixel intensities, the identity of each image (0,1,..., 9) quickly and accurately. This is a classification problem that decides whether a figure is a number.

An example of unsupervised approach can be seen in the domain of DNA Expression Microarrays (Hastie, Tibshirani and Friedman, 2008). The problem involves two main variables – genes and samples – and the question of how to cluster together the samples for an experiment. Figure 2 depicts an example of inference with machine learning in radiology.

Figure 2: probabilistic inference: after trained, the algorithm can detect malignant cells

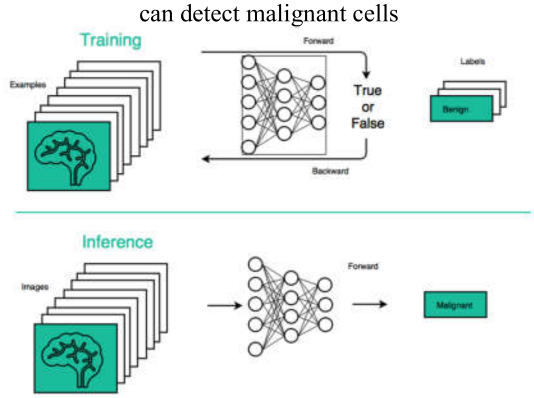

Source: Njenga (2016)

4. Discussion

Since we have presented a overview about human reasoning and classification, some essentials of AI and the operation of reasoning in both ontologies and ML, we are now ready to discuss some polemic topics. First of all, a brief comparison between the human reasoning (section 2) and machine reasoning (section 3) should be sketched.

As one can notice, the theories underlying human reasoning are somehow mirrored in the machines reasoning realm. At least, to human ruled-based approach one can correspond the ontological reasoning; and to human probabilistic approach one can correspond some modalities of probabilistic approaches in ML algorithms. In some sense, the human mental model approach is mimicked by the models that must be satisfied in the ruled-base approach, in which in all possible models, all given sentences have to be true.

These conclusions are not a original, but the endeavor to deep in these complex subjects is not a trivial task. We need revision and additions in future papers to include details and information. Accordingly, in the rest of this section, we only introduce some flaws and limitations of ML and ontologies that leads one to reflect whether reports we have seen in magazines, and even scientific papers, configure either hypes or a real situations.

Several reports from popular sources have showed the usefulness of ML in a diversity of situations. Markoff (2012) for example reminds us that deep learning (a modality of ML) has already been in our daily life in services like Apple's Siri virtual personal assistant and in Google's Street View. Lewis-Kraus (2016) offers a view in how ML is transforming Goggle Translate in working all the time even when humans are sleeping. On the other hand, there are also several studies rising flaws and complexities in using ML algorithms. Marcus (2018) poses that ML is just a statistical technique, and statistical techniques suffer from deviations about their own assumptions. A list of ten challenges faced by ML applications would be (Marcus, 2018):

- they are data hungry;
- they are shallow and have limited capacity in transfer tests (involving slightly modified scenarios);
- they have no natural way to deal with hierarchical structures;
- they have struggled with open-ended inferences;
- they have not sufficiently transparent, some call them "black-boxes";
- they have not been well integrated with prior knowledge;
- they cannot inherently distinguish causation from correlation;
- they presume a largely stable world, in problematic ways;
- they work as an approximation, but their answers cannot be fully trusted;
- they are difficult to engineer with.

The current paradigm of ML is dominated by what is called "agnostic deep machine learning" in which input-output pairs are employed to generate "stochastic models" in a process called "training" (see section 3.2). These models are: i) stochastic, because they maneuver data in a probabilistic sense, i.e., inputs are connected to outputs in a probabilistic fashion; ii) agnostic, because they do not rely on any prior knowledge about the task in which they are involved; iii) deep, in the sense that the architecture has multiple layers of networks of computational units (Landgrebe and Smith, 2019). To these models work properly several conditions must be satisfied, conditions that seems to be within the range just of governs and giant corporations (Landgrebe and Smith, 2019):

- A huge body of data must be available for training, in the form of input-output tuples connected according to the best human performance;
- The input-output data must be similar, i.e., inputs should lead to similar outputs, because ML requires patterns obtained in recurring processes;
- The data input must be abundant, since millions of records are need to represent the full variance.

Landgrebe and Smith (2019) still emphasize that agnostic models have no knowledge of linguistics, indeed, they have no knowledge of anything at all. Brochaussen (2018) presents a simple example of this lack of knowledge or semantics explaining the way ML algorithms manage representations of objects. Figure 3 shows three different common objects: a house, a table and a car.

However, in the lack of any semantics, ML algorithm sees these different objects are just “blue things”.

