Enrica Troiano

Where are Emotions in Text? A Human-based and Computational Investigation of Emotion Recognition and Generation

Abstract

Natural language processing (NLP) boasts a vibrant tradition of emotion studies, unified by the aim of developing systems that generate and recognize emotions in language. The computational approximation of these two capabilities, however, still faces fundamental challenges, as there is no consensus on how emotions should be processed, particularly in text: application-driven works often lose sight of foundational theories that describe how humans communicate what they feel, resulting in conflicting premises about the type of data best suited for modeling and whether this modeling should focus on textual meaning or style. My thesis fills in these theoretical gaps that hinder the creation of emotion-aware systems, demonstrating that a trans-disciplinary approach to emotions, which accounts for their extra-linguistic characteristics, has the potential to improve their computational processing. I investigate the human ability to detect emotions in written productions, and explore the linguistic dimensions that contribute to the emergence of emotions through text. In doing so, I clarify the possibilities and limits of automatic emotion classifiers and generators, also providing insights into where systems should model affective information.

1 Four Problems of Emotions for the NLP Researcher

"The world has changed far more in the past 100 years than in any other century in history". With these words Stephen Hawking (1999) praised the technological transformations that disrupted the 1900s. To name a few: the theory of general relativity shifted the understanding of the structure of the universe (Einstein, 1915); and discoveries on DNA sequencing allowed to clone mammals (Venter et al., 2001). Yet, while scientists cracked the principles of life at huge and microscopic scales, from the height of cosmic bodies down to the level of genome, their progress has been less conclusive on some seemingly simpler and human-sized things – things so deeply ingrained in us that they have fascinated the thinkers of all times, and that, in fact, everybody knows about. These things, which are still ripe for study today, are emotions, and they constitute the topic of my thesis.

The word *emotion* conceals hundreds of subjective experiences that differ in terms of how they feel and when they arise (e.g., awe, boredom, fear, and so on). It comes at no surprise, then, that research on the topic also shows great diversity. Emotion theories abound in many disciplines, from ethology to neuroscience (Tooby & Cosmides, 1990;

Wierzbicka, 1992, i.a.). What is remarkable, though, is that scholars converge on only a few insights, fundamentally disagreeing on what they are looking at. For example, the idea that this part of our evolutionary heritage serves to interact with the environment (Roseman, 1984) is pretty much undisputed, but it is claimed for disjunctive sets of emotions. Scientists even debate on the definition of their phenomenon of interest, to the extent that an article authored by psychologist Scherer in 2005 confronted an apparently basic and yet unresolved question: ultimately, what are emotions?

Emotions and Language. As part of this cross-disciplinary and ongoing debate, my thesis investigates emotions in relation to language. The link between the two modules of human intelligence opens up countless research opportunities, because emotions are not only felt, as personal episodes that influence perception (Brosch, Scherer, Grandjean, & Sander, 2013), behavior (Bach & Dayan, 2017), and judgment making (Nussbaum, 2004). They are also talked about, evoked and stirred up with words, thus pervading the sphere of inter-personal communication that exposes (more or less directly) what we or other individuals feel.

Notably, investigating verbal data gives reasons to reiterate Scherer's concern: what are emotions *in text*? However relevant for fields focused on both language and emotions, that question has been mostly neglected – in fact, outweighed by the study of emotions in other channels, like the body (Kleinsmith & Bianchi-Berthouze, 2012). I rise to the challenge of answering it in my thesis.

Approach and Contributions. My dissertation uses primarily the tools of computational emotion analysis in natural language processing, a research area striving to replicate the abilities that humans exhibit in their linguistic practices, via systems that, e.g., identify, quantify, and generate affective states. Combined with theories from linguistics and psychology, such an approach allows to conduct a critical inquiry of natural and artificial emotions, namely, to examine how emotions are written about and are recognized by humans, and in parallel, how NLP systems perform the same tasks for text generation and classification.

Like other emotion-centered disciplines, however, computational emotion analysis suffers from a lack of a unified framework, fragmented as it is into disparate approaches that engage in engineering rather than theoretical advances. I identify its crucial knowledge gaps and treat them as starting points for my discussion. The gaps concern: (1) the types of texts where emotions are communicated; (2) the way in which their interpretation, often subjective and variable, takes place; and the contribution of emotions in determining (3) how texts are written and (4) what texts say. Therefore, more than aiming to build powerful emotion-aware systems, I leverage the computational techniques of NLP as means to observe emotions in language, and to lay down a theoretically-informed foundation for future (more purely) computational endeavors.

