

Towards a Standardization of Computational Models of Affect: OWL and Machine Learning

Gianmarco Tuccini[†]
gianmarco.tuccini@stud.unifi.it

Luca Baronti[‡]
l.baronti@bham.ac.uk

Laura Corti^{*}
l.corti@unicampus.it

Roberta Lanfredini[†]
lanfredini@unifi.it

ABSTRACT

Computational models of affect (CMAS), in their most common form, cannot take into account the qualitative (phenomenal) dimension of affect itself. Their expressivity can be extended, thus promoting the much sought-after standardization in the most theory-neutral way, using OWL (Web Ontology Language) and machine learning techniques. OWL is an expressive formal language, as well as an established open standard, and can be used to describe the models, possibly including qualitative entities at the fundamental level. The supervised machine learning techniques allow the direct learning and application of models described as OWL ontologies. Thanks to human supervision (e.g. using datasets labeled by a human user), they can take into account the qualitative dimension of affect when the models warrant it. To further promote the aforementioned standardization, the task of classifying texts according to their affective content (known in computer science as "sentiment analysis") can be recommended as a way to assess the performance of the models. It is a multifaceted task, in which usually divergent fields as philosophy, psychology and computer science meet. Moreover, since it is a very current task in computer science, there are many resources available to facilitate the development of a standard benchmark for CMAS.

[†] Università degli Studi di Firenze, Italy.

[‡] University of Birmingham, UK.

^{*} Università Campus Bio-Medico, Rome, Italy.

1. Introduction

We live in a new era, characterized by an increasing pervasiveness of new technologies in our everyday life. Robots, AI systems and machines are going to interact with us in a deeper and more personal way, as we can imagine in the Japan societal transformation plan, called Society 5.0. This new pathway needs the development of new methods of analysis of the emotional dimension; the primary goals are to trace the emotional state of the stakeholders and to have a better interaction with the machine. Therefore, the possibility to insert emotions in AI is threefold: 1. recognition of emotions 2. Display of emotion in behaviour 3. the opportunity to build computational models of the affective domain. (Picard & others, 1995)

All in all, the general structure has three main blocks: in the first one, we aim to introduce the problem connected with affect and its artificial modelling; the second one aims to show how it is possible to create an ontology of affect; on the other hand, the last one is focused on a technical method, machine learning, which is able to take into account the complexity of the described phenomenon. In the final section, we introduce, as a case study, the analysis of texts from an affective point of view, that could be defined as sentiment analysis. We consider the text as a general word, that includes all the vehicles of content, as a piece of writing, image or speech; the text is a modality of expression in which the writer reveals affective states. (Jurafsky & Martin, 2019) Our intent is to underline not only the pure meaning of words but also the value presented in affective terms. We think that this case study is very compelling because of the increase of text interaction, in our everyday life, with AI systems like computers, vocal personal assistants, chatbot and robotics.

2. What is Affect?

Affect is a very complex phenomenon that is commonly experienced; it is an "old, Middle English (fourteenth century) word" (Ahmad, 2011, p. vii). In the dictionary, it is described as "1. The feeling of liking or loving sb/sth very much and caring about them" (Hornby et al., 2000). The definition represents the most common use but is very partial because it refers only to a positive feeling, while we consider the affect in a more general way. In the contemporary analysis, Affective computing, the field of research defined as "computing that releases to, arises from, or deliberately influences emotions", is born establishing an

ambiguity between emotion and affect because they “will be used interchangeably”; it is a comfortable assumption because the complexity of the term, so far not defined, is consigned saying that “affective will something be used in a broader sense than emotional” (Picard & others, 1995). Note that the last sentence is a shareable opinion; however, the meaning of adjective ‘broader’ is linked with an idea of spatiality, while, in order to underline the structure of the affective domain, the emphasis is on a different semantic area, using the adjective ‘general’. Consequently, we have to deeply analyze “What is affect”. This is an important question because Paul and Anne Kleinginna have found nearly a hundred definitions. (Kleinginna & Kleinginna, 1981) In our perspective, Affect can be studied as a mongrel concept, definable as macro category. This idea comes from Aristotle. He was the first philosopher who identified Pathos as one of the highest genera and exemplified it as “being cut, being burned” (Cat., 2a 3) (Ackrill & others, 1975) but also as “feeling pleasure or pain” (Cat., 11b 2) (Ackrill & others, 1975). The other kinds are, in order of appearance: substance, quantity, quality, relation, place, date, posture, state and action. In Aristotle’s reflection, we can find another book, *On the soul* (De Anima), in which the problem of affection is central both in body and in the soul. Starting from here, it is possible to find a few different affective dimensions (micro-categories) as, for example, emotion, sensation, feeling, mood. Other classifications, as WordNet-Affect, include different categories, as a cognitive state, personality, behaviour and mental attitude. (Strapparava et al., 2004) The differences between sub-categories are related to the duration, emotions are momentary while moods are enduring, and to the subjective position, emotion is linked with immediate expression and can be “single-track”, according to Deonna & Teroni (2012, 2012), while feeling is a deeper affective statement which is directed to something (or someone) outside and can be described as “multi-track”. In order to classify qualitative statements into sub-categories, emotions include joy, anger, sadness, surprise, and disgust; while happiness, fear, love, hate are feelings; otherwise, moods are, for examples, calm, anguish, arousal, activation and cheerfulness, while sensation refers to a variety of qualitative impressions made by perceptual senses. All these affective dispositions can be present at the same time; for example, I love eating ice-cream, but I felt disgusted when today I tasted a new flavour. In this case, I have a feeling, love, which is directed to something, the ice-cream; simultaneously I have an emotion, disgust, which is based on a sensation, the taste of the new flavour. Our mental life is full of multi-modal thoughts. The vagueness of the

main concept is strictly ingrained in word's etymology; nevertheless, it is an opportunity to compress the difficulties of several dimensional analysis in one that includes all even if maintaining, potentially, the differences and shades.

3. Three stances to investigate this concept

In addition, the main concept can be analyzed from a few different perspectives, from a bioregulatory aim (Damasio, 1999) to cognitive meaning (Lazarus & Lazarus, 1991). Generally, we distinguish at least two categories (Frankish & Ramsey, 2014): "low-level process model", which is based on the neurological mechanism of affect and "high-level process model", in which theories based on cognitive structure or appraisal are central. To apply our conceptual proposal to all these perspective, we suggest a more fine-grained distinction, based on Dennett's three stances. The philosopher Daniel Dennett introduced three different levels to comprehend an object's behaviour varying the abstraction (Dennett, 1989). The lower level is the physical stance, that is the concrete domain of physics and chemistry. All the things that happen in this level are ruled by physical or chemical rules. The second order, design stance, is more abstract and comprehends the domain of biology and engineering. In this point of view, the main idea is the teleological purpose of the system's design. The last level is the intentional stance, the most abstract domain of software and mind. The explanations, made in this section, are based on meaningful states of mind, which are vehicles of beliefs and desires. So we suggest to include the physical and design stances in the "low-level process model", while the intentional stance in the higher category. Nevertheless, the case study of affect proves the huge problem to understand the role and the position of the qualitative dimension of the affective domain. For this reason, it is necessary to speak about qualia, plural form of quale, a philosophical term that implies a subjective experience. In literature, other expressions describe this concept, as 'raw feel' or "what is it like to be" (Nagel, 1974). The main characteristics of this property are: intrinsic, qualitative, subjective, meaningful and not reducible, for an ontological issue, to a physical or neural reaction. So there is something more that we feel when we have experience. The best way to describe this untouchable, but the strong and crucial feeling is to add a very famous quotation from Jackson (1982, 1982). He wrote the Knowledge Argument for Qualia. In this hypothetical situation, a brilliant scientist specialized in the neurophysiology of vision, Mary, has always investigated the world "from a black and white room via a black and white

television monitor”. She acquired all the physical information, and she is able to use qualitative terms, like ‘yellow’, ‘blue’; “she discovers, for example, just which wavelength combinations from the sky stimulate the retina, and exactly how this produces via the central nervous system the contraction of the vocal cords and expulsion of air from the lungs that results in the uttering of the sentence ‘The sky is blue’ ”(Jackson, 1982, p. 130). Nevertheless, if she goes out or a color monitor is given, something will happen because, as Jackson said, “It seems just obvious that she will learn something about the world and our visual experience of it. But then is it inescapable that her previous knowledge was incomplete.” (Jackson, 1982, p. 130). In our opinion, Mary has two different kinds of knowledge: the first one is linked with an analysis of data, made during the acquisition through a television monitor, and the second one is strongly connected with experience, made in the real world, which produces an affective statement. It is not a question of monitor’s mediation if the processes of knowledge are different, but it relates to a deeper meaning. In the first case, we have a close system that produces output changing unknown variables; in contrast, the experience opens us not only to a set of data but also to interaction not predictable that changes the subjective structure; in other words, there is something that Mary feels when she sees the sky for the first time. To conclude, we do not want to give a judgement, axiologically connoted, but it is fundamental to distinguish the analysis and manipulation of data from experience. Is it possible to investigate our crucial topic with these stances? We respond affirmatively because the theories of affect are focused on different levels of abstractions. Scientific results prove that certain affective statements produce hormones, like Endorphins; other authors, as Damasio or James, find a biological meaning for emotional response. It is also obvious, according to what we have said before, that affect could be explained through intentional stance.

