

**Machine learning/natural language processing and emotions:
Some ethical considerations****Václav Ježek¹****Abstract**

This article examines how machine learning and natural language processing impact human emotional development. The main thesis is that a reciprocal relationship exists by programming machines with specific causal categories, we shape how they interact with us, and in turn, these machines influence human environments and emotional experiences. This mutual shaping can encourage reductionist views by both humans and machines. To address ethical concerns arising from this interaction, we emphasize the need to recognize the context and relativity of emotion recognition technology. Emotions are inherently indeterminate and cannot be fully captured by any single framework.

Keywords: emotion, machine learning, natural language processing

Introduction

In this article, we focus on the ethical and technical challenges at the intersection of machine learning, language processing, and emotion. Our main argument is that if designers of these systems do not thoroughly account for the complexity of human emotions, both technology development and people will be affected. Machine learning systems do more than process human information—they also actively shape human behavior and perceptions via feedback loops. Initially, machine learning treated human traits linearly, but modern approaches aim to be more holistic. We assert this relationship is fundamentally reciprocal and continually evolving as humans define machine processes, machines also influence human emotional experiences, with both shaping each other's development.

Machine learning and language processing

Two principal factors warrant consideration when analyzing online texts. First, searches may be intentional, where users seek specific information, or unintentional, when data present themselves without deliberate effort. Second, any search inherently involves interaction, leading to subsequent change. Over time, the distinction between the searcher and the object of the search becomes increasingly ambiguous.

The human mind uses association, a core element in both human thinking and machine learning. We understand and categorize things through association and context. This forms our emotional framework. Whether these things are real or imagined matters less. Their perceived impact on emotion is similar to how machines use associative data. It reflects how machines process associative data. For instance, Charles Darwin could not imagine emotion without its 'physical' expression and vice versa. He identified an association between a mentally ill woman and her 'disorderly' appearance (Darwin, 1897, p. 296), highlighting the link between mental states and physical manifestations.

In both machine learning and natural language processing—and in Darwin's view—there must be a connection between emotions and their physical manifestations. Without this, computation or analysis is not possible. One way to establish this connection is by devising emotional lexicons with defined emotional meanings. Language models can then capture common language patterns, such as those identified in BERT. Use definite emotional content for an emotional term. But this is not always accurate. Machine Learning methods, such as the bag-of-words representation, help machines analyze texts by assuming term frequency

¹ University of Prešov (Slovak Republic); email: vaclavjezek111@gmail.com; ORCID 0000-0001-9521-7963

indicating possible emotional content. At a basic level, things are more complex. Even terms such as “love” are not always “emotionally” charged. On the human level, terms or expressions acquire emotional content through complex associations or relativizations. From an ethical perspective, some sites use various ways to associate terms to provoke an emotional response. For instance, one website about Indian religious systems uses the title: “The Science of Self-Realization”.² Combining “Science” with “Self-Realization” creates more emotional impact than either word alone.

In this regard, the research of Hughes et al., who explored books from past centuries (using the Project Gutenberg Digital Library), has shown that “content-free words” acquire meaning through context, including conjunctions, articles, and pronouns. The combinations of “content-free” words and words with specific content display a vector of similarity (Hughes et al., 2012).

Other research methods can be used to assess other aspects in similar analyses. For example, the research of Twenge et al., using Google’s Ngram database, concluded that there was an increase in “individualistic words” at the expense of “communal words” in American books in the last decades of the twentieth century (Twenge et al., 2012). Later, Acerbi et al. used Google’s Ngram database to analyze the mood of books over the past century. They found periods of happiness and sadness. More importantly, the use of mood words decreased. This matches a decline in emotional words in literature. Acerbi concludes the words may have changed, not simply decreased (Acerbi et al., 2013).