The core chapters of the thesis deal with separate gaps. This way, I consider my driving question from four perspectives, which begin from a "distance" that includes

Chapter	Question: Emotions are ...	Task	Agents	Tools: NLP and ...
3	in what texts?	Generation, Recognition	Humans	Psychology
4	in what factors?	Recognition	Humans	Psychology
5	in textual style?	Generation, Recognition	Artificial	Linguistics
6	in textual meaning?	Recognition	Artificial	Psychology, Linguistics

Table 1: Structured overview of the thesis.

both the verbal material in which emotions emerge (i.e., texts) and the agents that process it (i.e., writers and readers), to then zoom in the texts alone:

- Chapter 3 asks if emotions are phenomena that people effectively extract from expressions without emotional words. (I find that this is the case.)
- Chapter 4 asks if emotions, as interpreted by readers, are dimensions of text that interact with elements exceeding language. (This idea proves correct.)
- Chapter 5 asks if emotions amount to a role of linguistic style and style only, as they affect the way in which things are said. (This assumption turns out unconvincing.)
- Chapter 6 asks the opposite, to understand if emotions are part of a text's semantics, in that they enter the very content that language communicates. (Experimental results provide supporting evidence.)

Overall, in the context of NLP, Scherer's question can be repurposed this way: where are emotions (generated and recognized) in text? Which is to ask, in what types of texts (Chapter 3), in what factors (Chapter 4), and at what linguistic level (Chapter 5 and 6) are they processed by people, and can accordingly be processed by automatic systems? Table 1 details the questions, agents, tasks and tools (for experimentation or discussion) that different chapters use to give an answer.

The corresponding findings form an all-round picture of emotions. (1) Emotions characterize even factual expressions that have no explicitly emotional tone, which merely describe events (hence, they can be legitimately studied from there). (2) Factors that extend to the readers' personal characteristics cause variability in how such emotions are interpreted (this requires us to carefully think about what it means for an emotion classifiers to be successful). Lastly, emotions are loaded (3) not in the shallow layer of a text wording but (4) in its meaning (particularly, and to close the circle, in the meaning of events, where they are to be scrutinized).

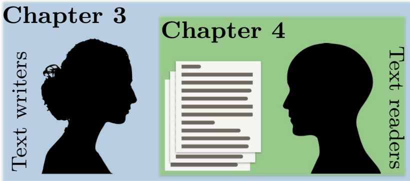

Figure 1: Experimental framework in Chapter 3 and 4.

2 Zooming Out: Emotions in a Relational Perspective

For NLP systems, sensing textual emotions like humans do is a demanding task, partly because emotions can be signalled covertly (e.g., "His face twisted into a grimace" suggests but does not mention an affective state), and partly because they are subjective (e.g., the example above could evoke fear, disgust, and other emotions). In this light, it is hard to set expectations for what a successful emotion classifier should learn. Facing the issue upstream, I question how good humans are at inferring emotions – and whether this question can be meaningfully posed at all: I focus mainly on covert (or implicit) expressions, to see (Chapter 3) if they are suitable data to investigate people's emotion recognition ability (and thus to conduct modeling tasks), and (Chapter 4) if these expressions convey emotions based solely on text, or on an interplay of factors both in language and out of it.

This part of my thesis not only serves as a pre-requisite to inform computational models, but also makes a methodological contribution. I set best practices to exploit human emotion knowledge in the creation and analysis of data, with a new experimental paradigm that approximates real-life communicative scenarios and puts emotions in their inter-personal, relational "habitat". Such a paradigm zooms out of data (cf. Figure 1), assuming a broader viewpoint that optimally includes all agents that create and interpret text (i.e., writers and readers, Chapter 3); alternatively, it looks at one side of this configuration of participants (i.e., readers, Chapter 4), but at the cost of also moving the research lenses from emotion judgments to information on who makes them.

2.1 Chapter 3

This chapter presents two crowdsourcing activities to collect short event descriptions, as extremely implicit expressions of emotions. Both activities follow a class of psychological models that seamlessly fit the study of these texts, because they link emotions to *appraisals*, i.e., evaluations of features of events, including (but not limited to) their suddenness, pleasantness, and relevance (Scherer, 2005). In essence, appraisal-based theories explain why certain emotions arise in certain circumstances – e.g., fear occurs when an event is perceived as threatening and sudden, joy when a pleasant event aligns with one's goals.