4. The possibility of CMAS, computational models of affect

Consequently, the main problem is how we can build a standard computational model, that is a closed procedure based on the Turing’s idea of effective calculability (Turing, 1937) for a dynamic concept that is analyzable from different points of view. Turing’s machine is supplied with a tape divided into squares; in each of them is written a symbol. At any moment there is only one square “in the machine” and the possible actions, according to a finite set of rules, are: changing position or changing symbols written. This is a symbolic

manipulation (Searle, 1980) that allows mimicking the human capacity for computing. Another important difference is between Strong and Weak Artificial Intelligence (Russell & Norvig, 2016). The first form, strong (or full or general) AI, is based on the idea that a product of AI replicates completely the human cognitive abilities, for, according to this perspective, a machine thinks. Otherwise, the weak (or narrow) A.I is linked with simulation and the assumption that “machines could possibly act intelligently (or, perhaps better, act as if they were intelligent).” Significant studies have been done on affect in many fields, like biology, psychology and philosophy; so is a computational model able to comprehend all the four levels identified before? The physical stance is dominated by rules that can be formalized into a computational model without problems. The second level, design stance, is based on teleological purpose, even if we consider the behaviour in accordance with conditions; so it is, however, possible to mutate in computational rules. The third level of abstraction, intentional stance, is linked with an assumption of beliefs and desires causally connected with the behaviour; it is also possible to write code based on causal properties. In this last stance, it is possible to include also the affective domain, which is connected with the analysis of behaviour but also with the “lived experience”. Therefore, because of the complexity of the affect, we need to find a new pathway to create a computational model. We propose machine learning as a method to assess the qualitative content of expressions. In fact, from the interaction between person and computer a quality, that produces an evaluative judgment, emerges. So rules are not created to run a program, but they are obtained by training. For these reasons, this new proposal is functional to maintain the complexity of affective statements instead of simplifying (Bosse et al., 2010) and it is possible only because of training of Human Intelligence. One possible theoretical framework is linked with the sensorimotor approach, a philosophical and psychological perspective developed in the 21st century by the philosopher, Alva Noë, and the psychologist, Kevin O’Regan. The most important element of this approach is a new conception of sight because they “propose that seeing is a way of acting” (O’Regan & Noë, 2001) instead of vision as an inner representation of the outside environment. They emphasize the explorative aspect of seeing as if we were in front of a peculiar examination by touching, suggesting an alternative way called “sensorimotor contingency theory”, in which the rules are regulated by “the fact that exploration is being done by the visual apparatus”. We conclude, in accordance to them, that “perception is a complex concept that implies different modalities, governed by

laws, for different sensory domains, as vision, smell, tasting etc.” Underlining active attitude of the subject in perception breaks the ordinary common sense of passive receiving of qualitative instances from the world, that becomes, in accordance with what we have said before, an external memory “that can be probed at will by sensory apparatus” (O’Regan & Noë, 2001). How is it possible to provide an explanation of qualia in the sensorimotor approach? “My own position is that regardless of whether qualia are a misguided notion, there’s no denying that a lot of people talk about their experiences as though there are ‘raw feels’ causing them. We need a scientific explanation for what people say about them, independently of whether qualia make sense.” (O’Regan, 2011). Kevin O’Regan identifies four mysteries for raw feeling: it feels like something, that cannot be reduced to a brain function in order to avoid the problem of infinite regress; it has different qualities, for instance, the cold of ice, the sound of the piano, the taste of pizza, the redness of red and so on; there is structure in difference, “in the sense that certain feels can be compared, but others cannot” ; and, raw feel is ineffable, because it cannot be communicated. These problems are solved in a sensorimotor approach which is based on four characteristics: bodiliness of the experience, richness, partial insubordinateness and grabbiness of the world. According to Kevin O’Regan, sensorimotor approach can be used to attribute qualitative statements to products of AI because the feeling is strictly integrated in interaction between agent and environment. In fact, “Imagine you are squeezing a sponge and experiencing the feel of softness. What brain mechanism might generate the softness feel? The word generate seems inapplicable here, and so this seems to be the wrong kind of question to ask about softness. Softness is a quality of the interaction you have with a soft object like a sponge. Feeling softness consists in the fact of your currently being in the process of noting that as you press the sponge, it squishes under your pressure. Having the feeling of softness does not occur in your brain; rather, it resides in your noting that you are currently interacting in a particular way with the sponge.” (O’Regan, 2011). All in all, the general structure has two main blocks: in the first one, the aim is to show how it is possible to create an ontology of affect; whereas the second is focused on a technical method, machine learning, which is able to take into account the complexity of the described phenomenon. One of the main practical applications is the analysis of texts from an affective point of view, sentiment analysis. As we have affirmed in the introduction, the text is a way to express the writer’s affective statement. According also to Valitutti et al. (Valitutti, 2004, p. 61) “text is an important modality for sensing

affect because the bulk of computer user interfaces today are textually based". Our intent is to underline not only the pure meaning of words but also the value presented in affective terms through the hermeneutic power of machines.

5. The case for OWL 2

At the moment there are many theories of emotions. *A fortiori*, there are even more theories of affect, in the sense defined above. Many of these theories are not yet described in a formal language and thus cannot be easily translated in computational terms, let alone implemented in computer programs. It is our contention that, if all the theories were described in the same formal language, researchers would quickly identify the overlap between them and soon would reduce the fragmentation of the theoretical landscape to an acceptable, fertile level: the methodological standardization would promote a partial content standardization. This would also make it easier to translate the theories in computational models and to choose the most adequate computational model to apply to any practical problem in which affect is involved, e.g. in human-robot interaction.

The aforementioned standardization requires at least the adoption of a very expressive formal language, in which all the existing theories – developed from the physical, design and intentional stance – can be adequately formalized. In our opinion, this language should preferably be an open, established standard, already known and widely used across disciplinary boundaries, in order to avoid vendor lock-in, reuse the existing tools and maximize the interdisciplinary collaboration.

OWL 2 (Web Ontology Language, from now on just OWL) (Motik, Patel-Schneider, & Parsia, 2012), which we'll describe¹ below, seems to satisfy all the requirements.

¹ Unless differently noted see Motik, Patel-Schneider, & Parsia (2012), *passim*, for reference on this section. See also Hitzler et al. (2012), which provides an accessible introduction to OWL 2. It's worth noting that OWL 2 belongs to the Semantic Web stack. Other technologies that belong to the same stack may be used to build formalised descriptions of theories of affect, such as unconstrained RDF and EmotionML (Schröder et al, 2011). We focus on OWL 2 because, as the rest of this section will show, it is expressive but strict in its semantics and thus it allows the creation of ontologies that are always comparable and may be used for reasoning without conversion or mapping steps.

6. OWL 2

OWL is a general purpose modelling language which is used to create ontologies: formal descriptions of the knowledge about a subject matter.

OWL ontologies contain entities, expressions and axioms.

Entities refer to something in the world, are named and come in four types. First, there are individuals: concrete objects² such as “John” or “this chair”. Second, there are classes: sets of individuals such as “people” or “chairs”. Third, there are datatypes: sets of literals³ such as “integers”. Fourth, there are properties: object properties, which are relations between two individuals such as “married to”; data properties, which are “relations” between individuals and literals such as “age”; and annotation properties, which are meta-descriptions of entities like “defined by”.

