An Application of the TCL Logic to Aerospace Missions

Antonio Lieto[1,2,†], Marco Orusa[3,†], Gian Luca Pozzato[3,*] and Stefano Zoia[3,†]

[1] *Cognition, Interaction and Intelligent Technologies Lab/DISPC, Università di Salerno, Via Giovanni Paolo II, 132 84084 Fisciano (SA), Italy*

[2] *Cognitive Systems for Robotics, ICAR-CNR Institute, Palermo, Italy*

[3] *Dipartimento di Informatica, Università degli Studi di Torino, c.so Svizzera 185, 10149 Turin, Italy*

Abstract

In this work we present a tool for the automatic attribution of emotions to astronauts speeches obtained from the NASA Mission Transcript Collection. While, at the current state of affairs, we do not have yet a quantitative evaluation of such tool, we present - as a proof of concept - some qualitative analysis showing the potential usefulness of our approach. The broader objective of this work is providing an intelligent system for monitoring the psycho-physical condition of an astronaut during an aerospace missions. This system exploits a commonsense reasoning framework based on the logic $\mathbf{T}^{\mathrm{CL}}$, a probabilistic extension of Description Logics of typicality able to deal with the conceptual combination of prototypical descriptions, i.e. commonsense representations of given concepts. Starting from an ontological formalization of emotions based on the Plutchik model, known as ArsEmotica, the system exploits the logic $\mathbf{T}^{\mathrm{CL}}$ to automatically generate novel commonsense semantic representations of compound emotions (e.g. Love as derived from the combination of Joy and Trust according to Plutchik). The generated emotions have then been applied for emotion attribution in the context of aerospace missions, in order to classify transcriptions of astronaut's speeches in the corresponding emotions.

Keywords

Description Logics, Commonsense Reasoning, Concept combination, Affective Computing

1. Introduction

The problem of assessing the physical condition and emotional states of astronauts during space missions has become critically important in recent years. What might appear to be a straightforward task can, in reality, prove to be unfeasible for a very simple reason: the methodology employed to evaluate the astronaut's psycho-physical state could itself contribute to a deterioration of that state. For example, an astronaut might feel overwhelmed by the excessive workload imposed by the tasks assigned to him at that moment: forcing him, for instance, to answer questions or interact with an app in order to monitor his condition could increase the level of stress he is already experiencing. It is therefore essential that such assessments be carried out without introducing additional stress for the subject. Moreover, the crew of an aerospace mission need to wear appropriate helmets and equipment, reducing the physical space available for sensors and measuring devices. In other words, the astronaut should be appropriately monitored in a non-intrusive manner that does not interfere with their activities. The system must, in effect, be transparent.

The tasks that astronauts need to continuously carry out are often physically or mentally challenging. This can have an impact on their psycho-physical state, leading to negative outcomes ranging from reduced performance to deadly risks. Previous works targeted the mental workload, stress and fatigue of the crew members as crucial aspects in order to anticipate critical performance drops and to improve training techniques [1, 2, 3, 4, 5, 6, 7]. We envision a multimodal system able to take into account as

CILC 2025: 40th Italian Conference on Computational Logic, June 25–27, 2025, Alghero, Italy

*Corresponding author.
†These authors contributed equally.
alieto@unisa.it (A. Lieto); marco.orusa911@edu.unito.it (M. Orusa); gianluca.pozzato@unito.it (G. L. Pozzato); stefano.zoia@unito.it (S. Zoia)
https://www.antoniolieto.net (A. Lieto); http://www.di.unito.it/~pozzato/ (G. L. Pozzato); https://informatica.unito.it/do/docenti.pl/Alias?stefano.zoia#tab-profilo (S. Zoia)
0000-0002-8323-8764 (A. Lieto); 0000-0002-3952-4624 (G. L. Pozzato)

many physiological and psychological data as possible and combine them in a representation of the state of the astronaut, which should be produced in real time and presented as an intuitive output. The tool we present in this paper is intended as a component of such a system: given the sentences uttered by the astronauts, it detects their emotional content. This analysis requires no other sensor than the microphones that are already on board.

Sentiment analysis, an important area in natural language processing, has evolved from traditional techniques to deep learning approaches, notably transformer-based models. The field now extends beyond polarity classification to the recognition of specific emotions like joy, anger, and fear, despite challenges such as ambiguity, cultural differences, and the scarcity of fine-grained annotated data. Emotions have long been recognized as fundamental to human experience across cultures. Two primary modeling approaches exist: dimensional models, which explore continuous and physiological aspects of emotions [8, 9, 10], and categorical models, which capture conscious emotional states [11, 12, 13]. Recent advances include deep neural networks for affective state detection, leading to the development of annotated corpora, ontologies, and lexicons [14, 15, 16]. Attention-based neural models, such as the bidirectional CNN-RNN with attention [17], have shown state-of-the-art results on datasets ranging from long reviews to short texts like tweets. The affective content of text encompasses emotion, sentiment, mood, personality, and interpersonal stance [18]. This complexity has driven recent work toward fine-grained analysis, including the prediction of emotion and sentiment intensity or activation [19, 20], often using ensemble methods that combine deep learning with traditional features.

Depending on the specific research goals addressed, one could be interested in issuing a discrete label describing the affective state expressed (frustration, anger, joy, etc.) to address different contexts of interaction and tasks. Both basic emotion theories, in the Plutchik-Ekman [12] tradition, and dimensional models of emotions, have provided a precious theoretical grounding for the development of lexical resources [21, 22, 23, 24, 25] and computational models for emotion extraction. However, there is a general tendency to move towards richer, finer-grained models, possibly including complex emotions, especially in the context of data-driven and task-driven approaches, where restricting the automatic detection to a small set of basic emotions would fall short to achieve the objective.

In [26, 27], the authors introduce a system for emotion attribution and recommendation, employing a white box approach to emotion classification based on the human-like conceptual combination framework proposed in the $\mathbf{T}^{\text{CL}}$ logic [28, 29]. In this work, we adopt a similar approach in order to perform emotion attribution in aerospace missions by exploiting the ontology of the Plutchik emotion model introduced in [26]. We refer to the ArsEmotica Ontology and to the NRC Emotion Intensity Lexicon [19], that provides a list of English words, each with real-values representing intensity scores for the eight basic emotions of Plutchik's theory. The generation process building the knowledge base of emotions in the $\mathbf{T}^{\text{CL}}$ logic unfolds in two main steps. First, the system constructs a prototypical description of basic emotions using the logic $\mathbf{T}^{\text{CL}}$, with typicality inclusions of the form $\mathbf{T}(BasicEmotion) \sqsubseteq Property$. These basic emotions are the eight defined in Plutchik's model. In the second step, the system analyzes an utterance to assign the most appropriate emotion. The utterance is first processed by a tokenizer, which extracts relevant words. The system then checks whether the utterance contains terms associated with basic emotions. If such terms are found, their individual scores are retrieved and summed for each of the eight basic emotions. This results in a score distribution that indicates the presence of specific emotions. A similar method is used to detect combined emotions: the system searches for terms derived from the combination of two basic emotions, and aggregates their scores accordingly.

The overall objective is to develop a system capable of assessing the astronaut's psycho-physical state and issuing appropriate alerts in the event that a potentially hazardous situation is detected—such as the astronaut being excessively stressed, frightened, angry, or disoriented—without introducing additional tasks or, consequently, additional stress for the astronaut during the mission. The use of wearable devices that monitor physiological parameters (e.g., blood pressure, heart rate, etc.) is thus complemented by an intelligent system capable of detecting the astronaut's emotional states.

2. A Typicality Description Logic for Concept Combination

The automatic generation of novel concepts within a knowledge base (also known as *knowledge invention* process) can be obtained, as happens in humans [30, 28], by exploiting a process of commonsense conceptual combination. Such ability is associated to creative thinking and problem solving, however, it still represents an open challenge in Artificial Intelligence [31]. Dealing with this problem, indeed, requires, from an AI perspective, the harmonization of two conflicting requirements that are hardly accommodated in symbolic systems [32]: the need for a syntactic and semantic compositionality (typical of logical systems) and the one concerning the exhibition of typicality effects. According to a well-known argument [33], in fact, prototypes (i.e. commonsense conceptual representations based on typical properties) are not compositional. The argument runs as follows: consider a concept like *pet fish*. It results from the composition of the concept *pet* and of the concept *fish*. However, the prototype of *pet fish* cannot result from the composition of the prototypes of a pet and a fish: e.g., a typical pet is furry and warm, a typical fish is grayish, but a typical pet fish is neither furry and warm nor grayish (typically, it is red). The *pet fish* phenomenon is a paradigmatic example of the difficulty to address when building formalisms and systems trying to imitate this combinatorial human ability.