The key distinction between machines and humans is that machines assume a strict correspondence between form and content, while humans engage more dynamically with meaning. In textual analysis, the relationship between meaning and form in machine analysis remains fixed by initial criteria. As a result, machine-generated emotional data reflects the parameters established at the outset, whereas humans interpret and adapt these relationships more flexibly. They are more attuned to human emotional “unexpectedness” and unpredictability. It uses various models, such as the Dimensional emotional model. This model represents emotions in three parameters: valence, arousal, and power (Denecke et al., 2024, p. 2). In the categorical model, emotions are defined discretely. Examples are anger, sadness, and fear (Denecke et al., 2024, p. 2). Emotional analysis is now a more complex method than popular Sentiment analysis. It does not focus solely on positivity, negativity, or neutrality, as Sentiment analysis does (Munezero et al., 2014). The growing appreciation for “emotional intelligence” makes things more complex. It can work at the level of complex neural networks (Serrat, 2017). These include variables that are difficult for machines to compute (Allahyari, 2020).

As these examples show, emotionality is not an inherent quality of seemingly emotional terms. Drawing from the theory of Special Relativity, emotions—like physical phenomena—are context-dependent and subject to change. Technically and ethically, inaccurate machine learning can affect human emotional development. Therefore, algorithms should prioritize contextual accuracy to prevent unintended negative effects since any functions can be limited to the variables and forms they are constructed with. On a basic level, for example, $2+2 = \text{four}$ is a valid rational assumption. However, it is limited in time and space to numbers 2 and 4, as well as its context, temporal goal, expression, and form. The equation is a valid node that can link to other valid, enclosed nodes. Another example is $3+3 = \text{six}$. This is a separate node, not necessarily connected to the first equation, if treated as another point in a node framework. In human terms, my grief is something I feel and therefore exists. But it is not *necessarily* linked to someone else’s happiness, which is just as real for them. Both emotions are valid, but they are not automatically related.

This is a fundamental aspect of human psychology. Human responses to emotional stimuli are subject to change; repeated exposure to emotionally provocative images, for example, may

² <https://vedabase.io/en/library/ssr/>

cultivate emotional resilience or alter emotional reactivity. Recent advancements in models, such as Large Language Models with anthropomorphic properties or “Emotional Alignment” (Huang et al., 2024), attempt to address heterogeneity in emotional responses; nonetheless, these models continue to encounter challenges in accurately representing the dynamic nature of human emotions, as such variables are difficult to algorithmically model.

A more serious challenge, however, is not only the inability of machine learning to fully capture the range of emotions. The potential effects of large volumes of emotional information on individuals have received little research attention. Only recently have studies begun to address these implications in professional environments, indicating the need to account for them in future work. Effective computing attempts to address potential flaws in emotional analysis by using broader datasets, incorporating psychological knowledge, and not only textual material but also video, audio, and related media. Here, multimodal sentiment analysis (MSA) can combine text and images, texts and audio, or images and audio, with combinations of all (Pandey & Vishwakarma, 2024). Recognizing emotion remains challenging, so more analytical methods are being developed. For sentiment analysis, aspect-level (feature-level) methods add more features for deeper review. Some approaches are used in mental disorder research or when working with convolutional neural networks (Monisha et al., 2023). However, preset mathematical models cannot always account for the role of chance in data.

There are projects that aim to circumvent the limits of predetermined categorization, such as the Automatic Sentiment Analysis in the Wild (SEWA-EC H2020), which seeks to offer a more complex analysis of human behavior by adopting specific machine learning and other techniques (www.sewaproject.eu).

Emotions and machine learning

As we have emphasized, the two-way process between the human being and the Machine entails also “emotional” consequences for the human being. The human being forms a language that Artificial intelligence can understand. This language is usually of a “rational” kind. However, Artificial intelligence does not necessarily adapt itself to human emotionality.

Some indication of this can be gained by, at least for now, growing clusters of sporadic inferences. Here, for example, Almashraee notes: “The majority of reviewers, e.g., Facebook users, bloggers, or Twitterers, express their emotional opinions using rational words rather than emotions” (Almashraee et al., 2016). These authors have developed a formula for computing rational statements as emotional ones and vice versa, $PMI(o-w,e-w) = \log$, where $P(o-w)$ is the probability of the opinion-word, $P(e-w)$ is the probability of the emotion, and $P(o-w,e-w)$ is the probability of both opinion-word and the emotion to co-occur (Almashraee et al., 2016).

Here, we may add that, generally, probability theory hinges on the N infinite number of possibilities. In emotion and its variations and combinations, this can also be a valid consideration.