Textual descriptions are generated by crowdworkers who recount a situation in which they felt a given emotion; later, other crowdworkers read the texts to decode that emotion: this way, I compile enISEAR, deISEAR, and crowd-enVENT. enISEAR and deISEAR together form the first event-centered multilingual corpus (1001 texts in English and German, respectively) labeled with 7 emotions, as experienced by writers and reconstructed by readers. crowd-enVENT contains 6000 event descriptions in English annotated with 13 emotions, 21 appraisals, and several personal factors (e.g., demographics) by text writers and readers.

This plentiful of annotation allows me to conduct multiple comparisons: between writers and readers, to see if their “interaction” causes a loss of emotion information; between the emotion and appraisal judgments of a text, to analyze if specific emotion labels correspond to specific event evaluations; and between languages, to appreciate their effect on emotion inferences.

An analysis of inter-annotator agreement (IAA) between writers and readers in crowd-enVENT (IAA=.49, as measured via F_1 scores) shows that factual descriptions do not fatally undermine the ability to infer emotions – an ability that is mirrored by classifiers trained and tested on the same corpus (e.g., for boredom, which is the best recognized emotion, a RoBERTa-based model achieves F_1=.84). The emotions that humans decode, however, do not always correspond to those that the writers referred to (irrespective of language).

This insight is expectable, given the subjectivity of the task. But it is only in virtue of the proposed experimental design (including all participants in the transmission of emotion signals) that one can measure *how well* humans recognize emotions, and assess if they are reliable sources of the very information from which automatic systems learn. Most importantly, the many layers of annotation prove useful in my communication-like framework. I cross-analyze appraisals with emotion annotations, finding that the former render transparent why the readers made specific emotion choices: there are regularities between patterns of appraisal and emotions. Hence, while emotions undergo changes in their transmission (e.g., original: fear, interpreted: joy), the way they are perceived appears pertinent, as motivated by specific underlying event evaluations that are different from the writers' but congruous with the chosen emotion label.

In sum, this chapter endows computational emotion analysis with good reasons to use implicit emotion expressions for its tasks, but also to adhere more closely to theories of emotions: if layers of appraisals reveal the behind-the-scene of the emotion judgments that people provide, they can also guide classifiers in identifying both what emotion is in a text and why.

2.2 Chapter 4

When many people read a text, they can end up formulating different emotion interpretations. That is natural, because these interpretations are due to idiosyncrasies like culture, personal values, and past knowledge (e.g., on how certain events feel). For computational emotion analysis, however, judgment diversity is troublesome. It reflects in low IAA scores that imply bad quality data and defy the learning of automatic systems. Chapter 4 faces this "curse of disagreement", proving that the quality of data must be assessed by asking if differences in emotion annotations are random (thus, symptomatic of unreliability) or consistent (in which case, they hint at worldviews or other factors characterizing people).

I study the emotion judgments of readers in a sample of the Corpus of Contemporary American English (Davies, 2015) and crowd-enVENT from Chapter 3. In both, I compute IAA for separate subsets of data, each annotated by readers sharing specific features (according to the information they gave at annotation time), to observe if and how such features influence judgment variability.

Far from being random, (dis)agreements turn out structured. Notable patterns include: readers tend to disagree on the perceived emotions when they attribute disparate qualities to the events described in text, or when they differ by demographic traits, like age. Vice versa, they achieve higher IAA on subsets of data where they perceive more intense emotions, or are more confident about the correctness of their emotion judgments (cf. the substantial boost of Fleiss'κ as annotation confidence increases in Table 2).

Confidence	IAA
Low	$-.001$
Medium	.03
High	.39

Table 2: Fleiss'κ for data subsets with different annotation confidence.

Not all possible factors that underlie disagreements lend themselves to easy examination, but my outcomes demonstrate that many can (in fact, should) be measured to complement emotion annotations. Learning when they systematically correlate with emotion choices is important. First, because that correlation shows what partakes in linguistic emotion judgments – e.g., as I find, many factors that also influence the recognition of physical emotions (Niedenthal, Halberstadt, Margolin, & Innes-Ker, 2000, i.a.). Second, because the entanglement between emotions and other information (about text and readers) makes low IAA lose its connotation of "unreliability", resolving (some) incompatible judgments as based on pre-annotation differences. Lastly, because that entanglement can determine how we train and evaluate classifiers: striving for coherent (ideally identical) annotations might be a missed opportunity to automatize the subjective core of our emotion abilities.