Expression is (optionally) named recursive combinations of entities and expressions. Expressions are created through a wide range of constructors. In particular, there are set-theoretic constructors, which allow creating expressions such as “The set ‘wives’ is the intersection of the sets ‘spouses’ and ‘women’”, “‘parents’ is the union of the sets ‘mothers’ and ‘fathers’” and “The set ‘even numbers’ is the complement of the set ‘odd numbers’”. There are also constructors which impose restrictions on the cardinality of individual properties, thus allowing to create expressions such as “Pet lovers have at least a pet”, “Dog lovers have only dogs as pets” and “Polyglots know at least two languages”. Table [tab:constructors] lists some of the available constructors.

OWL constructors

<i>ObjectIntersectionOf</i>	<i>[mothers is the] intersection of women and parents</i>
<i>ObjectUnionOf</i>	<i>[parents are the] union of mothers and fathers</i>
<i>ObjectComplementOf</i>	<i>[graduates is] the complement of undergraduates</i>
<i>ObjectSomeValuesFrom</i>	<i>[persons] have at least a child</i>
<i>ObjectAllValuesFrom</i>	<i>[good parents] have only happy children</i>
<i>ObjectHasValue</i>	<i>[John’s children] have John as parent</i>
<i>ObjectHasSelf</i>	<i>[narcisist] love themselves</i>
<i>ObjectMaxCardinality</i>	<i>[John has] at most 4 children</i>

² Anonymous individuals, which OWL allows in some context, are no entities.

³ Literals consist of strings such as “1”, called “lexical forms”, which are interpreted according to their datatype. Literals are not entities in themselves.

ObjectMinCardinality [John has] at least 2 children
ObjectExactCardinality [John has] exactly 3 children
ObjectOneOf [color is] on of black and white

Axioms are statements about the subject matters which the ontologies assert are true. In particular, axioms allow stating the relationships between individuals, as in "John is married to Mary"; the relationships between classes, as in "Dogs are a subclass of mammals"; the relationships between individuals and classes, as in "John is a person"; and the relationships between individuals and literals, as in "John is 18 years old". Table [tab:axioms] lists some of the available axioms (Power & Third, 2010).

OWL Axioms

SubClassOf	Every admiral is a sailor
EquivalentClasses	An admiral is defined as a person that commands a fleet
DisjointClasses	No sailor is a landlubber
ClassAssertion	Nelson is an admiral
ObjectPropertyAssertion	Nelson is victor of the Battle of Trafalgar
DataPropertyAssertion	The Battle of Trafalgar is dated 1805
ObjectPropertyDomain	If X commands Y, X must be a person
ObjectPropertyRange	If X commands Y, Y must be a fleet
SubObjectPropertyOf	If X is a child of Y, X must be related to Y
InverseObjectProperties	If X is a child of Y, Y must be a parent of X
TransitiveObjectProperty	If X contains Y and Y contains Z, X must contain Z

OWL ontologies describe the knowledge about subject matters both explicitly, as shown above, and implicitly, through entailment. That is if the axioms of an OWL ontology explicitly state that "English setters are dogs" and "dogs are mammals" it is also known, *qua* entailed, that "English setters are mammals". Entailment can be used in reasoning tasks, e.g. to check whether an individual is an instance of a certain class or expression. Entailment is strictly related to deduction. Since it is possible to look at computation as a form of mathematical deduction (Kripke, 2013), it follows that ontologies can be seen as declarative computational models.

OWL has two different semantics, which defines marginally different rules for entailment.

The RDF-Based semantics (Schneider, 2012), informally OWL Full, views OWL ontologies as RDF graphs⁴, and allows the unrestricted use of OWL constructs. Under the RDF-Based Semantics, though, there are undecidable reasoning tasks, e.g. it is not always possible to tell whether an individual is an instance of a certain class or expression.

The Direct-Semantics (Motik, Patel-Schneider, & Cuenca Grau, 2012), informally OWL DL, provides a meaning for OWL ontologies in the description logic *SROIQ(D)* (Horrocks et al., 2006),⁵ An expressive decidable fragment of first-order logic. The Direct-Semantics grants other useful computational properties to OWL Ontologies, apart from decidability, but imposes some restrictions on the use of OWL constructs. In particular, individuals and classes are to be treated as disjoint sets: no individual can also be a class and vice versa. All things considered, these restrictions are a price worth paying in order to gain the aforementioned useful computational properties. Moreover, there is an extensive literature on description logics, which by adopting the Direct-Semantics immediately becomes available for building better OWL ontologies. Therefore, from now on we'll assume the adoption of Direct-Semantics.

OWL has also three different profiles: subsets of OWL which trade expressivity for efficiency in specific tasks.

OWL EL (Motik, Cuenca Grau, et al., 2012) is based on the description logic *EL⁺⁺* (Baader et al., n.d.). It is rather efficient and thus it is particularly suitable for large ontologies, at the price of disallowing useful expression constructors such as union and property cardinality restrictions. We'll describe OWL EL in more details below.

OWL QL (Motik, Cuenca Grau, et al., 2012) is based on the description logic *DL – Lite* (Calvanese et al., 2007). It is suitable to build a semantic layer above relational databases, by rewriting OWL queries in SQL queries, thus promoting the integration with existing datasets. It also disallows useful expression constructors such as individual equality assertions such as “Hesperus is Phosphorus”.

OWL RL (Motik, Cuenca Grau, et al., 2012) is based on Description Logic Programs (Grosz et al., 2003), i.e. the intersection of OWL and logic programming. Logic programming is based on rules, also called clauses, which

⁴ It is outside the scope of this paper to introduce RDF. See Schreiber & Raimond (2014) for an accessible introduction to RDF

⁵ It is outside the scope of this article to introduce description logics. See Baader et al. (2010) for an accessible introduction.

can be read as logical implications: “ $father(x, y) \wedge brother(y, z) \Rightarrow uncle(x, z)$ ”. OWL RL builds a bridge to logic programming through a partial axiomatization of the RDF-Based semantics in the form of first-order implications. This axiomatization allows OWL RL to be implemented using rule-based technologies. It also allows OWL RL ontologies to be converted to RIF (Rule Interchange Format) (Morgenstern et al., 2013; Reynolds, 2013), which many rule systems can then use.

It is important to note, in this regard, that there exist alternative approaches to the integration of rules in OWL. SWRL (Semantic Web Rule Language), for instance, extends OWL to allow the direct use of rules (Horrocks et al., 2004). In a sense, while OWL RL is the intersection of OWL and logic programming, SWRL aims to be their union. The resulting language is very expressive but unfortunately is undecidable. SWRL becomes decidable (but also less expressive) only if some restrictions are imposed to its rules, e.g. by requiring every individual present in rules to be explicitly declared in the ontology (Motik et al., 2005). It is also possible to stretch the expressivity of OWL by encoding rules using advanced constructs such as property chains and keys. Current works show that $SROIQ(D)$ is expressive enough to allow the encoding of many rules using the aforementioned techniques (Hitzler et al., 2009; Krötzsch et al., 2008). Thus, this approach seems to be at least viable for many tasks, including the modelling of affect.

Finally, it’s worth mentioning that OWL ontologies are modular: an ontology can import and reuse other ontologies. In particular, it is common for OWL ontologies to import upper ontologies such as the Basic Formal Ontology (Arp et al., 2015) or the General Formal Ontology (Herre, 2010), which provide very general and common terms. This pattern promotes at least a minimal semantic interoperability between ontologies from different domains and disciplines.

7. Modelling with OWL 2

The expressivity of OWL, which already emerges from the above description, can also be shown by analyzing the ontologies in current use at the relevant level of abstractions. It is also appropriate to assess whether it is possible to represent

in OWL the most relevant anatomical, dimensional and appraisal theories of emotions.⁶

8. Physical stance

Although at the best of authors' knowledge, computational models of affect built at such a low level do not exist, it is possible in principle to use OWL ontologies from the physical stance.

For instance, the OPB (Ontology of Physics for Biology) (Cook et al., 2011) provides a faithful representation of physical entities and processes relevant for biology. Thus, it allows representing in an explicit way the thermodynamics and dynamics of physiological processes which occur in organisms.