In this work, we exploit the nonmonotonic Description Logic $\mathbf{T}^{\text{CL}}$ (typicality-based compositional logic), introduced in [28, 29], which is able to account for this type of human-like concept combination. Other works have already shown how such logic can be used to model complex cognitive phenomena [30], goal-directed creative problem solving [34, 35, 36] and to build intelligent applications for computational creativity [37]. In particular, we show how it can be used as a tool for the generation of novel compound emotions and, as a consequence, for the suggestion of novel emotion-related contents. In $\mathbf{T}^{\text{CL}}$, "typical" properties can be directly specified by means of a "typicality" operator $\mathbf{T}$ enriching the underlying Description Logic (from now on, DL for short), and a TBox can contain inclusions of the form $\mathbf{T}(C) \sqsubseteq D$ to represent that "typical Cs are also Ds". As a difference with standard DLs, in the logic $\mathbf{T}^{\text{CL}}$ one can consistently express exceptions and reason about defeasible inheritance as well. Typicality inclusions are also equipped by a real number $p \in (0.5, 1]$ representing the probability/degree of belief in such a typical property: this allows us to define a semantics inspired to the DISPONTE semantics [38] characterizing probabilistic extensions of DLs, which in turn is used in order to describe different *scenarios* where only some typicality properties are considered. Given a KB containing the description of two concepts C_H and C_M occurring in it, we then consider only *some* scenarios in order to define a revised knowledge base, enriched by typical properties of the combined concept $C \sqsubseteq C_H \sqcap C_M$ by also implementing a heuristics coming from the cognitive semantics.

By relying on $\mathbf{T}^{\text{CL}}$, as in [26], we first automatically build prototypes of existing *basic* emotions by extracting information about *concepts* or *properties* relying on the ArsEmotica ontology enriched with the NRC Emotion Intensity Lexicon [19]: this lexicon associates, in descending order of frequency, words to emotional concepts. In this setting, words with the highest frequencies of association to emotional concepts have been used as typical features of the basic emotions in the Plutchik model. Such prototypes of basic emotions have been formalized by means of a $\mathbf{T}^{\text{CL}}$ knowledge base, whose TBox contains both *rigid* inclusions of the form

$$BasicEmotion \sqsubseteq Concept,$$

in order to express essential desiderata but also constraints, as an example $Joy \sqsubseteq PositiveEmotion$ as well as *prototypical* properties of the form

$$p \ :: \ \mathbf{T}(BasicEmotion) \sqsubseteq TypicalConcept,$$

representing typical concepts of a given emotion, where p is a real number in the range $(0.5, 1]$, expressing the frequency of such a concept in items belonging to that emotion: for instance, 0.72 :: $\mathbf{T}(Surprise) \sqsubseteq Delight$ is used to express that the typical feature of being surprised contains/refers to the emotional concept *Delight* with a frequency/probability/degree of belief of the 72%.

The logic $\mathbf{T}^{\text{CL}}$ [30] combines three main ingredients. The first one relies on the DL of typicality $\mathcal{ALC} + \mathbf{T_R}$ introduced in [39], which allows to describe the *prototype* of a concept. In this logic, "typical" properties can be directly specified by means of a "typicality" operator $\mathbf{T}$ enriching the underlying DL, and a TBox can contain inclusions of the form $\mathbf{T}(C) \sqsubseteq D$ to represent that "typical Cs are also Ds". As a difference with standard DLs, in the logic $\mathcal{ALC} + \mathbf{T_R}$ one can consistently express exceptions and reason about defeasible inheritance as well. For instance, a knowledge base can consistently express that "normally, singers can sing", whereas "trappers usually cannot sing" by $\mathbf{T}(\mathit{Singer}) \sqsubseteq \mathit{CanSing}$ and $\mathbf{T}(\mathit{Trapper}) \sqsubseteq \neg \mathit{CanSing}$, given that $\mathit{Trapper} \sqsubseteq \mathit{Singer}$. The semantics of the $\mathbf{T}$ operator is characterized by the properties of *rational logic* [40], recognized as the core properties of nonmonotonic reasoning. $\mathcal{ALC} + \mathbf{T_R}$ is characterized by a minimal model semantics corresponding to an extension to DLs of a notion of *rational closure* as defined in [40] for propositional logic: the idea is to adopt a preference relation over $\mathcal{ALC} + \mathbf{T_R}$ models, where intuitively a model is preferred to another one if it contains less exceptional elements, as well as a notion of *minimal entailment* restricted to models that are minimal with respect to such preference relation. As a consequence, $\mathbf{T}$ inherits well-established properties like *specificity* and *irrelevance*: in the example, the logic $\mathcal{ALC} + \mathbf{T_R}$ allows us to infer $\mathbf{T}(\mathit{Singer} \sqcap \mathit{LongHair}) \sqsubseteq \mathit{CanSing}$ (having long hair is irrelevant with respect to being able to sing or not) and, if one knows that Tony is a typical trapper, to infer that he cannot sing, giving preference to the most specific information.

A second ingredient consists of a distributed semantics similar to the one of probabilistic DLs known as DISPONTE [41], which allows labeling inclusions $\mathbf{T}(C) \sqsubseteq D$ with a real number between 0.5 and 1, which represents its degree of belief/probability, under the assumption that each axiom is independent from each others. Degrees of belief in typicality inclusions allow defining a probability distribution over *scenarios*: roughly speaking, a scenario is obtained by choosing, for each typicality inclusion, whether it is considered as true or false. In a slight extension of the above example, we could have the need to represent that both the typicality inclusions about singers and trappers have a degree of belief of 80%, whereas we also believe that, normally, singers are also able to play an instrument with a higher degree of 95%, with the following KB:

(1) $\mathit{Trapper} \sqsubseteq \mathit{Singer}$
(2) $0.8\ ::\ \mathbf{T}(\mathit{Singer}) \sqsubseteq \mathit{CanSing}$
(3) $0.8\ ::\ \mathbf{T}(\mathit{Trapper}) \sqsubseteq \neg \mathit{CanSing}$
(4) $0.95\ ::\ \mathbf{T}(\mathit{Singer}) \sqsubseteq \mathit{PlayInstrument}$

In this case, we consider eight different scenarios, representing all possible combinations of typicality inclusion: as an example, $\{((2), 1), ((3), 0), ((4), 1)\}$ represents the scenario in which (2) and (4) hold, whereas (3) does not. Obviously, (1) holds in every scenario, since it represents a rigid property, not admitting exceptions. We equip each scenario with a probability depending on those of the involved inclusions: the scenario of the example has probability 0.8×0.95 (since 2 and 4 are involved) $\times(1-0.8)$ (since 3 is not involved) $= 0.152 = 15.2\%$. Such probabilities are then taken into account in order to choose the most adequate scenario describing the prototype of the combined concept.

As a third element of the proposed formalization is a method inspired by cognitive semantics [42] for the identification of a dominance effect between the concepts to be combined: for every combination, we distinguish a HEAD, representing the stronger element of the combination, and a MODIFIER. The basic idea is: given a KB and two concepts C_H (HEAD) and C_M (MODIFIER) occurring in it, we consider only *some* scenarios in order to define a revised knowledge base, enriched by typical properties of the combined concept $C \sqsubseteq C_H \sqcap C_M$.

Let us now present the logic $\mathbf{T}^{\text{CL}}$ more precisely. The language of $\mathbf{T}^{\text{CL}}$ extends the basic DL $\mathcal{ALC}$ by *typicality inclusions* of the form $\mathbf{T}(C) \sqsubseteq D$ equipped by a real number $p \in (0.5, 1]$ – observe that the extreme 0.5 is not included – representing its degree of belief, whose meaning is that "we believe with degree/probability p that, normally, Cs are also Ds" [1]

[1]The reason why we only allow typicality inclusions equipped with probabilities $p > 0.5$ is due to our effort of integrating two different semantics: typicality based logic and DISPONTE. In particular, as detailed in [30] this choice seems to be the only one compliant with both formalisms. On the contrary, it would be misleading to also allow low degrees of belief for typicality inclusions, since typical knowledge is known to come with a low degree of uncertainty.