3 Zooming In: Emotions Inside Linguistic Layers

If asked what dimension of language allows to infer emotion, early computational research (Strapparava & Mihalcea, 2008, i.a.) would reply: semantics. Many words have a prototypical affective connotation, and that is why, e.g., "win" evokes joy, while "darkness" fear. The correspondence between emotions and meaning is an idea that keeps taking hold of the field, but it is just a commonsense assumption. For that matter, also the alternative equivalence between emotions and style remains unverified. The goal of the second strand of research in my thesis is to learn if either of these potentially competing views, emotions as linguistic styles (Chapter 5) or emotions as meanings (Chapter 6), is valid.

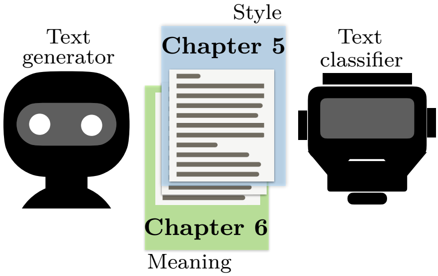

Figure 2: High-level overview of the basic components in Chapter 5 and 6.

I abandon the discussion of the relationship between people and emotions, narrowing the spotlight on data (cf. Figure 3) and leveraging computational systems (text generators and classifiers in Chapter 5, only classifiers in Chapter 6) for tasks that would require too intensive human labour. The results presented in both chapters have an applicative value for NLP. By identifying the linguistic level at which emotions arise, I determine whether computational tasks can model emotions as parts of style or word meanings. My contribution, however, is mainly theoretical: to explore the information needed to imbue and capture emotions in texts is to explore how emotions, which are primarily cognitive phenomena, turn into linguistic ones.

3.1 Chapter 5

To initiate a discussion of linguistic strata, Chapter 5 lays eyes on the "outer skin" of texts, namely style. Intuitively, style is an envelope that carries the text's content. It is formed by all sorts of syntactic and lexical choices used to communicate a meaning, and that could change without hampering its transmission (Bell, 1984). As a key component of language, style is played with every time humans reformulate the same idea in different ways. For instance, in some circumstances we prefer formal expressions to say things that could otherwise be phrased in jargon. This versatility applies to many aspects of language besides formality, but (so the current chapter observes) not to emotions.

I reach this conclusion proceeding from the premise that style is independent of meaning, and is treated as such by *style transfer* systems in NLP (Jin, Jin, Hu, Vechtomova, & Mihalcea, 2022). The task of style transfer requires automatic text generators to paraphrase an input while stirring its style towards a target attribute (e.g., given a formal expression, the goal is to translate it in jargon). The question is if emotions stand the test of style transfer: given a text loaded with joy, is it possible to generate another that is semantically similar but displays another emotion, like anger? If so, then the identity between emotions and style would be proven.

My approach builds upon backtranslation, a general purpose paraphrasing strategy that maps texts into a different language and then back, striving to preserve their meaning but not their surface form (i.e., style). Crucially, neural systems generate many candidate backtranslations under the hood. The one they return to the user is that which ranks the best according to criteria important for translation. I broaden these criteria, as to have emotions considered in the selection of the best text: the under-the-hood paraphrases are re-ranked based on information provided by an emotion classifier; accordingly, the text that proves to maximize a desired target emotion is returned. For example, the joyful “I was thrilled to see them” might be translated in German, and then back into “It was exciting to see them” and “Seeing them was shocking”. The second candidate would be outputted when targeting anger, since the classifier sees that emotion in there.

Subsequent experiments promise success. Without the intervention of my re-ranking algorithm, backtranslating softens input emotions (input: joy, output: less joy). I thus exploit these changes to the style transfer advantage, i.e., to counterbalance the loss of input emotions with an increase in the target ones. Helped by my classifier-informed re-ranking, I obtain paraphrases that consistently display different emotions than the input, such as texts with 44% more (target) shame, or 37% more (target) sadness than the corresponding inputs expressing guilt.

At a qualitative look, though, results appear less rosy. The supposedly re-styled texts contain target emotions to a greater extent than the inputs, but they fail to acquire a wholly new emotional profile. “Seeing them was shocking” sounds angrier, as it were, than “It was exciting to see them”, but does it express anger? Emotions are transferred only when the paraphrases stray away from the initial meaning (e.g., “It was outrageous to see them”). Thus, as long as style and meaning are considered separate axes of language, where changes in one leave the other unaffected, emotions cannot reside in style, at least not *just* in style, and not if one ignores that each emotion could be a *separate* axis, with its own space of possible linguistic operations (e.g., joy might be easily turned into surprise, not into anger).