Villanueva-Rosales & Dumontier (2007), instead, have provided an OWL ontology of functional groups for the classification of chemical compounds. As they write, "The ontology is suitably expressive to provide precise logic-based descriptions that match well defined chemical substructures, providing evidence that semantic web technologies are sufficient to represent and reason about the chemistry domain". The ontology is not intended as a mere taxonomy of functional groups. Instead, the authors have built it with the aim of allowing reasoning over the ontology entities.

9. Design stance

OWL ontologies are widely from the design stance, e.g. in biology, medicine and neuroscience.

Of special interest is the NIFSTD (Neuroscience Information Framework Standard Ontology) (Bug et al., 2008), which integrates many relevant ontologies in order to provide a unified, consistent terminology for the whole neuroscience domain. While some of the integrated ontologies provide just the terminology for specific subdomains, others, such as SAO (Larson et al., 2007),

⁶ In the rest of this section, we will mention existing OWL ontologies only insofar as they help to exemplify the various modelling stances, with no pretence at completeness: see e.g. Abaalkhail et al (2018) for a more systematic survey. Please note that, in particular, we will not mention notable and useful ontologies such as MARL (Westerski et al 2013) and Onyx (Sánchez-Rada et al 2016), that model affect with a focus on *opinions* and thus include classes such as "Algorithm" and "UserAccount".

exploit much of the expressivity of OWL and allow to draw useful inferences through reasoning.

It's worth mentioning that Damasio's anatomical theory of emotions (Damasio, 1999) has previously been formalized in (Bosse et al., 2008) using a temporal rule-based logic language (Bosse et al., 2005). We have seen above that it is possible to encode many rules in OWL using advanced constructs, even though there surely exist rules which cannot be encoded in this way and would require to extend the language. The situation is similar with time: there exist ontologies, such as the Time Ontology for OWL (Cox & Little, 2004) and the OMG Date-Time Ontology (Linehan et al., 2012) (see also (Welty & Fikes, 2006)), which allow to represent and reason over complex temporal relationships between entities. When they are not enough, extensions of the language such as tOWL (Milea et al., 2012) are needed. As far as we can see, it is possible to translate Bosse's formalization of Damasio's theory or to develop a new formalization using OWL constructs, at worst by adopting some of the aforementioned extensions. This is important, because Damasio's theory is surely one of the most relevant anatomical theory of emotions in affect literature.

It's also worth mentioning the Emotiono (OWL) Ontology (Chen et al., 2015), which is used for electroencephalogram-based emotion assessment. The ontology represent inference rules obtained from the analysis of raw EEG data and allow to map EEG data to emotional states, represented as valence-arousal levels. This is also relevant because it shows that it is possible to use OWL to formalize dimensional theories of affect.

10. Intentional stance

OWL ontologies are also used from the intentional stance.

Of special interest is the MFO (Mental Functioning Ontology) (Hastings, Smith, et al., 2012), which aims to describe human mental functioning in order to facilitate data aggregation across disciplinary boundaries and to allow reasoning over it. A submodule of the MFO, the MFO-Emotion (Hastings, Ceusters, et al., 2012), extends the ontology to the affective domain whilst keeping its useful properties.

It's also worth mentioning Broekens set-theoretic formalizations of cognitive-appraisal theories (Broekens et al., 2008), i.e. theories which assume that emotions are caused by beliefs about events which occur in the environment: by far the dominant theories in computational modelling of affect.

Broekens formalizations are relatively simple: that is, they only use basic constructs of set-theory. Since OWL classes are indeed sets and SROIQ(D) is an expressive fragment of first-order logic, in which the whole set-theory can be formalized, we think that Broekens formalizations can be encoded in OWL. Obviously, this should be proved rigorously. If it were true, anyway, it would mean that OWL is expressive enough to encode cognitive-appraisal theories: a significant result.

To further strengthen this point, it's worth noting that, according to Reizenzein et al. (2013), cognitive-appraisal theories can be formalized in agent logic, especially in those related to the BDI (belief-desire-intention) paradigm. One of the most famous agent logic related to BDI is AgentSpeak(L) (Rao, 1996), an agent logic based on first-order logic. Some years ago AgentSpeak(L) has been reformulated as AgentSpeak(DL) (Moreira et al., 2005), replacing first-order logic with the description logic *ALC*, a predecessor of *SROIQ(D)*. Indeed, the author aimed to allow the use of OWL ontologies as representation of the agents' knowledge.

Finally, Reizenzein et al. (2013) also propose to implement computational models of affect based on cognitive-appraisal theories in BDI frameworks. AgentOWL (Laclavik et al., 2012), which has successfully integrated ontologies in the BDI framework JADEX, shows that this approach is feasible whilst using OWL.

11. Beyond the intentional stance?

To the best of authors' knowledge, there aren't OWL ontologies built explicitly with the aim of going beyond the intentional stance to take in account quality. This is not worrying nor necessary: OWL ontologies built from the intentional stance can already use terms whose meaning is intended as irreducibly qualitative. The key to take into account quality, when we conceive it as irreducible, is using a bottom-up methodology, e.g. by learning the ontologies through machine learning techniques from labeled datasets. Even though OWL ontologies built bottom-up from the intentional stance can only approximate their subject matter in order to solve specific problems – in other words, they can only be *weak* computational models of affect – they naturally take into account the ability of humans to have irreducibly qualitative experiences and recognize their expression. Thus, we contend that the use of machine learning techniques to build bottom-up OWL ontologies is worth exploring.

12. Building Computational Affect Models Using Machine Learning

Supervised Machine Learning (Mitchell, 1997) techniques are usually used to approximate a *target* function on a given task. In this context, we use the term *function* in a broad way, namely a relation between the patterns (i.e. the examples in a given dataset) and their relative target values. Given a marked dataset, the aim of the supervised learning is to find a hypothesis (i.e. a mathematical model) using a learning algorithm, able to approximate the target function. As an approximation, these machine learning techniques are often used when the target function that originally generated the target values, given the patterns, is unknown, computationally expensive or not even possible to express in formal terms.

In our case, a Machine Learning technique aims to approximate the human affect expression motivated by qualitative experiences, which reflect the ontology by the markers. We believe that the problem of approximating the human affect expression, as a target function, is a good candidate to be tackled with machine learning problems, provided the markers in the ontology are a good representation of it.

13. The Ontology Learning Problem

Similar to the Inductive Logic Programming (Nienhuys-Cheng & De Wolf, 1997), the ontology learning problem finds a definition of a class *A* which on the one hand is specific enough to include most of similar instances in a given ontology, excluding at the same time most of the others, and on the other hand generic enough to include similar instances *not* present in the ontology.

An approach followed by many techniques defines the hypothesis space as a set of candidate classes, and learning the optimal class *A* combining a (i) refinement operator with (ii) a search heuristic. The former is used to build the search tree, whilst the latter is used to decide which node expands.

For instance, an OWL ontology can be coupled with datasets of examples, composed by individuals of that ontology, labeled as *positive* or *negative* instances of a given *task*. Machine Learning techniques can then be used to find an OWL class expression, of which most positive examples are instances whilst most negative examples are *not* instances. A key aspect is that the found class is able to effectively include *other* examples not present in the given datasets. The

model ability of generalizing the task is reflected in its accuracy in successfully classifying unknown data.

14. Learning Algorithms for OWL Ontologies

A family of learning algorithms (Lehmann & Hitzler, 2010a) approaches the problem as a search of the best concept definition in a *quasi-ordered* space (Σ, \preceq) with Σ set of concepts in the given ontology and \preceq reflexive and transitive ordering operator. In this approach, two *refinement operators* are considered:

$$\rho(\alpha) \subseteq \{\beta \in \Sigma \mid \beta \preceq \alpha\} \quad \delta(\alpha) \subseteq \{\beta \in \Sigma \mid \alpha \preceq \beta\} \quad \forall \alpha \in \Sigma$$

with ρ and δ *downward* and *upward* operators, used to traverse toward more specific and general, respectively, concepts. In particular, downward refinement operators are commonly (Nienhuys–Cheng & De Wolf, 1997) used to generate progressively more specialized hypothesis in the search space. Therefore, the properties of the downward operator ρ are of particular interest in our case.