Definition 2.1 (Language of $\mathbf{T}^{\text{CL}}$). *We consider an alphabet of concept names* $\mathtt{C}$*, of role names* $\mathtt{R}$*, and of individual constants* $\mathtt{O}$*. Given $A \in \mathtt{C}$ and $R \in \mathtt{R}$, we define:*

$$C, D := A \mid \top \mid \bot \mid \neg C \mid C \sqcap C \mid C \sqcup C \mid \forall R.C \mid \exists R.C$$

We define a knowledge base $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ where: (i) $\mathcal{R}$ is a finite set of rigid properties of the form $C \sqsubseteq D$; (ii) $\mathcal{T}$ is a finite set of typicality properties of the form $p\ ::\ \mathbf{T}(C) \sqsubseteq D$ where $p \in (0.5, 1] \subseteq \mathbb{R}$ is the degree of belief of the typicality inclusion; (iii) $\mathcal{A}$ is the ABox, i.e. a finite set of formulas of the form either $C(a)$ or $R(a, b)$, where $a, b \in \mathtt{O}$ and $R \in \mathtt{R}$.

A model $\mathcal{M}$ in $\mathbf{T}^{\text{CL}}$ extends standard $\mathcal{ALC}$ ones by a preference relation among domain elements as in the DL of typicality [39]. In this respect, $x < y$ means that x is "more normal" than y, and that typical members of a concept C are the minimal elements of C with respect to $<$. An element $x \in \Delta^{\mathcal{I}}$ is a *typical instance* of C if $x \in C^{\mathcal{I}}$ and there is no C-element in $\Delta^{\mathcal{I}}$ *more normal* than x. Formally:

Definition 2.2 (Model of $\mathbf{T}^{\text{CL}}$). *A model $\mathcal{M}$ is any structure $\langle \Delta^{\mathcal{I}}, <, .^{\mathcal{I}} \rangle$ where: (i) $\Delta^{\mathcal{I}}$ is a non empty set of items called the domain; (ii) $<$ is an irreflexive, transitive, well-founded and modular (for all x, y, z in $\Delta^{\mathcal{I}}$, if $x < y$ then either $x < z$ or $z < y$) relation over $\Delta^{\mathcal{I}}$; (iii) $.^{\mathcal{I}}$ is the extension function that maps each atomic concept C to $C^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, and each role R to $R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, and is extended to complex concepts as follows: $(\neg C)^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$; $(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$; $(C \sqcup D)^{\mathcal{I}} = C^{\mathcal{I}} \cup D^{\mathcal{I}}$; $(\exists R.C)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \exists (x, y) \in R^{\mathcal{I}}$ such that $y \in C^{\mathcal{I}}\}$; $(\forall R.C)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \forall (x, y) \in R^{\mathcal{I}}$ we have $y \in C^{\mathcal{I}}\}$; $(\mathbf{T}(C))^{\mathcal{I}} = Min_<(C^{\mathcal{I}})$, where $Min_<(C^{\mathcal{I}}) = \{x \in C^{\mathcal{I}} \mid \nexists y \in C^{\mathcal{I}}$ s.t. $y < x\}$.*

A model $\mathcal{M}$ can be equivalently defined by postulating the existence of a function $k_{\mathcal{M}} : \Delta^{\mathcal{I}} \longmapsto \mathbb{N}$, where $k_{\mathcal{M}}$ assigns a finite rank to each domain element [39]: the rank of x is the length of the longest chain $x_0 < \cdots < x$ from x to a minimal x_0, i.e. such that there is no x' such that $x' < x_0$. The rank function $k_{\mathcal{M}}$ and $<$ can be defined from each other by letting $x < y$ if and only if $k_{\mathcal{M}}(x) < k_{\mathcal{M}}(y)$.

Definition 2.3. *Let $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ be a KB. Given a model $\mathcal{M} = \langle \Delta^{\mathcal{I}}, <, .^{\mathcal{I}} \rangle$, we assume that $.^{\mathcal{I}}$ is extended to assign a domain element $a^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$ to each individual constant a of* $\mathtt{O}$*. We say that: (i) $\mathcal{M}$ satisfies $\mathcal{R}$ if, for all $C \sqsubseteq D \in \mathcal{R}$, we have $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$; (ii) $\mathcal{M}$ satisfies $\mathcal{T}$ if, for all $q\ ::\ \mathbf{T}(C) \sqsubseteq D \in \mathcal{T}$, we have that [2] $\mathbf{T}(C)^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, i.e. $Min_<(C^{\mathcal{I}}) \subseteq D^{\mathcal{I}}$; (iii) $\mathcal{M}$ satisfies $\mathcal{A}$ if, for each assertion $F \in \mathcal{A}$, if $F = C(a)$ then $a^{\mathcal{I}} \in C^{\mathcal{I}}$, otherwise if $F = R(a, b)$ then $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$.*

Even if the typicality operator $\mathbf{T}$ itself is nonmonotonic (i.e. $\mathbf{T}(C) \sqsubseteq E$ does not imply $\mathbf{T}(C \sqcap D) \sqsubseteq E$), what is inferred from a KB can still be inferred from any KB' with KB $\subseteq$ KB', i.e. the resulting logic is monotonic. As already mentioned, in order to perform useful nonmonotonic inferences, in [39] the authors have strengthened the above semantics by restricting entailment to a class of minimal models. Intuitively, the idea is to restrict entailment to models that *minimize the atypical instances of a concept*. The resulting logic corresponds to a notion of *rational closure* on top of $\mathcal{ALC} + \mathbf{T_R}$. Such a notion is a natural extension of the rational closure construction provided in [40] for the propositional logic. This nonmonotonic semantics relies on minimal rational models that minimize the *rank of domain elements*. Informally, given two models of KB, one in which a given domain element x has rank 2 (because for instance $z < y < x$), and another in which it has rank 1 (because only $y < x$), we prefer the latter, as in this model the element x is assumed to be "more typical" than in the former. Query entailment is then restricted to minimal *canonical models*. The intuition is that a canonical model contains all the individuals that enjoy properties that are consistent with KB. This is needed when reasoning about the rank of the concepts: it is important to have them all represented.

Given a KB $\mathcal{K} = \langle \mathcal{R}, \mathcal{T}, \mathcal{A} \rangle$ and given two concepts C_H and C_M occurring in $\mathcal{K}$, the logic $\mathbf{T}^{\text{CL}}$ allows defining a prototype of the combined concept C as the combination of the HEAD C_H and the MODIFIER

[2] It is worth noticing that here the degree q does not play any role. Indeed, a typicality inclusion $\mathbf{T}(C) \sqsubseteq D$ holds in a model only if it satisfies the semantic condition of the underlying DL of typicality, i.e. minimal (typical) elements of C are elements of D. The degree of belief q will have a crucial role in the application of the distributed semantics, allowing the definition of scenarios as well as the computation of their probabilities.

C_M, where the typical properties of the form $\mathbf{T}(C) \sqsubseteq D$ (or, equivalently, $\mathbf{T}(C_H \sqcap C_M) \sqsubseteq D$) to be ascribed to the concept C are obtained by considering blocks of scenarios with the same probability, in decreasing order starting from the highest one. We first discard all the inconsistent scenarios, then:

- we discard those scenarios considered as *trivial*, consistently inheriting all the properties from the HEAD from the starting concepts to be combined. This choice is motivated by the challenges provided by task of commonsense conceptual combination itself: in order to generate plausible and creative compounds, it is necessary to maintain a level of surprise in the combination. Thus both scenarios inheriting all the properties of the two concepts and all the properties of the HEAD are discarded, since they prevent this surprise;
- among the remaining ones, we discard those inheriting properties from the MODIFIER which are in conflict with properties that could be consistently inherited from the HEAD;
- if the set of scenarios of the current block is empty, i.e. all the scenarios have been discarded either because trivial or because the MODIFIER is preferred, we repeat the procedure by considering the block of scenarios having the immediately lower probability.