3.2 Chapter 6

This chapter lifts the blanket of style to reach meaning. Intuitive as it may be, the idea that emotions nest in this layer of language raises big issues. First and foremost, not all formalisms to represent meaning in NLP might fully capture emotions. For example, thinking that individual words are the basic emotion units (e.g., “win” denotes joy, in a dictionary-like manner) ignores that their interpretation relies on world knowledge and the context in which words occur (e.g., if the “win” is undesirable for some reasons, it likely denotes anger).

It turns out that emotions are part of semantics if considered in the U-semantic perspective (Fillmore, 1985) of FrameNet (Baker, Fillmore, & Lowe, 1998). As a source of lexical abstractions, FrameNet describes meaning with a combination of predicates (i.e., frames) and arguments reflecting the structural properties of events – e.g., the frame BEAT_OPPONENT evoked by the word "prevail" comprises the arguments of winner and loser, and presupposes a WIN_PRIZE situation. In this sense, all frames account for the interplay between words, their context, and their interpretation.

I find that there is an emotional side to hundreds of them. An analysis of the link between emotions (obtained with an emotion classifier) and frames (obtained with a frame identification tool) in ≈44M sentences from the Corpus of Contemporary American English (Davies, 2015), testifies the presence of an emotion for 204 frames (i.e., the two variables have a strong positive correlation, pointwise mutual information $\geq$0.23), among the 818 unique frames found in the corpus. Simply put, the emotional import of a text is latently carried by the frames that it evokes.

Not all frames, however, bring the same type of affective information, as they represent different *components* of emotions. A qualitative inspection shows that some emotional frames denote events (e.g., PROTEST), others event evaluations (e.g., RISKY_SITUATION), and yet others the effects that events can cause in humans (e.g., FACIAL_EXPRESSION). In this light, frame semantics not only captures emotions, but it grasps their core properties, namely, the factors that appraisal theories in psychology (e.g., Scherer, 2005) claim to elicit, underlie or manifest event-ensuing emotions in the physical world.

All in all, for computational researchers in emotion analysis, this chapter reveals how FrameNet suits the study of emotions, in an event-based theoretical framework that holds potential to improve automatic text interpretation and generation. For frame semantics, it advocates the need to consider emotions as an integral part of word meanings, one where the folk understanding of worldly experiences (seized by frames) meets the experts' understanding (found in theories and definitions) of affective experiences.

4 The Four Problems, Revised

My thesis opens a channel of communication between NLP studies and other disciplines based on emotions. I find that (1) emotions can be modeled in the guise of "untold things" which are extracted from text, though imperfectly, because (2) their interpretation rests on factors in and out of language, (3) with linguistic style playing only a marginal role, and (4) meaning (which can be fuzzy and subject to many potential reads) being the primary source of their emergence.

Hence, *where are emotions (generated and recognized) in text?* They are in a text semantics: in the semantics of events. By reviewing the discussion in retrospect, it appears clear that this finding gradually peeps out from each of the core chapters. Chapter 3 discusses emotions in factual descriptions, all with the same highly-structured, factual style (hardly could it contribute to their linguistic realization). Chapter 4 proves that differences in the emotions recognized in text are backed up by differences in interpretations of the described events (i.e., in people's background knowledge about the affective meaning of events). Chapter 5 corroborates that generating text with a given emotion is unfeasible if one only cares about the visible layer of language (looking at a deeper level is necessary). Lastly, Chapter 6 uses frame semantics to study meaning irrespective of the lexical units that instantiate it (i.e., to peel off style), declaring frames the basic units of emotion, which capture the link with events and event judgments that, according to psychology, lies at the very heart of this complex phenomenon.

Emotions in Today's NLP. One year after my thesis defense, large language models have unprecedented linguistic skills. As ChatGPT recognizes emotions with enormous accuracy (Elyoseph, Hadar-Shoval, Asraf, & Lvovsky, 2023) and generates them on demand (Koptyra, Ngo, Radliński, & Kocoń, 2023), it comes natural to ask if the dialogue between technologies and theories that I advised is still relevant. I believe it is. In fact, my cross-disciplinary approach now has even grater potential to push research forward: NLP tasks can still inform theories, thanks to systems that have learned their emotion abilities from gigantic amounts of verbal data, and that give the chance to search the "rules" of emotions, by observing how they use them in text. Vice versa, theories (e.g., from psychology) can keep enhancing computational models, especially since these models interact with the general population today. Emotions serve to convey meanings and to carry out natural sounding conversations. Thus, it is still important to tackle the issue of what it means for a system to be emotionally proficient, considering that such a proficiency deals with a fundamental subjectivity (e.g., of event evaluations) that no speaker, human or artificial, can escape.