The computation of emotions cannot automatically mirror all the possibilities. Deilamsalehy observes that “in some cases, NLP models are so good that we can confuse consciousness with intelligence. One thing to keep in mind is that these models have seen the entire corpus of the internet. However, data does not automatically demonstrate “happiness” (Dooley, 2025).

In assessing the role of emotions in machine learning, we encounter a fundamental problem that has long been a stumbling block in general psychology: what emotions are. For our purposes, suffice it to say that while an individual can experience emotion, we assume that emotions entail two sides. A human being cannot experience emotions in an autogenic way for long.

Defining emotions is notoriously difficult. Ekman identifies some emotions such as “amusement, anger, contempt, contentment, disgust, embarrassment, excitement, fear, guilt, pride in achievement, relief, sadness/distress, satisfaction, sensory pleasure, and shame”.

However, his scheme depends on the categories he chose (Ekman, 1999, p. 56). His models have been criticized as limited, due to his view of human behavior (Coppini et al., 2023). Current research has developed broader categories, such as the Dimensional theory, which, for example, assesses human emotional interaction with landscapes (Bakker et al., 2014).

The possibilities of emotions are related to the eye of the beholder. Emotions can include touch, blood flow, blood pressure, facial expressions, infrared aspects of the body, body movement (proprioceptive data), gait, haptic data (force sensors), speech expression, nervous reactions, sexual arousal, and others. The most important consideration for our purposes is that emotions are dynamic, not static. Thus, they change and evolve over time, and each person can experience them with varying intensity and character. Obviously, this dynamism creates problems for Machine Learning. For example, each text can be emotionally experienced differently by different people and change over time. Further, there is no automatic correspondence between facial expressions, other expressions, and internal emotionality.

Emotions can be characterized more reductively in relation to basic brain functions. Barrett observes:

A brain implements an internal model of the world with concepts because it is metabolically efficient to do so... So the brain constructs an online concept of happiness, not in absolute terms, but with reference to a particular goal in the situation (to be with friends, to enjoy a meal, to accomplish a task), all in the service of allostasis. This implies that “happiness” has a specific meaning, but its specific meaning changes from one instance to the next (Barrett, 2017, p. 11).

Barrett’s suggested contextualization of brain function again confirms the complexities of emotional variables and our stress on the relativity and contextual variability of emotions.

One of the fundamental conclusions that stems from Machine Learning is that emotion and reason are partners and are related to each other. For example, if I am presented with a nice book on the World Wide Web, I can decide to buy it based on my rational assessment that it is, in fact, good, and on my emotional disposition, which leads me to feel a consequent emotional positivity toward this book. Or I can feel emotion when I see a war image, which leads me to the rational conclusion that war is bad. The symbiosis of reason and emotion was already important in ancient Greek thought; suffice it to mention Thucydides (Visvardi, 2012).

Now this traditional symbiosis is being rediscovered through Appraisal theory, which sees a relationship between cognitive functions and emotion. Reason can stimulate emotion, and emotion can stimulate reason. My current emotion may result from pure reason, or it may be stimulated by an emotional predisposition. Here, emotions can paradoxically have all the features of rationality and vice versa (Du Toit, 2014). There is something called “reasonable emotion”. The well-known neuroscientist António Damásio sees a link between functional rationality and emotions. Without emotions there essentially cannot be a rational decision (Damásio, 2020).

All these considerations are important for new possibilities in calculating and understanding the relationships among image, perception, expression, and so on. In the nineteenth century, William James proposed a new view of emotions, suggesting that the relationship between cognition/rationality and emotion could be bidirectional, in which an experienced emotion can have rational consequences. He writes:

Our natural way of thinking about these standard emotions is that the mental perception of some fact excites the mental affection called the emotion, and that this latter state of mind gives rise to the bodily expression. My thesis, on the contrary, is that *the bodily changes follow the PERCEPTION of the exciting fact directly, and that our feeling of the same changes as they occur IS the emotion* (James, 1884, p. 190).