Remaining scenarios are those selected by the logic $\mathbf{T}^{\text{CL}}$. The ultimate output of our mechanism is a knowledge base in the logic $\mathbf{T}^{\text{CL}}$ whose set of typicality properties is enriched by those of the compound concept C. Given a scenario w satisfying the above properties, we define the properties of C as the set of inclusions $p \ :: \ \mathbf{T}(C) \sqsubseteq D$, for all $\mathbf{T}(C) \sqsubseteq D$ that are entailed from w in the logic $\mathbf{T}^{\text{CL}}$. The probability p is such that:

- if $\mathbf{T}(C_H) \sqsubseteq D$ is entailed from w, that is to say D is a property inherited either from the HEAD (or from both the HEAD and the MODIFIER), then p corresponds to the degree of belief of such inclusion of the HEAD in the initial knowledge base, i.e. $p \ : \ \mathbf{T}(C_H) \sqsubseteq D \in \mathcal{T}$;
- otherwise, i.e. $\mathbf{T}(C_M) \sqsubseteq D$ is entailed from w, then p corresponds to the degree of belief of such inclusion of a MODIFIER in the initial knowledge base, i.e. $p \ : \ \mathbf{T}(C_M) \sqsubseteq D \in \mathcal{T}$.

The knowledge base obtained as the result of combining concepts C_H and C_M into the compound concept C is called C*-revised* knowledge base, and it is defined as follows:

$$\mathcal{K}_C = \langle \mathcal{R}, \mathcal{T} \cup \{p \ : \ \mathbf{T}(C) \sqsubseteq D\}, \mathcal{A} \rangle,$$

for all D such that either $\mathbf{T}(C_H) \sqsubseteq D$ is entailed in w or $\mathbf{T}(C_M) \sqsubseteq D$ is entailed in w, and p is defined as above. As an example, consider the following version of the above mentioned *Pet-Fish* problem. Let KB contains the following inclusions:

$$Fish \sqsubseteq LivesInWater \tag{1}$$
$$0.6 \ :: \ \mathbf{T}(Fish) \sqsubseteq Greyish \tag{2}$$
$$0.8 \ :: \ \mathbf{T}(Fish) \sqsubseteq Scaly \tag{3}$$
$$0.8 \ :: \ \mathbf{T}(Fish) \sqsubseteq \neg Affectionate \tag{4}$$
$$0.9 \ :: \ \mathbf{T}(Pet) \sqsubseteq \neg LivesInWater \tag{5}$$
$$0.9 \ :: \ \mathbf{T}(Pet) \sqsubseteq LovedByKids \tag{6}$$
$$0.9 \ :: \ \mathbf{T}(Pet) \sqsubseteq Affectionate \tag{7}$$

representing that a typical fish is greyish (2), scaly (3) and not affectionate (4), whereas a typical pet does not live in water (5), is loved by kids (6) and is affectionate (7). Concerning rigid properties, we have that all fishes live in water (1). The logic $\mathbf{T}^{\text{CL}}$ combines the concepts Pet and $Fish$, by using the latter as the HEAD and the former as the MODIFIER. The prototypical Pet-Fish inherits from the prototypical fish the fact that it is scaly and not affectionate, the last one by giving preference to the HEAD since such a property conflicts with the opposite one in the modifier (a typical pet is affectionate). The scenarios in which all the three typical properties of a typical fish are inherited by the combined concept are considered as trivial and, therefore, discarded, as a consequence the property having the lowest degree ($Greyish$ with degree 0.6) is not inherited. The prototypical Pet-Fish inherits from the prototypical pet only property (6), since (5) conflicts with the rigid property (1), stating that all fishes (then, also pet fishes) live in water, whereas (7) is blocked, as already mentioned, by the HEAD/MODIFIER heuristics. Formally, the $Pet \sqcap Fish$-revised knowledge base contains, in addition to the above inclusions, the following ones:

$$0.8 \ :: \ \mathbf{T}(Pet \sqcap Fish) \sqsubseteq Scaly \tag{3'}$$
$$0.8 \ :: \ \mathbf{T}(Pet \sqcap Fish) \sqsubseteq \neg Affectionate \tag{4'}$$
$$0.9 \ :: \ \mathbf{T}(Pet \sqcap Fish) \sqsubseteq LovedByKids \tag{6'}$$

In [30] it has been also shown that reasoning in $\mathbf{T}^{\text{CL}}$ remains in the same complexity class of standard $\mathcal{ALC}$ Description Logics.

3. The Ontological Model of Emotions

In this work, we have applied the $\mathbf{T}^{\text{CL}}$ reasoning framework to the generation of new compound emotions by starting from the affective ontological knowledge base named ArsEmotica, as well as for attributing the most adequate emotion to astronaut's speeches. Intuitively, given the ArsEmotica knowledge base equipped with the prototypical descriptions of basic emotions, we exploit the reasoning capabilities of the logic $\mathbf{T}^{\text{CL}}$ in order to generate new *derived* emotions as the result of the creative combination of two (or even more) basic or derived ones. Moreover, an item of the tested dataset belongs to the new generated emotion if it contains all the rigid properties as well as at least the 30% of the typical properties of such a derived emotion.

More in detail, the affective knowledge leveraged by the $\mathbf{T}^{\text{CL}}$ logic is encoded in an ontology of emotional categories based on Plutchik's psychological circumplex model [11], called ArsEmotica[3] and includes also concepts from the Hourglass model [43]. The ontology structures emotional categories in a taxonomy, which currently includes 32 emotional concepts. The design of the taxonomic structure of emotional categories, of the disjunction axioms and of the object and data properties mirrors the main features of Plutchik model. As already mentioned, such model can be represented as a wheel of emotions (see Figure 1) and encodes the following elements:

- Basic or primary emotions: *Joy*, *Trust*, *Fear*, *Surprise*, *Sadness*, *Disgust*, *Anger*, *Anticipation*; in the color wheel, this is represented by differently colored sectors.
- Opposites: basic emotions can be conceptualized in terms of polar opposites: *Joy* versus *Sadness*, *Anger* versus *Fear*, *Trust* versus *Disgust*, *Surprise* versus *Anticipation*.
- Intensity: each emotion can exist in varying degrees of intensity; in the wheel, this is represented by the vertical dimension.
- Similarity: emotions vary in their degree of similarity to one another; in the wheel, this is represented by the radial dimension.
- Complex emotions: a complex emotion is a composition of two basic emotions; the pair of basic emotions involved in the composition is called a dyad. Looking at the Plutchik wheel, the eight emotions in the blank spaces are compositions of similar basic emotions, called primary dyads. Pairs of less similar emotions are called secondary dyads (if the radial distance between them is 2) or tertiary dyads (if the distance is 3), while opposites cannot be combined.

We have chosen to encode the Plutchik model in the ontology for several reasons: (i) it is well-grounded in psychology and general enough to guarantee a wide coverage of emotions; (ii) the Plutchik wheel of emotions is perfectly compliant with the generative model underlying the $\mathbf{T}^{\text{CL}}$ logic; (iii) it encodes interesting notions, e.g. emotional polar opposites, which can be exploited for finding novel, non obvious relations among contents.