Figure 3 – three different objects recognized just as “blue things”

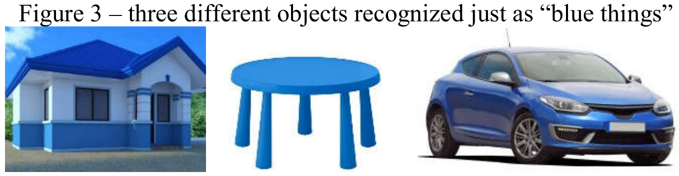

Source: Brochaussen (2018)

Some criticisms suggest that costs of these automatic processes are prohibitive, if one is waiting for accurate results. After training on large amounts of data, conventional ML approaches are frozen in time because, unlike humans, they cannot continue to learn and receive more information all the time. This causes the need of constant resources to training the algorithms again and again (Kanan, 2019).

The case of the costs seems to be the main complaint about reasoning in ontologies. The need of experts to build ontologies, which only after completely axiomatized and populated, can offer some inference, have burden projects using this technology in spite of notable exceptions (Ashburner et al., 2010). Another well-known constraint is the limited capacity of representing the world in logics and the balance between expressiveness and computability (Rector et al. 2019).

It is also no clear at this time that inferences in ontologies can be used in daily applications, for example, in the well-know decision support systems. It seems that the syllogistic reasoning they can provide is more useful for the engineers that are building them than for a lay person or an expert. Some difficulties in using inferences in ontologies are (Rector et al. 2019):

- Ontologies began as unstructured rule sets, and rule sets proved difficult to scale up;
- The reasoning path, or the so-called “justification” process, leading to inferences in OWL is very difficult to work out by manual inspection.

There are attempts to modify and to add constructs to OWL language, so that it can be more amenable for experts (Horridge et al., 2011). At least, ontologies do not suffer from the lack of semantics observed in ML algorithms.

5. Final remarks

In this paper we introduce the operation of machine reasoning and human reasoning for purposes of classification. We present some essentials of AI and of human reasoning, trying to critically understand the accomplishments of machines running AI algorithms. After this clarification, we discuss some limitations of AI approaches, admitting that much still have to be done. We admit there are many papers presenting the advantages and issues of these technologies which try to substitute human tasks, but no much initiatives in the scope of LIS. This paper intends to collaborate in this sense. We hope that the developments of AI can be useful and do not repeat what happened in the 20th, when the novelties did not prove to be feasible, frustrating the expectations of society regarding the future of the field.

Regarding issues of classification, because of all we presented so far, we believe that AI is a good assistant to human beings in a context in that large amounts of data surpass the biological capacity of the human brain. It is really difficult to imagine that, with the current technology, we can find intelligence in machines, but a brute-force approach originated in developments in memory and capacity of processing. As Searle (1980) have already proved in the last century with a simple situation, computers are far from human intelligence:

“Imagine a native English speaker who knows no Chinese locked in a room full of boxes of Chinese symbols (a database) together with a book of instructions for manipulating the symbols (the program). Imagine that people outside the room send in other Chinese symbols which, unknown to the person in the room, are questions in Chinese (the input). And imagine that by following the instructions in the program the man in the room is able to pass out Chinese symbols which are correct answers to the questions (the output). The program enables the person in the room to pass the Turing Test for understanding Chinese but he does not understand a word of Chinese.”

6. References

Apté, C., Damerau, F., Weiss, S. M. (1994). Automated learning of decision rules for text categorization. *ACM Transactions on Information Systems*, Vol. 12, n. 3, 233–251.

Ashburner, M. et al. (2000). Gene Ontology: tool for the unification of biology. *Nat Genet.* Vol. 25, No.1, 25-29.

Baaeder, F. et al. (1992). *Terminological Knowledge Representation*: A Proposal for a Terminological Logic. Retrieved March 2, 2015 from http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.41.8116

Bittner, T., Smith, B. (2008). A Theory of Granular Partitions. In Katharine Munn, Barry Smith (Eds.). *Applied Ontology: an introduction,* Ontos-Verlag, Berlin.

Braine, M. D., O’Brien, D. P. (1991). A theory of if: A lexical entry, reasoning program, and pragmatic principles. *Psychological Review*, Vol. 98, No.2, 182.

Brochhausen, M. (2018). Lecture Notes. University of Arkansas For Medical Science. https://humanities.uams.edu/faculty/mathias-brochhausen-ph-d/

Ereshefsky, M. (2004). *The Poverty of the Linnaean Hierarchy*: A Philosophical Study of Biological Taxonomy. Cambridge University Press, Cambridge.

Genesereth, M.R., Nilsson, J. (1988). *Logical foundations of artificial intelligence*. Morgan Kaufmann , San Francisco.

Gorman, M., Sanford, J.J. (2004). Aristotelian Tradition. In Michael Gorman and Jonathan J. Sanford. (Eds.). *Categories*: Historical and Systematic Essays. The Catholic Univ of America Press, Washington D.C.

Gorman, M. (2004). Categories and Normativity. In Michael Gorman and Jonathan J. Sanford. (Eds.). *Categories*: Historical and Systematic Essays. The Catholic Univ of America Press, Washington D.C.

Hastie, T., Tibshirani, R., Friedman, J. (2008). *The Elements of Statistical Learning*: Data Mining, Inference, and Prediction. New York, Springer.