Here, the basic criticism of modern technology that it creates “non-reality” is not the problem. This is because, in human psychology, “unreal”/“real” constructions existed before. For example, I can fall in love with an unreal mental construction (falling in love with a female I met but no longer see). For technological assessment, all my real or unreal concepts that I form are factors to be computed and have no qualitative differential in terms of Machine approaches. However, “true reality” is related to consciousness, which is what really differentiates us from machines at this point.

The main ethical problem that emerges in our context is that a distortion between emotion and rational assessment can create a condition of emotional numbness, similar to that of psychiatric patients with various conditions. This can lead to an inability to identify one’s emotions. The Toronto Alexithymia Scale identifies the difficulties people have in identifying their emotions.

If there is any positivity in Machine Learning regarding emotions, it is confirmation that the best way forward is to see emotions as essentially located between two sides or two people. This allows for appreciating the complexities of the possibilities involved.

Broader ethical concerns

Unsurprisingly, ethics in Artificial Intelligence is becoming increasingly important in academic literature (Ciuverca, 2024). In terms of broader emotional or sentiment analysis we must ask the question of what its underlining goals are. As we have seen, problems arise when machine learning ceases to be an observer and becomes a producer of emotions and other human characteristics.

Further, as we have implied contextuality and relativization in machine learning technology does not enable one single ethical approach. It is a false premise to assume that there is some kind of infallible ethics or a general ethics for all circumstances. This is a fact recognized by many authors. One of the solutions proposed by Mohammed is to create multiple ethical sheets. He observes: “Multiple ethics sheets can be created (by different teams and approaches) to reflect multiple perspectives, viewpoints, and what is more important to different groups of people. We should be wary of the world with single authoritative ethics sheets per task and no dissenting voices” (Mohammad, 2022, p. 241).

Here, ethics to some extent resemble the natural characteristics of Machine Learning itself, a rapidly and dynamically developing field. As Townsend et al. observe: “As argued by the Association of Internet Researchers (2012), no set of internet research guidelines can be static, because technologies and the way that technologies are used are constantly changing. Consequently, conversations on ethical standards in social media research need to be dynamic too” (Townsend & Wallace, 2018, p. 196).

Here, for example, the so-called Digital Ethics Canvas

is a visual tool for assessing ethical risks from six ethical perspectives specific to the digital domain: beneficence, non-maleficence, privacy, fairness, sustainability, and empowerment... Depending on whether the canvas is used in the design or development phase or at runtime, mitigation strategies can be artifact- or context-oriented. That is, the strategies could include changing the technological artifact (e.g., avoid collecting personal data that is not needed to reduce privacy risk) or changing the usage context (e.g., ask users to provide a nickname rather than their actual name) (Denecke & Gabarron, 2024, p. 4).

Apart from general ethical concerns, there is the more concrete problem of machine flaws, which result in invalid analyses and assessments. We may argue that such flaws can disrupt the development of human emotions. Flaws include Bias, which can be constructed, inadvertent, or the result of chance. This relates to bias made through “training data collection, pre-

processing, or algorithm” (Karoo & Chitte, 2023, p. 2992).

The problems with individual emotional assessment are addressed by machine learning and by aggregate statistical methods, which we may term pattern analysis. This can be understood by applying a statistical spiral to, for example, a hundred emails. Through this type of analysis, a “pattern” of emotionality can emerge. However, this method has its own flaws, as it reduces the uniqueness of each person to more general considerations and patterns, limiting the freedom of such approaches to account for personal variations and to utilize them for further assessment. Here, logistic regression can serve as a more discriminative classifier than, for example, the naïve Bayes classifier, which is generative (Jurafsky & Martin, 2008).

Of course, a basic ethical concern is privacy. The misuse of personal data and so on. However, as we have seen, the main problem here is that, even if personal data is not misused for other purposes, it remains the fundamental building block of all media interaction. It is almost impossible to avoid being emotionally or otherwise labeled through data processing. All our searches and choices are automatically stored and further structure the information that we receive. Here, the problem is not with anonymization but with causality and predetermination. Paradoxically, we form the machine which then forms us. Further, personal data collected can be misused or used in the future in ways that are now inconceivable. Here, medical assessments can be problematic (Saraff et al., 2020). Chatbots can be used to treat psychological medical conditions. However, data here can be flawed, as can the conclusions (Denecke & Gabarron, 2024, p. 7).