References

Bach, D. R., & Dayan, P. (2017). Algorithms for survival: a comparative perspective on emotions. *Nature Reviews Neuroscience*, *18*(5), 311–319.

Baker, C. F., Fillmore, C. J., & Lowe, J. B. (1998). The Berkeley FrameNet project. In *Coling-acl '98: Proceedings of the conference* (p. 86-90). Montreal, Canada.

Bell, A. (1984). Language style as audience design. *Language in society*, *13*(2), 145–204.

Brosch, T., Scherer, K. R., Grandjean, D., & Sander, D. (2013). The impact of emotion on perception, attention, memory, and decision-making. *Swiss medical weekly*, *143*(w13786).

Davies, M. (2015). *Corpus of Contemporary American English (COCA).* Harvard Dataverse. Retrieved from `https://doi.org/10.7910/DVN/AMUDUW` doi: 10.7910/DVN/AMUDUW

Einstein, A. (1915). *Die feldgleichungen der gravitation.*

Elyoseph, Z., Hadar-Shoval, D., Asraf, K., & Lvovsky, M. (2023). Chatgpt outperforms humans in emotional awareness evaluations. *Frontiers in Psychology*, *14*, 1199058.

Fillmore, C. J. (1985). Frames and the semantics of understanding. *Quaderni di semantica*, *6*(2), 222–254.

Hawking, S. (1999). *A brief history of relativity: What is it? how does it work? why does it change everything? an easy primer by the world's most famous living physicist.*

Jin, D., Jin, Z., Hu, Z., Vechtomova, O., & Mihalcea, R. (2022, March). Deep learning for text style transfer: A survey. *Computational Linguistics*, *48*(1), 155–205. Retrieved from https://aclanthology.org/2022.cl-1.6 doi: 10.1162/coli_a_00426

Kleinsmith, A., & Bianchi-Berthouze, N. (2012). Affective body expression perception and recognition: A survey. *IEEE Transactions on Affective Computing*, *4*(1), 15–33.

Koptyra, B., Ngo, A., Radliński, Ł., & Kocoń, J. (2023). Clarin-emo: Training emotion recognition models using human annotation and chatgpt. In *International conference on computational science* (pp. 365–379).

Niedenthal, P. M., Halberstadt, J. B., Margolin, J., & Innes-Ker, Å. H. (2000). Emotional state and the detection of change in facial expression of emotion. *European journal of social psychology*, *30*(2), 211–222.

Nussbaum, M. (2004). *Emotions as judgments of value and importance.* Oxford University Press.

Roseman, I. J. (1984). Cognitive determinants of emotion: A structural theory. *Review of personality & social psychology*(5), 11–36.

Scherer, K. R. (2005). What are emotions? And how can they be measured? *Social Science Information*, *44*(4), 695–729. Retrieved from http://journals.sagepub.com/doi/10.1177/0539018405058216 doi: 10.1177/0539018405058216

Strapparava, C., & Mihalcea, R. (2008). Learning to identify emotions in text. In *Proceedings of the 2008 acm symposium on applied computing* (pp. 1556–1560). New York, NY, USA: ACM. Retrieved from http://doi.acm.org/10.1145/1363686.1364052 doi: 10.1145/1363686.1364052

Tooby, J., & Cosmides, L. (1990). The past explains the present: Emotional adaptations and the structure of ancestral environments. *Ethology and sociobiology*, *11*(4-5), 375–424.

Venter, J. C., Adams, M. D., Myers, E. W., Li, P. W., Mural, R. J., Sutton, G. G., ... al. et (2001). The sequence of the human genome. *Science*, *291*(5507), 1304–1351.

Wierzbicka, A. (1992). Talking about emotions: Semantics, culture, and cognition. *Cognition & Emotion*, *6*(3-4), 285–319.

Correspondence

Enrica Troiano

HK3Lab
Rovereto, Italy
enrica.troiano@hk3lab.ai

Thesis Information

Doctoral thesis defended on February 16, 2023, at the Institute for Natural Language Processing (IMS), University of Stuttgart. Supervised by Prof. Roman Klinger and Prof. Sebastian Padó. Full text available in the University of Stuttgart Electronic Library: https://elib.uni-stuttgart.de/bitstream/11682/13671/1/Troiano.pdf.