Amongst the desirable properties that a downward operator ρ should have, it is possible to define ρ as:

finite

iff its application $\rho(\alpha)$ to any α is finite;

proper

when the operator applied to a concept always generate another different concept, namely $\rho(\alpha) = \beta \Rightarrow \alpha \not\equiv \beta$;

complete

iff $\beta < \alpha$ implies that exists a *refinement chain* of repeated downward applications $\rho^*(\alpha) = \beta'$ that leads to $\beta' \equiv \beta$;

non-redundant

iff for any $\alpha, \beta \in \Sigma$ such as $\beta < \alpha$ exists only one refinement chain that leads to β from α , namely $\rho^*(\alpha) = \beta$;

From a theoretical standpoint, the definition of those operators faces limitations (Lehmann & Hitzler, 2010b) to the useful properties that they can hold. For this reason, learning algorithms that actually implement refinement operators to traverse the ontology, and narrow down the set of candidate concepts, make use of a set of heuristics and algorithmic strategies to mitigate this problem.

The OWL Class Expression Learner (OCEL) (Lehmann, 2010) makes use of a complete and proper refinement operator to build the search tree, using heuristics to decide which node to expand. These heuristics are tuned with score values associated with each concept and updated dynamically. More advanced techniques, like the Class Expression Learning for Ontology Engineering (CELOE) (Lehmann et al., 2011) extend the OCEL heuristics to make the learning process more suited for classification problems. It also implements strategies to promote the finding of short class expressions, with the aim to adapt the algorithm to the ontology engineering scenario. A refinement operator proved to be finite, proper and complete for the *EL* description logic is used by the *EL* Tree Learner (ELTL) (Lehmann & Haase, 2009), which also adopt strategies that make the generation of refinements very efficient.

The description logic *EL* was designed to trade (Baader et al., 2005) a lesser expressive power for useful computational properties. In its simplest formulation, *EL* relies (Kazakov et al., 2014) on the top concept \top , conjunctions $C \sqcap D$, and existential restrictions $\exists RC$ as concept constructors. Using these constructors it's possible to build ontologies in many cases, however the absence of more advanced constructs may pose serious limitations in dealing with more complex domains. The major advantage of *EL* is that the reasoning can be performed in polynomial time, making the application of many operators a tractable problem. Despite its limits in the expressive power, *EL* or slight extensions can be successfully used to encode ontologies in many important cases, such as the Systematized Nomenclature of Medicine (Spackman, 2000), the Gene Ontology (Ashburner et al., 2000) and the Galen medical knowledge (Rector & Horrocks, 1997). One of it's most known extension *EL*⁺⁺ (Baader et al., 2005) is one of the main profiles of the OWL 2 standard.

Finally, the Inductive Statistical Learning of Expressions (ISLE) (Bühmann et al., 2014) technique, extends the ETEL technique using text corpus as external source. The key heuristic used in ISLE is based on the idea that the more two words are related, the more they tend to concur in the texts present in the corpus. Statistical information about the words in the corpus is then used among the other search heuristics to improve the algorithm performance.

The two key aspects that differentiate this family of algorithms are the type of refinement operators and heuristics used. Solutions like ELTL and its variants restrict the scope of the semantics on which they operate (i.e. *EL*) in order to use an ideal operator. Other algorithms are able to deal with more expressive

semantics albeit with less ideal operators. Generally, the more a semantic is expressive, the more is difficult to find a refinement operator with useful properties. However, depending on the language logic of interest, other operators can be designed to deal with different types of constructs, present in different variants of OWL. As for the heuristics variants used to expand the nodes of the search tree, the score value used is based on the accuracy, precision or recall of the given class on the ontology.

15. DL-Learner

A framework that implements all the techniques described, among others not mentioned here, is DL-Learner (Bühmann et al., 2016), an open source tool specifically designed for class expression learning.

Its multi-purpose nature can be used to tackle different scenario. DL-learner can be used to make suggestions of similar class definitions, based on the classes already present in a given ontology, to be added as *owl:equivalentClass* axiom. When used with the CELOE learning algorithm, shorter and more easily readable definitions are expected to be found more frequently. The found class definitions can also be used to retrieve similar instances efficiently. Finally, the class definitions can be used to classify unknown instances (i.e. assess if an instance not present in the known ontology is part of a certain class).

DL-Learner is able to acquire an ontology either by file or by a SPARQL endpoint. Its internal structure is modular, composed by the reasoner (e.g. Pellet), the learning algorithm and operator as separate components. The tool is designed to be easily extended, adding new features such as new learning algorithms or new refinement operators which are integrated seamlessly in the framework.

The classification feature has been successfully tested in a problem of carcinogenesis prediction, where the tool achieved (Lehmann & Hitzler, 2010b) competitive results compared state of the art algorithms. In another case, DL-Learner was used (Salguero & Espinilla, 2016) in the sentiment analysis domain, where the tool was used to learn class expressions describing documents with a positive opinion, on a given topic. This field is of particular interest in our case, being closely related to the affect domain, and one application case of wide interest among different fields, and it will be briefly illustrated in the next section.

16. Example of Use Case: Sentiment Analysis

One of the most natural ways (Bellegarda, 2014) for a human to interact with an AI is through natural language. Several frameworks (e.g. (Brooks et al., 2012; Williams et al., 2015)) have been devised to allow an interaction able to effectively convey the intentions of a human operator, in the form of specific instructions, that are translated into computable operations. These intentions are often partially expressed with affective elements embedded in the language used. Therefore, the NLP (Natural Language Processing) and the analysis of the emotions expressed in sentences can be of primary importance to discern the true meaning of human-issued instructions. Similarly, including affective elements in AI-provided replies can benefit the interaction, being perceived as more natural for the human operator.

Sentiment Analysis, also called Opinion Mining, is “a field of study that aims to extract opinions and sentiments from natural language text using computational methods.” (Liu, 2015, p. XI). In the most frequent sense, sentiment analysis is intended as the extraction of sentiment, which is “the positive or negative orientation that a writer expresses toward some object” (Jurafsky & Martin, Naive Bayes and Sentiment Classification, 2019) This research had grown since early 2000 when social media and web 2.0 became a relevant phenomenon.⁷

In order to analyze the affective dimension of the text, an accurate analysis should include two different steps: a psychological (or philosophical) inquiry about the affective dimensions and linguistic analysis of text structure to create a computational program. Nevertheless, the complexity of human expression reveals NLP problems particularly. (Cambria, Poria, Gelbukh, & Thelwall, 2017) According to Cambria, Poria, et al. it is possible to find 15 NLP problems divided into three layers: syntactics, semantics, and pragmatics.

⁷ In the rest of this section, we will describe sentiment analysis, its problems and its techniques from a somewhat abstract and traditional point of view. We aim only to show a) how and why OWL ontologies and sentiment analysis may be intertwined, and b) why sentiment analysis may prove itself useful to compare computational models of affect represented through OWL ontologies. We want to note, though, that in the latest years deep-learning approaches based on transformers (Vaswani et al, 2017) have proved themselves performant in tasks such as machine translation. Models based on these approaches learn a very general representation of language and have been specialized (“fine-tuned”) in order to perform different tasks, such as predicting the sentiment of a text, with state-of-art accuracy

The syntactics level analyzes the linguistic structure of the sentence; in particular, it provides five steps: 1) the normalization of microtext, characterized by “relaxed spelling and reliance on abbreviation, acronyms and emoticons” (Cambria, Poria, Gelbukh, & Thelwall, 2017); 2) the SBD (Sentence Boundary Disambiguation), a phase that deconstructs texts into sentences; 3) the POS (Part-of-Speech Tagging), which labels each word in parts of speech (e.g. nouns, verb, adjectives); 4) Text Chunking finds the “nonoverlapping groups of words that represent a syntactic unit” (Cambria, Poria, Gelbukh, & Thelwall, 2017); 5) Lemmatization, the last syntactics phase in which the words are converted into a base form.

The semantic layer is responsible for the analysis of concepts obtained by the normalized text. In particular, it improves the accuracy of sentiment classification in the following aspects: 1) the Word Sense Disambiguation because different meanings could emerge in different contexts; 2) Concept Extraction, that is the ability to deconstruct the text in concept, a “key for a semantic-aware analysis of text” (Cambria, Poria, Gelbukh, & Thelwall, 2017); 3) NER (Named Entity Recognition) which organizes the different named entities into categories; 4) the Anaphora Resolution in which is determined the antecedent, or the entity to which an anaphor refers; 5) Subjectivity Detection, which locates and isolates the neutral content of a text.