Hjørland, B. (2012). Is classification necessary after Google? Retrieved March 30, 2019 from http://www.iskoiberico.org/wp-content/uploads/2014/09/19-30_Hjorland.pdf.

Horridge et al. (2011). *The Cognitive Complexity of OWL Justifications*. Retrieved August 13, 2012, from http://ceur-ws.org/Vol-745/

Jansen, L. Classifications. In Katharine Munn, Barry Smith (Eds.). *Applied Ontology: an introduction,* Ontos-Verlag, Berlin.

Jacob, E. K. (2004). Classification and Categorization: A Difference that Makes a Difference. Retrieved March 3, 2018, from https://www.ideals.illinois.edu/handle/2142/1686

Jansen, L. Classifications. In Katharine Munn, Barry Smith (Eds.). *Applied Ontology: an introduction,* Ontos-Verlag, Berlin.

Johnson-Lard, P.N., Bara, B. G. (1984). Cognition, *Syllogistic inference*, Vol.16, l-61.

Kanan, C. (2019). *Incremental Learning in Deep Neural Networks using Memory Replay*. Cognitive Science Colloquium. State University at Buffalo.

Kelleher, J. D., Mac Namee, B., D'Arcy, A. (2015). *Fundamentals of machine learning for predictive data analytics*: algorithms, worked examples, and case studies. Cambridge, The MIT Press.

Levesque, H. J. (1989). Philosophical Logic and Artificial Intelligence. In Richmond H. Thomason (Ed.). *Logic and the Complexity of Reasoning*, Springer, Dordrecht.

Lewis-Kraus, G. (2016). *The Great A.I. Awakening*. Retrieved December 22, 2018, from https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html

Marcus, G. (2018). *Deep Learning*: A Critical Appraisal. Retrieved December 22, 2018, from https://arxiv.org

Markoff, J. (2012). *Scientists See Promise in Deep-Learning Programs*. Retrieved April 13, 2016, from https://www.nytimes.com/2012/11/24/science/scientists-see-advances-in-deep-learning-a-part-of-artificial-intelligence.html

Mclay, R. (2018). *Managing the rise of Artificial Intelligence*. Retrieved November 10, 2018, from https://tech.humanrights.gov.au/sites/default/files/inline-files/100%20-%20Ron%20McLay.pdf.

Njenga, P. W. (2016). *Adapting Deep Learning to Medicine*. Retrieved April 26, 2016, from https://aws.amazon.com/blogs/startups/adapting-deep-learning-to-medicine-with-behold-ai/

Rector et al. (2019). On beyond Gruber: "Ontologies" in Today's Biomedical Information Systems and the Limits of OWL. Forthcoming: *Journal of Biomedical Informatics.*

Russell, S., Norvig, P. (2017). *Artificial Intelligence: modern approach*. New York Prentice Hall

Oaksford, M., Chater, N. (2001). The probabilistic approach to human reasoning. *Trends in Cognitive Sciences*. Vol.5 No.8, 349-357.

Saracevic, T. (1996). *Interdisciplinary Nature of Information Science*. Retrieved August 30, 2017 from http://www.brapci.inf.br/_repositorio/2010/03/pdf_dd085d2c4b_0008887.pdf

Sattler, U., Stevens, R., Lord, P. (2014). *How Does a Reasoner Work*? Retrieved August 12, 2016 from http://ontogenesis.knowledgeblog.org/1486

Searle, J. (1980). Minds, Brains and Programs. *Behavioral and Brain Sciences*, Vol. 3, 417–57

Smith, B., Brogaard, B. (2003). Sixteen Days. *The Journal of Medicine and Philosophy*, Vol. 28, 45–78.

Stalnaker, R. (1968): A theory of conditionals. In: N. Rescher (ed.), *Studies in Logical Theory*. Blackwell, Oxford.

Stenning, K., Van Lambalgen, M. (2008). Human reasoning and cognitive science. MIT Press.

Stenning, K., Yule, P. (1997). Image and language in human reasoning: a syllogistic illustration. Cognitive Psychology, 34, 109–159.

Studtmann, P. (2013). Aristotle's Categories. Retrieved March 30, 2019 from https://plato.stanford.edu/entries/aristotle-categories/

Tennis, J. T. (2016). Structure of Classification Theory: On Foundational and the Higher Layers of Classification Theory. Proceedings of thr *13th Annual Conference on Knowledge Organization*. Rio de Janeiro, Brazil.

The Guardian (2018). Facebook labels declaration of independence as 'hate speech'. Retrieved March 29, 2019 from https://www.theguardian.com/world/2018/jul/05/facebook-declaration-of-independence-hate-speech

Zielinski, T. A., Goodwin, G. P., Halford, G. S. (2010). Complexity of categorical syllogisms: an integration of two metrics. *European Journal of Cognitive Psychology*, Vol. 22,No 3, 391–421.

W3Consortium. *Web Ontology Language*. Retrieved January 3, 2019 from https://www.w3.org/OWL/