Chatbots in emotional exchanges are fully fledged emotional partners (of course, with their technical limits), and their benefit or not depends on the level of the person. Emotional attachment to the Chatbot on a personal level presents a challenge. In 2023, a Belgian man committed suicide by speaking intensively with Chatbot Eliza. Here, as we have indicated, there is no point in assessing whether emotions were real or imagined; the point is that all these have consequences. As we have further implied, the problem of personal information being lost or misused is not as great a concern as the problem of how the machine interacts with the individual. Personal information about someone else can just as negatively impact the machine’s interaction with me as it can with the person in question.

As we have further shown, information related to emotional disposition is dynamic. For example, I may write on social media that I do not experience love from anyone. The machine incorporates this information. In the meantime, I find someone who loves me. However, for all purposes, my data could still show a state of “unlove”, which paradoxically could entail problems for me or my relationships. Anonymization of personal data can alleviate the situation, but not necessarily. Further, all automatic emotion recognition systems operate with specific goals in mind. It is not a philanthropic activity. There is a purpose and goal.

Conclusions

Anyone traveling on public transport can confirm the impact of media on people who are often fully immersed in their mobile devices. On a basic level, we may argue that nothing serious is happening. However, these people who are commuting in public are redirecting their emotional investment from their immediate physical surroundings (the people around them) to something that is definitely of a different qualitative emotional meaning and value (the mobile phone). This situation is remarkably like the ethical considerations related to emotions in Machine learning generally, since at first glance, nothing challenging is occurring.

As we have indicated, one of the important missing assumptions in Machine Learning and Artificial Intelligence is the fact that there is a two-way process between the Machine and the human being. It is not only the human being who programs the machine through algorithms, programs, and so on; this process then reverses itself, and the machine paradoxically expects the human to behave according to the way human beings have presented their “human behavior”

to the machine.

In terms of emotions, I can program the machine to offer emotional responses or more physical responses. However, I will always receive what I expect, not what is necessarily objective or beneficial to my emotional development. An obvious example would be the increasing problem of marital crisis, which relates to social media. I can simply choose a partner on social media solely for my own egoistic emotional needs. This is not only illusory but destructive. This leads us to another conclusion: that any emotional programming of the machines is problematic not only because of the possible flaws in such algorithms, but also because of the inherent, unexplored value, processes, and meaning of emotions among human beings themselves. The problem of individualism or communion is another issue. The machine is programmed to give what you expect, but not necessarily what you need.

As we have seen, causality and pre-determinative structures, coupled with machine flaws, can create an artificial environment of emotionality in machine learning contexts. The nature of human emotions, their expression, and prevalence do not allow for easy categories. While contemporary systems allow for a more “natural construction” of emotional recognition, problems remain (Denecke, 2023). Here, the axiom is that what you seek is what you get. And further, what you found is what will form you in the future.

Another problem is the dissecting of human characteristics. Here, it would be more useful to operate on the level of the Person. The Person as a category would allow for a greater contextuality than simply assessing individual human components. Thus, just because we see somebody behaving emotionally in one way or another does not usually tell us much about the person. However, if I know the person, I can contextualize their emotions more objectively.

Further, as we have seen, there can be an algorithmic problem in any emotional analysis. If emotion is the result of interaction, and I analyze that interaction, my analysis will always be “limited” to that interaction. This is because in any given interaction, only a limited expression of emotions is available or possible. By virtue of the “sides” involved. For example, what is valid in my emotional interaction with my neighbor may not be valid in my interaction with my wife, and so on. However, any individual interaction can serve as a basis for general node analysis. Unless, however, this unique node is properly contextualized, it can become a closed node.

Emotions in the interaction then continue to form another node in a similar analysis. The same process continues until I reach a limited understanding of the possibilities of emotional expression. I started with an open, infinite number of possibilities and ended up with a limited number. This has consequences because it begins to limit the freedom of emotional expression, understood as an indeterminate mathematical construction. As we have seen, this kind of unpredictability must be understood by the machine to avoid pitfalls. See figure 1.