The last layer, the pragmatic one, aims to recollect the previous analysis investigating the complexity of texts. It is divided into five sections: 1) Personality recognition, which determines the type of personality “detecting the presence or absence” of psychological traits; 2) Sarcasm Detection, because sarcasm can completely change the polarity and the meaning of a text; 3) Metaphor Understanding, that is necessary to the evaluation of polarity and aspect extraction; 4) Aspect Extraction, which reveals different polarities and opinion in a text; 5) Polarity Detection, which is considered as the “most popular sentiment analysis task” (Cambria, Poria, Gelbukh, & Thelwall, 2017) that starts from a binary classification in positive and negative and is going to a more detailed analysis, which measures the intensity of polarity.

For instance, note the example made before “I love eating ice-cream, but I felt disgusted when today I tasted new flavor”; it is a very complex expression because there are two main complete sentences, ‘I love eating ice-cream’ and ‘I felt disgusted’, and coordinating conjunction, ‘but’, that joins sentences that are opposite in idea. To classify it, we need to consider the multi-modal hypothesis made in the first section; otherwise, we have no chance to assess the

relevance of each part of the sentence. We evaluate “love” as an enduring feeling and “disgust” as a temporary emotion based on a new experience. So, it is more relevant to the perturbation than the background condition. Consequently, the overall polarity is negative because the central part of the sentence is “I felt disgusted” and has a strong negative connotation. The disgust or love is not included in this sentence, but we are, however, able to evaluate the content. This is possible because expression emerges and is strictly associated with a qualitative statement even if there is no causal reason which produces an effect, the expression of disgust. We can adopt the same approach to the analysis of images. For example, we are in front of a picture of a child, which is laughing and crying at the same time. In this case, our interpretation is based on folk psychology, in which human behavior is explained by intentional states, as belief or desire. We assume that probably the child feels the situation hilarious; therefore, the polarity is positive.

In our opinion, sentiment analysis is a good case of study because it includes different shades: the necessity to consider the meaning in a wider sense, connected to the pure semantics but also to value; the possibility to create a computational model of affective dimension and the ability to judge and predict. On the other hand, “The automatic analysis of online opinions, however, involves a deep understanding of natural language text by machines, from which we are still very far. Online information retrieval, in fact, is still mainly based on algorithms relying on the textual representation of web pages. Such algorithms are very good at retrieving texts, splitting them into parts, checking the spelling, and counting their words. But when it comes to interpreting sentences and extracting useful information for users, their capabilities are still very limited.” (Cambria & Hussain, 2012).

The Sentiment Analysis is usually applied on different levels, such as document, sentence or aspect level. The latter aims to perform classification using a specific aspect of the entities, requiring identifying the entities and their aspects. In contrast, there is no fundamental difference (Liu, 2012) between the document and sentence analysis. One of the first steps in the Sentiment Analysis is to define and select the feature used by the classifier in the next steps. Examples of features considered (Aggarwal & Zhai, 2012) are the presence/frequency of the terms, the Part of Speech (e.g. division in nouns, verbs, adjectives, of particular interest in this context, and so on), opinion keywords or phrases and negation terms. Main approaches in the feature selection can be divided into lexicon-based methods, which heavily rely on

human work support and supervision, and statistical methods, performed mostly automatically. Amongst the most used statistical methods is possible to find Point-wise Mutual Information (Cover & Thomas, 2012) and Chi-square (Aggarwal & Zhai, 2012), which are two different ways to assess the correlation between the words and a set of categories. An unsupervised statistical feature selection technique that aims to reduce the dimensionality of data is the Latent Semantic Indexing (LSI) (Deerwester et al., 1990). This technique uses the Principal Component Analysis (Jolliffe, 1986) to transform the text space into a new system, retaining the feature with the highest value variation in the dataset.

The actual classification in sentiment analysis is commonly done using different approaches. The Machine Learning approach use well known techniques of the field such as Naive Bayes Classifier (Kang et al., 2012), Neural Networks (Ng et al., 1997; Ruiz & Srinivasan, 1999), Support Vector Machines (Chen & Tseng, 2011; Li & Li, 2013) and Decision Tree Classifiers (Lewis & Ringuette, 1994) on examples encoded with the features, found in the previous step, marked according to a set of classes. The Lexicon-based classifier approaches reduce the problem in finding the opinion lexicon, namely building a list of opinion words, phrases and idioms that are regularly used in affect instances. Usually, similar lists are generated starting from a small set of known words, then expand the list with iterative searches of synonyms or antonyms in known corpora, like thesaurus (Mohammad et al., 2009) or WordNet (Miller et al., 1990). An alternative method (Hatzivassiloglou & McKeown, 1997) allows the building of opinion lists contextual to a given corpus starting from a small set of opinion words and looking in the given corpus for additional terms related to the opinion words with the use of connectives (e.g. and, or, either, and so on). These connectives are then used to form a graph where a clustering algorithm is used to divide the words in positive/negative.

In our opinion, the integration of sentiment analysis with computational models of affect would benefit both sides. First, the corpora currently used with supervised techniques are often annotated with the polarity of the texts. There are also many corpora (especially for the english language) annotated with Ekman's basic emotions, with the triple valence-arousal-dominance commonly used in dimensional theories (Buechel & Hahn, 2017) or according to

Plutchik’s wheel of emotions (Plutchik 2001).⁸ These corpora are widely used by researchers. An analysis of computational models of affect, nonetheless, could suggest new relevant annotations, representative of researchers’ expertise, which could further improve the performance of supervised machine learning techniques in current use. Second, if the relevant annotations were available in the corpora used in sentiment analysis – or, at least, if it were possible to define a mapping between the annotations used in current corpora and the bottom classes of computational models of affect – it would be possible to plug these annotations in computational models of affect to predict the known target values. Thus, sentiment analysis would provide a way to assess the performance of alternative computational models of affect. We already noted that “text is an important modality for sensing affect because the bulk of computer user interfaces today are textually based” (Valitutti, 2004). Indeed, it is important to engage with texts – in the broad sense – because people express a lot of their emotions, moods and so on, through words (and, nowadays, also through media such as images and videos), especially on the web. In our opinion, a computational model of affect, to be deemed performant, should be performant also in sentiment analysis. In other words, sentiment analysis could be used to establish a baseline for the performance of computational models of affect.

17. Conclusions

To achieve the goal we have presented some premises from several areas like psychology, philosophy and computer science. In particular, we have defined ‘affect’, and we have underlined the relevance of the qualitative (phenomenal) dimension of affect itself. Then we have described OWL 2, and we have shown that it is expressive enough to formalize a wide range of affect theories. We have also argued that it should be adopted as a standard language for computational modelling of affect. Subsequently we have shown how machine learning techniques can be used to learn ontologies from labeled datasets, thus taking into account the qualitative dimension of affect of texts as recognized by humans in light of their ability to have qualitative experiences. Finally, we have

⁸ Sometimes it is possible to map the annotation made according to a model to annotations compatible with other models, albeit with some loss of information (e.g. about the intensity of affective states). Cambria’s hourglass of emotions (Cambria et al, 2012b), for instance, may be intended as a four-dimensional version of Plutchik’s wheel of emotions.

introduced the task of sentiment analysis, and we have argued that it is useful in assessing the performance of alternative computational models of affect. In light of our work, we propose that the combination of OWL and machine learning techniques is the most promising method to create CMAS, Computational Models of Affect. On the one hand, OWL can be used for the standardization of CMAS. On the other supervised machine learning techniques can be used to better take into account the qualitative dimension of affect, often neglected in modelling.