[3]The ArsEmotica ontology is available here: http://130.192.212.225/fuseki/ArsEmotica-core and queryable via SPARQL endpoint at: http://130.192.212.225/fuseki/dataset.html?tab=query&ds=/ArsEmotica-core

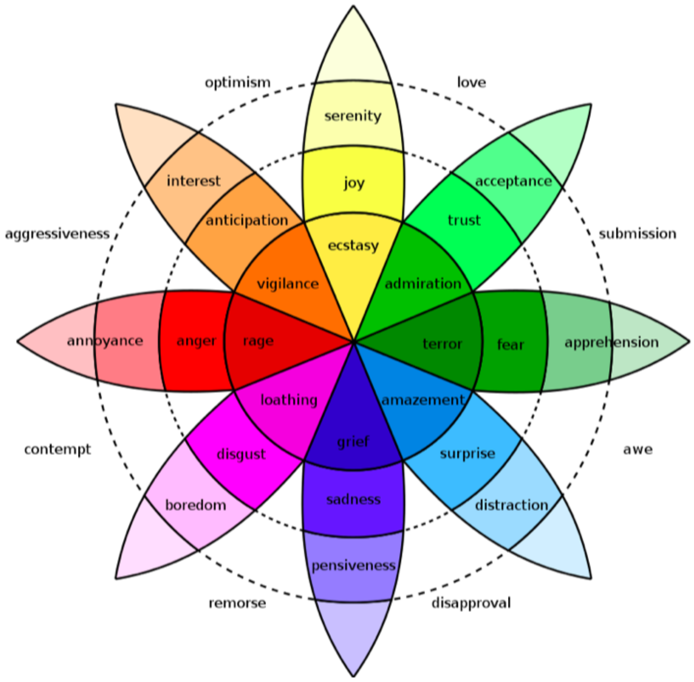

Figure 1: The Wheel of Emotion of the Plutchik Model

Within the ArsEmotica ontology, the class Emotion is the root for all the emotional concepts. The Emotions hierarchy includes all the 32 emotional categories presented as distinguished labels in the model. In particular, the Emotion class has two disjoint subclasses: BasicEmotion and ComplexEmotion. Basic emotions of the Plutchik model are direct sub-classes of BasicEmotion. Each of them is specialized again into two subclasses representing the same emotion with weaker or stronger intensity (e.g. the basic emotion *Joy* has *Ecstasy* and *Serenity* as sub-classes). Therefore, we have 24 emotional concepts subsumed by the BasicEmotion concept. Instead, the class CompositeEmotion has 24 subclasses, corresponding to the primary (*Love, Submission, Awe, Disapproval, Remorse, Contempt, Aggressiveness e Optimism*), secondary (*Hope, Guilt, Curiosity, Despair, Unbelief, Envy, Cynicism e Pride*) and tertiary (*Anxiety, Delight, Sentimentality, Shame, Outrage, Pessimism, Morbidness, Dominance*) dyads. Other relations in the Plutchik model have been expressed in the ontology by means of object properties: the hasOpposite property encodes the notion of polar opposites; the hasSibling property encodes the notion of similarity and the isComposedOf property encodes the notion of composition of basic emotions. Moreover, a data type property hasScore was introduced to link each emotion with an intensity value mapped into the Hourglass model.

The devised model allows inferring complex emotions from basic ones by exploiting simple SWRL rules (i.e. if-then clauses) allowing to infer, from the isComposedOf property connecting Basic and Composite Emotions, the fact that if an agent feels two emotions (suppose for a given item), and if these emotions jointly constitute a Composite Emotion, then the latter emotion will be automatically assigned to the agent in order to better describe his/her aesthetic experience.

Due to the need of modeling the links between words in a language and the emotions they refer to, the ArsEmotica Ontology is also integrated with the ontology framework LExicon Model for ONtologies (LEMON) [44]. In particular, such integration allows differentiating explicitly between the language level (lexicon-based) and the conceptual one in representing the emotional concepts [45]. Within this enriched framework, it is possible to associate a plethora of emotional words, with the encoding of language information, to the corresponding emotional concepts. In this work, we have used the ArsEmotica model of emotional concepts with the NRC Emotion Intensity Lexicon mentioned above [19]. Such lexicon provides a list of English words, each with real-values representing intensity scores for the eight basic emotions of Plutchik's theory. The lexicon includes close to $10,000$ words including terms already known to be associated with emotions as well as terms that co-occur in Twitter posts that convey emotions. The intensity scores were obtained via crowdsourcing, using best-worst scaling annotation scheme. For our purposes, we considered the most frequent terms available in such lexicon (and associated to the basic emotions of the Plutchik wheel) as typical features of such emotions. In this way, once the prototypes of the basic emotional concepts were formed, the $\mathbf{T}^{\mathsf{CL}}$ reasoning framework was used to generate the compound emotions.

Let us now provide some details about the construction of the ontology of emotions in $\mathbf{T}^{\mathsf{CL}}$. As in [26], prototypes generation proceeds in two steps: in the first one, the system builds a prototypical description of basic emotions in the language of the logic $\mathbf{T}^{\mathsf{CL}}$, in order to describe their typical properties. In detail, a knowledge base in the logic $\mathbf{T}^{\mathsf{CL}}$ characterized by typicality inclusions of the form

$$p \ :: \ \mathbf{T}(BasicEmotion) \sqsubseteq Property$$

where $BasicEmotion$ is one of the eight basic emotions of the Plutchik model: Joy, Trust, Fear, Surprise, Sadness, Disgust, Anger, and Anticipation. As an example, consider the basic emotion Joy. The words having the highest scores are happiness (0.98), bliss (0.97), to celebrate (0.97), jubilant (0.97), ecstatic (0.95), and euphoria (0.94). Therefore, the knowledge base generated will contain, among others, the following inclusions:

$$\begin{array}{l}
Joy \sqsubseteq \neg Holocaust \\
0.98 \ :: \ \mathbf{T}(Joy) \sqsubseteq Happiness \\
0.97 \ :: \ \mathbf{T}(Joy) \sqsubseteq Bliss \\
0.97 \ :: \ \mathbf{T}(Joy) \sqsubseteq Celebrating \\
0.97 \ :: \ \mathbf{T}(Joy) \sqsubseteq Jubilant \\
0.95 \ :: \ \mathbf{T}(Joy) \sqsubseteq Ecstatic \\
0.94 \ :: \ \mathbf{T}(Joy) \sqsubseteq Elation
\end{array}$$

As a second step, the system exploits the above described reasoning mechanism of such a Description Logic in order to combine the prototypical descriptions of pairs of basic emotions, generating the prototypical description of compound emotions by using the same logical procedure of the *pet-fish* problem. As an example, let us consider the combination of the above basic emotion Joy with $Fear$, whose prototypical description is as follows:

$$\begin{array}{l}
0.96 \ :: \ \mathbf{T}(Fear) \sqsubseteq Kill \\
0.95 \ :: \ \mathbf{T}(Fear) \sqsubseteq Annihilate \\
0.95 \ :: \ \mathbf{T}(Fear) \sqsubseteq Terror \\
0.98 \ :: \ \mathbf{T}(Fear) \sqsubseteq Torture \\
0.97 \ :: \ \mathbf{T}(Fear) \sqsubseteq Terrorist \\
0.97 \ :: \ \mathbf{T}(Fear) \sqsubseteq Horrific
\end{array}$$

In order to obtain a description of the compound emotion $Guilt$ as the result of the combination of the two basic emotions ($Joy \sqcap Fear$) in the logic $\mathbf{T}^{\mathsf{CL}}$, the system combines the two basic emotions by implementing a variant of CoCoS [46], a Python implementation of reasoning services for the logic $\mathbf{T}^{\mathsf{CL}}$ in order to exploit efficient DLs reasoners for checking both the consistency of each generated scenario and the existence of conflicts among properties, following the line of the system DENOTER [47] and DEGARI [26]. CoCoS generates scenarios and chooses the selected one(s) by exploiting the translation of an $\mathcal{ALC} + \mathbf{T_R}$ knowledge base into standard $\mathcal{ALC}$ introduced in [39] and adopted by the system RAT-OWL [48]. CoCoS makes use of the above mentioned library `owlready2`[4], which allows relying on the services of efficient DL reasoners, e.g. the HermiT reasoner.

CoCoS is embedded in DEGARI and allows one: i) to include the logical descriptions of the concepts to be combined; ii) to select which among the concepts has to be intended as HEAD and as MODIFIER(s); iii) to choose how many typical properties one wants to inherit in the scenarios that will be selected by

[4]https://pythonhosted.org/Owlready2/

$\mathbf{T}^{\text{CL}}$. In addition to presenting the selected scenario with typical properties of the combined concept, CoCoS also allows the users to select alternative scenarios, ranging from more trivial to more surprising ones. More in detail, the system considers both the available choices for the HEAD and the MODIFIER, and it allows restricting its concern to a given and fixed number of inherited properties. The combined emotion $Guilt$ has the following $\mathbf{T}^{\text{CL}}$ description (concept $Joy \sqcap Fear$):

$$
\begin{array}{l}
0.98 \;::\; \mathbf{T}(Joy \sqcap Fear) \sqsubseteq Happiness \\
0.97 \;::\; \mathbf{T}(Joy \sqcap Fear) \sqsubseteq Celebrating \\
0.97 \;::\; \mathbf{T}(Joy \sqcap Fear) \sqsubseteq Bliss \\
0.98 \;::\; \mathbf{T}(Joy \sqcap Fear) \sqsubseteq Torture \\
0.97 \;::\; \mathbf{T}(Joy \sqcap Fear) \sqsubseteq Terrorist \\
0.97 \;::\; \mathbf{T}(Joy \sqcap Fear) \sqsubseteq Horrific
\end{array}
$$

Obviously, rigid properties of basic emotions (if any) are inherited by the compound emotion (in the example, $Joy \sqcap Fear \sqsubseteq \neg Holocaust$), and this retain the system from considering any inconsistent typical properties even if they have the highest probability.