Last but not least, emotional recognition can be dangerous in totalitarian frameworks. Here we may quote an innocent sentence from one article: “Through the analysis of emotional tendency in online public opinion, governments and enterprises can stabilize people’s emotions more effectively and maintain social stability” (Xiaojie Li, 2021, p. 1).

Acknowledgement

The paper was the result and supported by the project MEYS, ERC CZ program, grant No. LL2325. Grant of the Mathematics Physics faculty of the Charles University (year 2024).

References

- ALLAHYARI, E. (2020): Application of artificial neural network in predicting EI. In: *BioMedicine*, 10(3), pp. 18–24. doi: 10.37796/2211-8039.1029
- ACERBI, A., LAMPOS, V., GARNETT, P. & BENTLEY, R. A. (2013): The expressions of emotions in 20th century books. In: *PLoS ONE*, 8(3): e59030, pp. 1–6. doi:

10.1371/journal.pone.0059030

ALMASHRAEE, M., MONETT, D. & PASCHKE A. (2016): Emotion level sentiment analysis: The Affective Opinion Evaluation”, Conference paper, Semantic Sentiment Analysis, ESWC, 2018, pp. 1–12. [online] [Retrieved December, 2024] Available at: https://ceur-ws.org/Vol-1613/paper_1.pdf

BAKKER, I., de BOON, J., THEO, J. M., VOORDT, Van der & VINK, P. (2014): Pleasure, Arousal, Dominance: Mehrabian and Russel revisited. In: *Current Psychology*, 33(3) pp. 405–421, doi: 10.1007/s12144-014-9219-4

BARRETT, L. F. (2018): The theory of constructed emotion: an active inference account of interoception and categorization. In: *Social Cognitive, and Affective Neuroscience*, 12(1), pp. 1–23, doi: 10.1093/scan/nsw154.

COPINI, S., LUCIFORA, C., VICARIO, C. M. & GANGEMI, A. (2023): Experiments on real-life emotions, challenge Ekman’s model. In: *Scientific reports*, 13, 9511, pp. 1–12. [online] [Retrieved October 11, 2024] Available at: <https://doi.org/10.1038/s41598-023-36201-5>

DAMÁSIO, A. (2020): Interview. Without emotions you always make the wrong decisions, In: *Philonomist*. [online] [Retrieved December 12, 2024] Available at: <https://www.philonomist.com/en/interview/without-emotions-you-always-make-wrong-decisions-2024>.

DENECKE, K. (2023): *Sentiment analysis in the medical domain*. Cham: Springer Nature.

DENECKE, K. & GABARRON, E. (2024): The ethical aspects of integrating sentiment and emotion analysis in chatbots for depression intervention. In: *Frontiers in Psychiatry*, 15, pp. 1–11. doi: 10.3389/fpsy.2024.1462083.

DARWIN, C. (1897): *The expression of the emotions in man and animals*. New York: D. Appleton and Company.

DOOLEY, J. (2025): How NLP Is changing everything about modern search engines. Coveo.com. [online] [Retrieved November 10, 2024], Available at: <https://www.coveo.com/blog/nlp-search-engine/>.

CIUVERCA, A. C. D. (2024): Sentiment analysis of research on AI Ethics: A web based study. In: *Proceedings of the 7 th International Conference on Economics and Social Sciences*. Bucharest: Bucharest University of Economic Studies, pp. 672–681. doi: 10.24818/ICESS/2024/060

EKMAN, P. (1999): Basic emotions. In: T. Dalgleish & M. J. Power (eds.): *Handbook of cognition and emotion*. New York: John Wiley & Sons, pp. 45–60. Also available at: <https://doi.org/10.1002/0470013494.ch3>

FLEMMING, J. E. (2012): *Passions and emotions*. New York: New York University Press.

HUANG, Jen-tse, MAN, Ho A. LAM, JOHN, E. Li, REN, Shujie, WANG W., JIAO W., ZHAOPENG, Tu & LYU, R. M. (2023): Emotionally numb or empathetic? Evaluating how LLMS feel using emotion bench. In: *ArXiv*. abs/2308.03656, [online] [Retrieved November 8, 2024], doi:10.48550/arXiv.2308.03656 Available at: <https://api.semanticscholar.org/CorpusID:260682960>

HUGHES, J. M., FOTI, N. J., KRAKAUER, D. C. & ROCKMORE, D. N. (2012): Quantitative patterns of stylistic influence in the evolution of literature. In: *Proceedings of the National Academy USA*, 109(20), pp. 7682–7686.