REFERENCES

- Abaalkhail, R., Guthier, B., Alharthi, R., & El Saddik, A. (2018). Survey on ontologies for affective states and their influences. *Semantic Web*, 9(4), 441–458
- Ackrill, J. L., & others. (1975). *Categories and de interpretatione*. Clarendon Press.
- Aggarwal, C. C., & Zhai, C. (2012). *Mining text data*. Springer Science & Business Media.
- Ahmad, K. (2011). *Affective computing and sentiment analysis: Emotion, metaphor and terminology* (Vol. 45). Springer Science & Business Media.
- Arp, R., Smith, B., & Spear, A. D. (2015). *Building ontologies with basic formal ontology*. The MIT Press.
- Ashburner, M., Ball, C. A., Blake, J. A., Botstein, D., Butler, H., Cherry, J. M., Davis, A. P., Dolinski, K., Dwight, S. S., Eppig, J. T., & others. (2000). Gene ontology: Tool for the unification of biology. *Nature Genetics*, 25(1), 25.
- Baader, F., Brandt, S., & Lutz, C. (2005). Pushing the envelope. *IJCAI*, 5, 364–369.
- Baader, F., Calvanese, D., McGuinness, D. L., Nardi, D., & Patel-Schneider, P. F. (2010). *The description logic handbook: Theory, implementation and applications*. Cambridge University Press.
- Bellegarda, J. R. (2014). Spoken language understanding for natural interaction: The siri experience. In *Natural interaction with robots, knowbots and smartphones* (pp. 3–14). Springer.
- Bosse, T., Gratch, J., Hoorn, J. F., Portier, M., & Siddiqui, G. F. (2010). Comparing three computational models of affect. In *Advances in practical applications of agents and multiagent systems* (pp. 175–184). Springer.
- Bosse, T., Jonker, C. M., & Treur, J. (2008). Formalization of damasio's theory of emotion, feeling and core consciousness. *Consciousness and Cognition*, 17(1), 94–113.

- Bosse, T., Jonker, C. M., Van Der Meij, L., & Treur, J. (2005). LEADSTO: A language and environment for analysis of dynamics by simulation. *German Conference on Multiagent System Technologies*, 165–178.
- Broekens, J., Degroot, D., & Kusters, W. A. (2008). Formal models of appraisal: Theory, specification, and computational model. *Cognitive Systems Research*, 9(3), 173–197.
- Brooks, D. J., Lignos, C., Finucane, C., Medvedev, M. S., Perera, I., Raman, V., Kress-Gazit, H., Marcus, M., & Yanco, H. A. (2012). Make it so: Continuous, flexible natural language interaction with an autonomous robot. *Workshops at the Twenty-Sixth AaaI Conference on Artificial Intelligence*.
- Buechel, S., & Hahn, U. (2017). EMOBANK: Studying the impact of annotation perspective and representation format on dimensional emotion analysis. *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, 2, 578–585.
- Bug, W. J., Ascoli, G. A., Grethe, J. S., Gupta, A., Fennema-Notestine, C., Laird, A. R., Larson, S. D., Rubin, D., Shepherd, G. M., Turner, J. A., & Martone, M. E. (2008). The nifstd and birnlex vocabularies: Building comprehensive ontologies for neuroscience. *Neuroinformatics*, 6(3), 175–194.
- Bühmann, L., Fleischhacker, D., Lehmann, J., Melo, A., & Völker, J. (2014). Inductive lexical learning of class expressions. *International Conference on Knowledge Engineering and Knowledge Management*, 42–53.
- Bühmann, L., Lehmann, J., & Westphal, P. (2016). DL-learner—a framework for inductive learning on the semantic web. *Web Semantics: Science, Services and Agents on the World Wide Web*, 39, 15–24.
- Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., & Rosati, R. (2007). Tractable reasoning and efficient query answering in description logics: The dl-lite family. *Journal of Automated Reasoning*, 39(3), 385–429.
- Cambria, E., & Hussain, A. (2012). *Sentic computing: Techniques, tools, and applications* (Vol. 2). Springer Science & Business Media.
- Cambria, E., Livingstone, A., & Hussain, A. (2012b). The hourglass of emotions. In *Cognitive behavioural systems* (pp. 144–157). Springer, Berlin, Heidelberg.
- Cambria, E., Poria, S., Gelbukh, A., & Thelwall, M. (2017). Sentiment analysis is a big suitcase. *IEEE Intelligent Systems*, 32(6), 74–80.
- Chen, C. C., & Tseng, Y.-D. (2011). Quality evaluation of product reviews using an information quality framework. *Decision Support Systems*, 50(4), 755–768.

- Chen, J., Hu, B., Moore, P., Zhang, X., & Ma, X. (2015). Electroencephalogram-based emotion assessment system using ontology and data mining techniques. *Applied Soft Computing*, *30*, 663–674.
- Cook, D. L., Bookstein, F. L., & Gennari, J. H. (2011). Physical properties of biological entities: An introduction to the ontology of physics for biology. *PLoS One*, *6*(12), e28708.
- Cover, T. M., & Thomas, J. A. (2012). *Elements of information theory*. John Wiley & Sons.
- Cox, S., & Little, C. (2004). *Time ontology in owl*. <https://www.w3.org/TR/owl-time/>
- Damasio, A. R. (1999). *The feeling of what happens: Body and emotion in the making of consciousness*. Harcourt Brace.
- Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, *41*(6), 391–407.
- Dennett, D. C. (1989). *The intentional stance*. MIT press.
- Deonna, J., & Teroni, F. (2012). *The emotions: A philosophical introduction*. Routledge.
- Frankish, K., & Ramsey, W. M. (2014). *The Cambridge handbook of artificial intelligence*. Cambridge University Press.
- Grosov, B. N., Horrocks, I., Volz, R., & Decker, S. (2003). Description logic programs: Combining logic programs with description logic. *Proceedings of the 12th International Conference on World Wide Web*, 48–57.
- Hastings, J., Ceusters, W., Mulligan, K., & Smith, B. (2012). *Annotating affective neuroscience data with the emotion ontology*.
- Hastings, J., Smith, B., Ceusters, W., Jensen, M., & Mulligan, K. (2012). Representing mental functioning: Ontologies for mental health and disease. *ICBO 2012: 3rd International Conference on Biomedical Ontology*, 1–5.
- Hatzivassiloglou, V., & McKeown, K. R. (1997). Predicting the semantic orientation of adjectives. *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics and Eighth Conference of the European Chapter of the Association for Computational Linguistics*, 174–181.
- Herre, H. (2010). General formal ontology (gfo): A foundational ontology for conceptual modelling. In *Theory and applications of ontology: Computer applications* (pp. 297–345). Springer.

- Hitzler, P., Krötzsch, M., Parsia, B., Patel-Schneider, P. F., & Rudolph, S. (2012). *OWL 2 Web Ontology Language: Primer (Second Edition)*. <http://www.w3.org/TR/owl2-primer/>
- Hitzler, P., Krtzsch, M., & Rudolph, S. (2009). *Foundations of semantic web technologies*. Chapman & Hall/CRC.
- Hornby, A. S., Wehmeier, S., & Ashby, M. (2000). *Oxford advanced learner's dictionary*. Oxford university press Oxford.
- Horrocks, I., Kutz, O., & Sattler, U. (2006). *The even more irresistible sroiq*.
- Horrocks, I., Patel-Schneider, P. F., Boley, H., Tabet, S., Grosfand, B., & Dean, M. (2004). *SWRL: A semantic web rule language combining OWL and RuleML*. <http://www.w3.org/Submission/SWRL/>
- Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32(127), 127–136.
- Jolliffe, I. T. (1986). Principal component analysis and factor analysis. In *Principal component analysis* (pp. 115–128). Springer.
- Jurafsky, D., & Martin, J. H. (2019). Lexicons for Sentiment, Affect, and Connotation in Speech and Language Processing. In D. Jurafsky, & J. H. Martin, *Speech and Language Processing*
- Jurafsky, D., & Martin, J. H. (2019). Naive Bayes and Sentiment Classification. In D. Jurafsky, & J. H. Martin, *Speech and Language Processing*
- Kang, H., Yoo, S. J., & Han, D. (2012). Senti-lexicon and improved naive bayes algorithms for sentiment analysis of restaurant reviews. *Expert Systems with Applications*, 39(5), 6000–6010.
- Kazakov, Y., Krötzsch, M., & Simancik, F. (2014). The incredible elk. *Journal of Automated Reasoning*, 53(1), 1–61.
- Kleinginna, P. R., & Kleinginna, A. M. (1981). A categorized list of emotion definitions, with suggestions for a consensual definition. *Motivation and Emotion*, 5(4), 345–379.
- Kripke, S. A. (2013). The church-turing thesis as a special corollary of godel's completeness theorem. In B. J. Copeland, C. Posy, & O. Shagrir (Eds.), *Computability: Turing, gödel, church, and beyond*. MIT Press.
- Krötzsch, M., Rudolph, S., & Hitzler, P. (2008). Description logic rules. *ECAI 2008 - 18th European Conference on Artificial Intelligence, Patras, Greece, July 21-25, 2008, Proceedings*, 80–84.
- Laclavik, M., Balogh, Z., Babik, M., & Hluchý, L. (2012). Agentowl: Semantic knowledge model and agent architecture. *Computing and Informatics*, 25(5), 421–439.