It is worth noticing that the properties of the derived emotion are still expressed in the language of the logic $\mathbf{T}^{\text{CL}}$, therefore the combined emotion, $Guilt$ in the example, can be further combined with another emotion, in order to iterate the procedure.

4. The Tool for Emotion Attribution in Aerospace Missions

In this section we describe the system exploiting the logic $\mathbf{T}^{\text{CL}}$ on the ArsEmotica knowledge base in order to detect the emotions of astronauts involved in aerospace missions. In order to detect the most adequate emotion to assign to a given utterance, the systems computes the following steps:

- Utterances are first processed by a tokenizer, whose objective is to extract relevant words. The first step is about determining whether an utterance contains (or not) terms used to describe basic emotions and identifying them.
- When the system finds some of these terms in an utterance, it considers the score each term has paired, and it sums up all the scores for each of the eight basic emotions. This way we get indicative values concerning which basic emotions are found in each utterance. It's possible to implement a more suitable function rather than the sum, for example a weighted mean considering the total number of words in the utterance we are analyzing.
- A very similar process is carried out in order to identify combined emotions: instead of looking for the basic emotion's terms, this time the system looks for the terms obtained from the combination of the two basic emotions, considering their scores.

Transcriptions belong to real missions and are taken from `NASA Mission Transcript Collection`[5]. It is worth mentioning that we classify some terms used to look for emotions as "technical", meaning the pilots using those specific terms are referring to concepts that we do not consider interesting in our analysis: as an example, the word "Buzz" appears in the description of the emotion *Anticipation*, but in the mission Apollo 11 Buzz Aldrin was one of the astronauts involved, so in this specific case we avoid this term in the emotion attribution. To fix this, we have adopted a list of terms (it can also be empty) for the system to ignore during the analysis.

Preliminary results seem to be promising. On the one hand, we observed that the system is able to correctly identify the emotions to be attributed to each processed text. On the other hand, we have found that all the derived emotions, obtained through combinations, are used to label at least one utterance, which supports the idea that the proposed approach may be useful for the intended purpose.

As an example, let us consider the transcription of an utterance by the Lunar module pilot (LMP) of the Apollo 11 mission. The system produces the following json. The first three fields exhibit the

[5]https://historycollection.jsc.nasa.gov/JSCHistoryPortal/history/mission_trans/mission_transcripts.htm

information extracted by the transcription under consideration (“time”, “speaker” and “text”), then the “occurrences” field shows the relevant terms found in the text. Next, we have a simple word counter regarding the “text” field. After that, the system shows the score obtained for the utterance by each of the eight basic emotions, as well as by combined emotions, along with the matched words:


```
{
    "time": "00 01 25 44",
     "speaker": "LMP",
     "text": "Sun is bright, isn't it? ... It's a pretty nice camera,
                    to tell you the truth.",
     "occurrences": ["sun", "pretty", "truth"],
     "word_count": 15,
     "anger-NRC-Emotion-Intensity-Lexicon-v1": 0.0,
     "joy-NRC-Emotion-Intensity-Lexicon-v1": 0.735,
     "fear-NRC-Emotion-Intensity-Lexicon-v1": 0.0,
     "anticipation-NRC-Emotion-Intensity-Lexicon-v1": 0.859,
     "sadness-NRC-Emotion-Intensity-Lexicon-v1": 0.0,
     "surprise-NRC-Emotion-Intensity-Lexicon-v1": 0.172,
     "trust-NRC-Emotion-Intensity-Lexicon-v1": 1.649,
     "disgust-NRC-Emotion-Intensity-Lexicon-v1": 0.0,
     "combined_emotions": [
      {
        "emotion": "anticipation-trust",
        "matched_words": [{"word": "truth", "score": 0.844}]
      },
      {
        "emotion": "trust-anticipation",
        "matched_words": [{"word": "truth", "score": 0.844}]
      }
    ]
  }
```


In this example, the basic emotion with the highest score is *Trust*, whose representation in ArsEmotica contains the following inclusions :

$$
\begin{aligned}
0.906 &:: \mathbf{T}(\mathit{Trust}) \sqsubseteq \mathit{Truthfulness}\\
0.883 &:: \mathbf{T}(\mathit{Trust}) \sqsubseteq \mathit{Trusted}\\
0.844 &:: \mathbf{T}(\mathit{Trust}) \sqsubseteq \mathit{Truth}\\
0.438 &:: \mathbf{T}(\mathit{Trust}) \sqsubseteq \mathit{Sun}\\
0.367 &:: \mathbf{T}(\mathit{Trust}) \sqsubseteq \mathit{Pretty}
\end{aligned}
$$

It is easy to understand that the presence of words “sun”, “truth” and “pretty” in the analyzed text are responsible of an high score given the description here above, where the last three typicality inclusions exhibit such properties, even with high probabilities. In this example, the fact that also the score for the emotion *Anticipation* is significant, it follows that the two compound emotions suggested by the system are *Anticipation-Trust* and *Trust-Anticipation*, whose descriptions in ArsEmotica follow:

$$
\begin{aligned}
0.859 &:: \mathbf{T}(\mathit{Trust} \sqcap \mathit{Anticipation}) \sqsubseteq \mathit{Anticipation}\\
0.859 &:: \mathbf{T}(\mathit{Trust} \sqcap \mathit{Anticipation}) \sqsubseteq \mathit{Excited}\\
0.820 &:: \mathbf{T}(\mathit{Trust} \sqcap \mathit{Anticipation}) \sqsubseteq \mathit{Excitement}\\
0.820 &:: \mathbf{T}(\mathit{Trust} \sqcap \mathit{Anticipation}) \sqsubseteq \mathit{Anticipate}\\
0.844 &:: \mathbf{T}(\mathit{Trust} \sqcap \mathit{Anticipation}) \sqsubseteq \mathit{Truth}
\end{aligned}
$$

$$
\begin{aligned}
0.906 &:: \mathbf{T}(\mathit{Anticipation} \sqcap \mathit{Trust}) \sqsubseteq \mathit{Truthfulness}\\
0.883 &:: \mathbf{T}(\mathit{Anticipation} \sqcap \mathit{Trust}) \sqsubseteq \mathit{Trusted}\\
0.867 &:: \mathbf{T}(\mathit{Anticipation} \sqcap \mathit{Trust}) \sqsubseteq \mathit{Trustworthy}\\
0.844 &:: \mathbf{T}(\mathit{Anticipation} \sqcap \mathit{Trust}) \sqsubseteq \mathit{Honor}\\
0.844 &:: \mathbf{T}(\mathit{Anticipation} \sqcap \mathit{Trust}) \sqsubseteq \mathit{Truth}
\end{aligned}
$$

As a further example, consider the following transcription by the Command module pilot (CMP):


```
{
  "time": "05 11 04 25",
   "speaker": "CMP",
   "text": "Well, you know, - well, just looking at that one sample,
          it was - I'm surprised you didn't have a lot more dust. Now you saw
          dust during descent, I think, around 40 feet, something like that,
         30 feet maybe.",
   "occurrences": ["surprised","descent","like"],
   "word_count": 38,
   "anger-NRC-Emotion-Intensity-Lexicon-v1": 0.0,
   "joy-NRC-Emotion-Intensity-Lexicon-v1": 0.0,
   "fear-NRC-Emotion-Intensity-Lexicon-v1": 0.391,
   "anticipation-NRC-Emotion-Intensity-Lexicon-v1": 0.0,
   "sadness-NRC-Emotion-Intensity-Lexicon-v1": 0.377,
   "surprise-NRC-Emotion-Intensity-Lexicon-v1": 0.867,
   "trust-NRC-Emotion-Intensity-Lexicon-v1": 0.484,
   "disgust-NRC-Emotion-Intensity-Lexicon-v1": 0.0,
   "combined_emotions": [
    {
      "emotion": "anticipation-surprise",
      "matched_words": [
        {
          "word": "surprised",
          "score": 0.867
        }
      ]
    },
    {
      "emotion": "surprise-anticipation",
      "matched_words": [
        {
          "word": "surprised",
          "score": 0.867
        }
      ]
    }
  ]
}
```


In this example, the compound emotions with the highest scores are *Anticipation-Surprise* and *Surprise-Anticipation*, even if *Fear*, *Sadness* and *Trust* have higher scores among basic ones. This is a consequence of the application of reasoning mechanisms of the logic $\mathbf{T}^{\mathsf{CL}}$, which allows to obtain the following description of such compound emotions.

$$
\begin{aligned}
0.906 &:: \mathbf{T}(\mathit{Anticipation} \sqcap \mathit{Surprise}) \sqsubseteq \mathit{Explode}\\
0.906 &:: \mathbf{T}(\mathit{Anticipation} \sqcap \mathit{Surprise}) \sqsubseteq \mathit{Flabbergast}\\
0.898 &:: \mathbf{T}(\mathit{Anticipation} \sqcap \mathit{Surprise}) \sqsubseteq \mathit{Explosion}\\
0.883 &:: \mathbf{T}(\mathit{Anticipation} \sqcap \mathit{Surprise}) \sqsubseteq \mathit{Ambush}\\
0.867 &:: \mathbf{T}(\mathit{Anticipation} \sqcap \mathit{Surprise}) \sqsubseteq \mathit{Surprised}
\end{aligned}
$$

$$
\begin{aligned}
0.906 &:: \mathbf{T}(\mathit{Surprise} \sqcap \mathit{Anticipation}) \sqsubseteq \mathit{Explode}\\
0.906 &:: \mathbf{T}(\mathit{Surprise} \sqcap \mathit{Anticipation}) \sqsubseteq \mathit{Flabbergast}\\
0.898 &:: \mathbf{T}(\mathit{Surprise} \sqcap \mathit{Anticipation}) \sqsubseteq \mathit{Explosion}\\
0.883 &:: \mathbf{T}(\mathit{Surprise} \sqcap \mathit{Anticipation}) \sqsubseteq \mathit{Ambush}\\
0.867 &:: \mathbf{T}(\mathit{Surprise} \sqcap \mathit{Anticipation}) \sqsubseteq \mathit{Surprised}
\end{aligned}
$$

Running on the transcription of the Apollo 11 flightcrew communications as recorded on the command module, containing a total of 4914 utterances, our system identified at least one of the eight basic emotions in 1665 utterances. The system also identified complex emotions in 7 utterances.

5. Conclusions and Future Works

In this work, we introduced a tool for the automatic identification of emotions characterizing astronauts' speeches during missions. The system is based on the Description Logic of typicality $\mathbf{T}^{\text{CL}}$, introduced in [28, 29] as a probabilistic extension of the basic formalism able to deal with the combination of concepts. Intuitively, the system checks words/concepts belonging to astronaut's utterances with the properties characterizing both basic and derived emotions described in the language of the $\mathbf{T}^{\text{CL}}$ logic, by also considering probabilities equipping each typicality inclusion as suitable weights of such properties. The ontology with typicality and probabilities describing emotions and related properties, called ArsEmotica [26], is based on Plutchik's psychological circumplex model [11].

The study is supported by a very preliminary evaluation conducted on reports from real missions contained in the dataset `NASA Mission Transcript Collection`[6], although the results are decidedly encouraging. Our goal is to conduct a more structured evaluation of the system, on the one hand by comparing the system's classification with a zero-shot classification using well-known Large Language Models, and on the other by performing further tests using the results of mission simulations.

This work is intended as the first step towards an intelligent system for the autonomous detection of the psycho-physical state of astronauts and aerospace crew members in general. Our goal is to develop a multimodal system able to take into account different physiological signals (heart rate, heart rate variability, eye tracking, keystroke dynamics, speech, etc.). The present work only uses textual data, taking advantage of the availability of official transcripts from space missions and the wide experience in textual analysis of the Natural Language Processing community. Since our envisioned system is intended to be used in real time, other aspects of the astronauts' speech could be used to feed the analysis: prosodic, spectral and wavelet features can be informative about the intentions of the speaker and of their emotional state [49].

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] G. Froger, C. Blättler, E. Dubois, C. Camachon, N. Bonnardel, Time-Interval Emphasis in an Aeronautical Dual-Task Context: A Countermeasure to Task Absorption, Human Factors: The Journal of the Human Factors and Ergonomics Society 60 (2018) 936–946. Publisher: SAGE Publications.

[2] G. Borghini, L. Astolfi, G. Vecchiato, D. Mattia, F. Babiloni, Measuring neurophysiological signals in aircraft pilots and car drivers for the assessment of mental workload, fatigue and drowsiness, Neuroscience & Biobehavioral Reviews 44 (2014) 58–75. Publisher: Elsevier BV.

[3] A. Hamann, N. Carstengerdes, Investigating mental workload-induced changes in cortical oxygenation and frontal theta activity during simulated flights, Scientific Reports 12 (2022). Publisher: Springer Science and Business Media LLC.

[4] R. J. Gentili, J. C. Rietschel, K. J. Jaquess, Li-Chuan Lo, C. M. Prevost, M. W. Miller, J. M. Mohler, Hyuk Oh, Ying Ying Tan, B. D. Hatfield, Brain biomarkers based assessment of cognitive workload in pilots under various task demands, in: IEEE 2014, IEEE, Chicago, IL, 2014, pp. 5860–5863.

[5] J. A. Blanco, M. K. Johnson, K. J. Jaquess, H. Oh, L.-C. Lo, R. J. Gentili, B. D. Hatfield, Quantifying Cognitive Workload in Simulated Flight Using Passive, Dry EEG Measurements, IEEE Transactions on Cognitive and Developmental Systems 10 (2018) 373–383.

[6] S. Grissmann, J. Faller, C. Scharinger, M. Spüler, P. Gerjets, Electroencephalography Based Analysis of Working Memory Load and Affective Valence in an N-back Task with Emotional Stimuli, Frontiers in Human Neuroscience 11 (2017). Publisher: Frontiers Media SA.

[6] https://historycollection.jsc.nasa.gov/JSCHistoryPortal/history/mission_trans/mission_transcripts.htm