- Larson, S. D., Fong, L. L., Gupta, A., Condit, C., Bug, W. J., & Martone, M. E. (2007). A formal ontology of subcellular neuroanatomy. *Frontiers in Neuroinformatics*, 1, 3.
- Lazarus, R. S., & Lazarus, R. S. (1991). *Emotion and adaptation*. Oxford University Press on Demand.
- Lehmann, J. (2010). *Learning owl class expressions* (Vol. 6). IOS Press.
- Lehmann, J., Auer, S., Bühmann, L., & Tramp, S. (2011). Class expression learning for ontology engineering. *Web Semantics: Science, Services and Agents on the World Wide Web*, 9(1), 71–81.
- Lehmann, J., & Haase, C. (2009). Ideal downward refinement in the EL description logic. *International Conference on Inductive Logic Programming*, 73–87.
- Lehmann, J., & Hitzler, P. (2010a). Concept learning in description logics using refinement operators. *Machine Learning*, 78, 203–250.
- Lehmann, J., & Hitzler, P. (2010b). Concept learning in description logics using refinement operators. *Machine Learning*, 78(1-2), 203.
- Lewis, D. D., & Ringuette, M. (1994). A comparison of two learning algorithms for text categorization. *Third Annual Symposium on Document Analysis and Information Retrieval*, 33, 81–93.
- Li, Y.-M., & Li, T.-Y. (2013). Deriving market intelligence from microblogs. *Decision Support Systems*, 55(1), 206–217.
- Linchan, M. H., Barkmeyer Jr, E. J., & Hendryx, S. (2012). The date-time vocabulary. *Seventh International Conference on Formal Ontologies in Information Systems (Fois 2012)*.
- Liu, B. (2012). Sentiment analysis and opinion mining. *Synthesis Lectures on Human Language Technologies*, 5(1), 1–167.
- Liu, B. (2015). *Sentiment analysis: Mining opinions, sentiments, and emotions*. Cambridge University Press.
- Milea, V., Frasinca, F., & Kaymak, U. (2012). TOWL: A temporal web ontology language. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 42(1), 268–281.
- Miller, G. A., Beckwith, R., Fellbaum, C., Gross, D., & Miller, K. J. (1990). Introduction to wordnet: An on-line lexical database. *International Journal of Lexicography*, 3(4), 235–244.
- Mitchell, T. (1997). *Machine Learning*. McGraw Hill.

- Mohammad, S., Dunne, C., & Dorr, B. (2009). Generating high-coverage semantic orientation lexicons from overtly marked words and a thesaurus. *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 2-Volume 2*, 599–608.
- Moreira, A. F., Vieira, R., Bordini, R. H., & Hübner, J. F. (2005). Agent-oriented programming with underlying ontological reasoning. *International Workshop on Declarative Agent Languages and Technologies*, 155–170.
- Morgenstern, L., Welty, C., Boley, H., & Hallmark, G. (2013). *RIF Primer (Second Edition)*. <http://www.w3.org/TR/rif-primer/>
- Motik, B., Cuenca Grau, B., Horrocks, I., Wu, Z., Fokoue, A., & Lutz, C. (2012). *OWL 2 Web Ontology Language: Profiles (Second Edition)*. <http://www.w3.org/TR/owl2-profiles/>
- Motik, B., Patel-Schneider, P. F., & Cuenca Grau, B. (2012). *OWL 2 Web Ontology Language: Direct Semantics (Second Edition)*. <http://www.w3.org/TR/owl2-direct-semantics/>
- Motik, B., Patel-Schneider, P. F., & Parsia, B. (2012). *OWL 2 Web Ontology Language: Structural Specification and Functional-Style Syntax (Second Edition)*. <http://www.w3.org/TR/owl2-syntax/>
- Motik, B., Sattler, U., & Studer, R. (2005). Query answering for owl-dl with rules. *Web Semantics: Science, Services and Agents on the World Wide Web*, 3(1), 41–60.
- Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, 83(4), 435–450.
- Ng, H. T., Goh, W. B., & Low, K. L. (1997). Feature selection, perceptron learning, and a usability case study for text categorization. *ACM Sigir Forum*, 31, 67–73.
- Nienhuys-Cheng, S.-H., & De Wolf, R. (1997). *Foundations of inductive logic programming* (Vol. 1228). Springer Science & Business Media.
- O'Regan, J. K. (2011). *Why red doesn't sound like a bell: Understanding the feel of consciousness*. Oxford University Press.
- O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5), 939–973.
- Picard, R. W., & others. (1995). *Affective computing*. Perceptual Computing Section, Media Laboratory, Massachusetts Institute of Technology.
- Power, R., & Third, A. (2010). Expressing owl axioms by english sentences: Dubious in theory, feasible in practice. *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*, 1006–1013.

- Plutchik, R. (2001). The nature of emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice. *American scientist*, 89(4), 344-350.
- Rao, A. S. (1996). AgentSpeak (I): BDI agents speak out in a logical computable language. *European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, 42–55.
- Rector, A., & Horrocks, I. (1997). Experience building a large, re-usable medical ontology using a description logic with transitivity and concept inclusions. *Proceedings of the Workshop on Ontological Engineering, Aaai Spring Symposium (Aaai'97), Stanford, ca*, 321–325.
- Reizenzein, R., Hudlicka, E., Dastani, M., Gratch, J., Hindriks, K., Lorini, E., & Meyer, J. J. C. (2013). Computational modeling of emotion: Toward improving the inter-and intradisciplinary exchange. *IEEE Transactions on Affective Computing*, 4(3), 246–266.
- Reynolds, D. (2013). *OWL 2 RL in RIF (Second Edition)*. <http://www.w3.org/TR/rif-owl-rl/>
- Ruiz, M. E., & Srinivasan, P. (1999). Hierarchical neural networks for text categorization. *Proceedings of the 22nd Annual International Acm Sigir Conference on Research and Development in Information Retrieval*, 281–282.
- Russell, S. J., & Norvig, P. (2010). *Artificial intelligence: A modern approach*. New Jersey, Pearson Education.
- Salguero, A., & Espinilla, M. (2016). Description logic class expression learning applied to sentiment analysis. *Sentiment Analysis and Ontology Engineering, an Environment of Computational Intelligence*, 93–111.
- Sánchez-Rada, J. F., & Iglesias, C. A. (2016). Onyx: A linked data approach to emotion representation. *Information Processing & Management*, 52(1), 99-114.
- Schneider, M. (2012). *OWL 2 Web Ontology Language: RDF-Based Semantics (Second Edition)*. <http://www.w3.org/TR/owl2-rdf-based-semantics/>
- Schröder, M., Baggia, P., Burkhardt, F., Pelachaud, C., Peter, C., & Zovato, E. (2011). EmotionML—an upcoming standard for representing emotions and related states. *International Conference on Affective Computing and Intelligent Interaction*, 316–325.
- Schreiber, G., & Raimond, Y. (2014). *RDF 1.1 Primer*. <http://www.w3.org/TR/rdf11-primer/>

- Scarle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417–424.
- Spackman, K. (2000). Managing clinical terminology hierarchies using algorithmic calculation of subsumption: Experience with snomed-rt. *J. Of the Amer. Med. Informatics Ass.*
- Strapparava, C., Valitutti, A., & others. (2004). Wordnet affect: An affective extension of wordnet. *Lrec*, 4, 1083–1086.
- Turing, A. M. (1937). On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, 2(1), 230–265.
- Valitutti, A., Strapparava, C., & Stock, O. (2004). Developing affective lexical resources. *PsychNology Journal*, 2(1), 61–83.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. and Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems* (pp. 5998–6008).
- Villanueva-Rosales, N., & Dumontier, M. (2007). Describing chemical functional groups in owl-dl for the classification of chemical compounds. *OWLED*, 258.
- Welty, C., & Fikes, R. (2006). A reusable ontology for fluents in owl. *Frontiers in Artificial Intelligence and Applications*, 226–236.
- Westerski, A., & Sánchez-Rada, J. F. (2013). *Marl Ontology Specification*, V1. 0 May 2013.
- Williams, T., Briggs, G., Oosterveld, B., & Scheutz, M. (2015). Going beyond literal command-based instructions: Extending robotic natural language interaction capabilities. *Twenty-Ninth Aaai Conference on Artificial Intelligence*.