[7] D. Tao, H. Tan, H. Wang, X. Zhang, X. Qu, T. Zhang, A Systematic Review of Physiological Measures of Mental Workload, International Journal of Environmental Research and Public Health 16 (2019) 2716. Publisher: MDPI AG.
[8] J. A. Russell, A circumplex model of affect., J. personality and social psychology 39 (1980) 1161.
[9] D. Watson, L. A. Clark, A. Tellegen, Development and validation of brief measures of positive and negative affect: the PANAS scales., Journal of personality and social psychology 54 (1988) 1063.
[10] A. Mehrabian, Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in temperament, Current Psychology 14 (1996) 261–292.
[11] R. Plutchik, A general psychoevolutionary theory of emotion, in: Theories of emotion, Elsevier, 1980, pp. 3–33.
[12] P. Ekman, Basic emotions, in: M. J. P. Tim Dalgleish (Ed.), Handbook of cognition and emotion, volume 98, Wiley Online Library, 1999, pp. 45–60.
[13] T. Bänziger, K. R. Scherer, Introducing the geneva multimodal emotion portrayal (gemep) corpus, Blueprint for affective computing: A sourcebook 2010 (2010) 271–94.
[14] D. Jurafsky, J. H. Martin, Lexicons for sentiment, affect, and connotation, in: Speech and language processing: an introduction to natural language processing, computational linguistics, and speech recognition, 3rd Edition, 2019.
[15] M. Nissim, V. Patti, Chapter 3 - semantic aspects in sentiment analysis, in: F. A. Pozzi, E. Fersini, E. Messina, B. Liu (Eds.), Sentiment Analysis in Social Networks, 2017, pp. 31 – 48.
[16] Z. Wang, S. Ho, E. Cambria, A review of emotion sensing: categorization models and algorithms, Multimedia Tools and Applications 79 (2020).
[17] M. E. Basiri, S. Nemati, M. Abdar, E. Cambria, U. R. Acharya, Abcdm: An attention-based bidirectional cnn-rnn deep model for sentiment analysis, Future Generation Computer Systems 115 (2021) 279–294.
[18] S. M. Mohammad, Sentiment analysis: Detecting valence, emotions, and other affectual states from text, CoRR abs/2005.11882 (2020). arXiv:2005.11882.
[19] S. Mohammad, Word affect intensities, in: Proceedings of LREC 2018, European Language Resources Association (ELRA), 2018.
[20] M. S. Akhtar, A. Ekbal, E. Cambria, How intense are you? predicting intensities of emotions and sentiments using stacked ensemble [application notes], IEEE Computational Intelligence Magazine 15 (2020) 64–75.
[21] C. Strapparava, R. Mihalcea, Semeval-2007 task 14: Affective text, in: Proc. of SemEval ’07, The Association for Computer Linguistics, 2007, pp. 70–74.
[22] S. M. Mohammad, P. D. Turney, Crowdsourcing a word-emotion association lexicon, Computational Intelligence 29 (2013) 436–465.
[23] S. Mohammad, Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 English words, in: Proc. ACL, 2018, pp. 174–184.
[24] E. Cambria, Y. Li, F. Z. Xing, S. Poria, K. Kwok, Senticnet 6: Ensemble application of symbolic and subsymbolic ai for sentiment analysis, in: M. d’Aquin, S. Dietze, C. Hauff, E. Curry, P. Cudré-Mauroux (Eds.), CIKM ’20, ACM, 2020, pp. 105–114.
[25] C. Strapparava, A. Valitutti, WordNet affect: an affective extension of WordNet, in: Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC’04), European Language Resources Association (ELRA), Lisbon, Portugal, 2004.
[26] A. Lieto, G. L. Pozzato, S. Zoia, V. Patti, R. Damiano, A commonsense reasoning framework for explanatory emotion attribution, generation and re-classification, Knowledge Based Systems 227 (2021) 107166.
[27] A. Lieto, G. L. Pozzato, M. Striani, S. Zoia, R. Damiano, DEGARI 2.0: A diversity-seeking, explainable, and affective art recommender for social inclusion, Cog. Syst. Res. 77 (2023) 1–17.
[28] A. Lieto, G. L. Pozzato, A description logic of typicality for conceptual combination, in: M. Ceci, N. Japkowicz, J. Liu, G. A. Papadopoulos, Z. W. Ras (Eds.), Foundations of Intelligent Systems - 24th International Symposium, ISMIS 2018, Limassol, Cyprus, October 29-31, 2018, Proceedings, volume 11177 of *Lecture Notes in Computer Science*, Springer, 2018, pp. 189–199.

[29] A. Lieto, G. L. Pozzato, A description logic framework for commonsense conceptual combination integrating typicality, probabilities and cognitive heuristics, Journal of Experimental and Theoretical Artificial Intelligence 32 (2020) 769–804.

[30] A. Lieto, G. L. Pozzato, A description logic framework for commonsense conceptual combination integrating typicality, probabilities and cognitive heuristics, Journal of Experimental & Theoretical Artificial Intelligence 32 (2020) 769–804.

[31] M. A. Boden, Creativity and artificial intelligence, Artificial Intelligence 103 (1998) 347–356.

[32] M. Frixione, A. Lieto, Representing and reasoning on typicality in formal ontologies, in: C. Ghidini, A. N. Ngomo, S. N. Lindstaedt, T. Pellegrini (Eds.), Proceedings of the 7th International Conference on Semantic Systems, ACM International Conference Proceeding Series, ACM, 2011, pp. 119–125.

[33] D. N. Osherson, E. E. Smith, On the adequacy of prototype theory as a theory of concepts, Cognition 9 (1981) 35–58.

[34] A. Lieto, G. L. Pozzato, F. Perrone, E. Chiodino, Knowledge capturing via conceptual reframing: A goal-oriented framework for knowledge invention, in: M. Kejriwal, P. A. Szekely, R. Troncy (Eds.), Proceedings of K-CAP 2019, Marina del Rey, ACM, 2019, pp. 109–114.

[35] A. Lieto, F. Perrone, G. L. Pozzato, E. Chiodino, Beyond subgoaling: A dynamic knowledge generation framework for creative problem solving in cognitive architectures, Cognitive Systems Research 58 (2019) 305–316.

[36] E. Chiodino, A. Lieto, F. Perrone, G. L. Pozzato, A goal-oriented framework for knowledge invention and creative problem solving in cognitive architectures, in: Proc. of ECAI 2020, volume 325, IOS Press, 2020, pp. 2893–2894.

[37] A. Lieto, G. L. Pozzato, Applying a description logic of typicality as a generative tool for concept combination in computational creativity, Intelligenza Artificiale 13 (2019) 93–106.

[38] F. Riguzzi, E. Bellodi, E. Lamma, R. Zese, Probabilistic description logics under the distribution semantics, Semantic Web 6 (2015) 477–501.

[39] L. Giordano, V. Gliozzi, N. Olivetti, G. L. Pozzato, Semantic characterization of Rational Closure: from Propositional Logic to Description Logics, Artificial Intelligence 226 (2015) 1–33.

[40] D. Lehmann, M. Magidor, What does a conditional knowledge base entail?, Artificial Intelligence 55 (1992) 1–60.

[41] F. Riguzzi, E. Bellodi, E. Lamma, R. Zese, Reasoning with probabilistic ontologies, in: Proc. of IJCAI 2015, AAAI Press, 2015, pp. 4310–4316.

[42] J. A. Hampton, Inheritance of attributes in natural concept conjunctions, Memory & Cognition 15 (1987) 55–71.

[43] E. Cambria, A. Livingstone, A. Hussain, The hourglass of emotions, in: A. Esposito, A. M. Esposito, A. Vinciarelli, R. Hoffmann, V. C. Müller (Eds.), COST 2102, volume 7403 of *Lecture Notes in Computer Science*, Springer, 2012, pp. 144–157.

[44] J. P. McCrae, D. Spohr, P. Cimiano, Linking lexical resources and ontologies on the semantic web with lemon, in: G. Antoniou, M. Grobelnik, E. P. B. Simperl, B. Parsia, D. Plexousakis, P. D. Leenheer, J. Z. Pan (Eds.), ESWC 2011, volume 6643 of *Lecture Notes in Computer Science*, Springer, 2011, pp. 245–259.

[45] V. Patti, F. Bertola, A. Lieto, Arsemotica for arsmeteo. org: Emotion-driven exploration of online art collections, in: I. Russell, W. Eberle (Eds.), The Twenty-Eighth International Florida Artificial Intelligence Research Society Conference (FLAIRS 2015), Association for the Advancement of Artificial Intelligence, AAAI Press, 2015, pp. 288–293.

[46] A. Lieto, G. Pozzato, A. Valese, COCOS: a typicality based COncept COmbination System , in: M. Montali, P. Felli (Eds.), Proceedings of the 33rd Italian Conference on Computational Logic (CILC 2018), CEUR Workshop Proceedings, Bozen, Italy, 2018, pp. 55–59.

[47] E. Chiodino, D. D. Luccio, A. Lieto, A. Messina, G. L. Pozzato, D. Rubinetti, A knowledge-based system for the dynamic generation and classification of novel contents in multimedia broadcasting, in: Proc. ECAI 2020, volume 325 of *FAIA*, IOS Press, 2020, pp. 680–687.

[48] L. Giordano, V. Gliozzi, G. L. Pozzato, R. Renzulli, An efficient reasoner for description logics of typicality and rational closure, in: Proc. of DL 2017, volume 1879 of *CEUR Workshop Proceedings*,

CEUR-WS.org, 2017.
[49] M. B. Akçay, K. Oğuz, Speech emotion recognition: Emotional models, databases, features, preprocessing methods, supporting modalities, and classifiers, Speech Communication 116 (2020) 56–76.