JAMES, W. (1884): What is an emotion? In: *Mind*, 9(34), pp. 188–205.

JUNGHANNS, G. & KERSTEN, N. (2020): Informationsüberflutung am Arbeitsplatz: Gesundheitliche Konsequenzen. In: *Zentralblatt für Arbeitsmedizin, Arbeitsschutz und Ergonomie*, 70(1), pp. 8–17. doi: 10.1007/s40664-019-00370-w

JURAFSKY, D. & MARTIN, J. H. (2008): *Speech and language processing*. Hoboken, NJ: Prentice Hall.

KAROO, K. & CHITTE, V. (2023): Ethical considerations in sentiment analysis: Navigating the complex landscape”, In: *International research journal, of Modernization in Engineering*,

Technology and Science, pp. 1234–1240. doi: <https://www.doi.org/10.56726/IRJMETS46811>).

MOHAMMAD, S. M. (2022): Ethics sheet for automatic emotion, recognition and sentiment analysis. In: *Computational linguistics*, 48(2), pp. 239–278.

MOORS, A. (2017): Appraisal Theory of Emotion, In: V. Zeigler-Hill & T. K. Shackelford (eds.): *Encyclopedia of Personality and Individual Differences*. Cham: Springer, pp. 1–9, doi:10.1007/978-3-319-28099-8_493-1

MONISHA, G. S., YOGASHREE, G. S., BAGHYALAKSMI, R. & HARITHA, P. (2023): Enhanced Automatic recognition of human emotions using machine learning techniques. In: *Procedia Computer Science*, 218, pp. 375–382. [online] [Retrieved October 10, 2024] Available at: <https://doi.org/10.1016/j.procs.2023.01.020>.

MUNEZERO, M. D., MONTER, C.S., SUTINEN, E. & PAJUNEN, J. (2014): Are they different? Affect, feeling, emotion, sentiment, and opinion detection in text. In: *IEEE Transactions on Affective Computing*, 5(2), pp. 101–111. doi: 10.1109/TAFFC.2014.231787.

PANDEY, A. & VISHWAKARMA, D. K. (2024): Progress, achievements, and challenges in multimodal sentiment analysis using deep learning: A survey. In: *Applied Soft Computing*, 152(2), 111206, [online] [Retrieved December 12, 2024] Available at: <https://doi.org/10.1016/j.asoc.2023.111206>.

SERRAT, O. (2017): *Knowledge solutions: Tools, methods, and approaches to drive organizational performance*. Cham: Springer Open. [online], [Retrieved December 5, 2024] Available at: doi:10.1007/978-981-10-0983-9_37.

SARAFF, S., TARABAN, R., RISHIPAL, R. B., KEDAS, S. & GUPTA, S. (2020): Application of Sentiment analysis in understanding human emotions and behaviour. In: *EAI Endorsed transactions, on Smart Cities*, 5(13), pp. 1–8. [online] [Retrieved December 2, 2024]. Available at: doi: 10.4108/eai.6-10-2020.166547.

TOWNSEND, L. & WALLACE, C. (2018): Social media research: A guide to ethics. The ethics of using social media data in research: A new framework. In: K. Woodfield (ed.): *The ethics of online research*. Leeds: Emerald publishing, pp. 189–207.

TOIT, Du (2014): Emotions and the affective turn: towards an integration of cognition and affect in real life experience. In: *HTS Toelgiese Studies/Theological Studies*, 70(1), 2692, pp. 1–9. Available at: <http://dx.doi.org/10.4102/hts.v70i1.2692>.

TWENGE, J., CAMPBELL, K. W. & GENTILE, B. (2012): Increases in individualistic words and phrases in American books, 1960–2008. In: *PLoS ONE* 7(7): e40181, pp. 1–5.

VISVARDI, E. (2012): Collective Emotion in Thucydides. In: *CHS Research Bulletin*, 1(1). [online] [Retrieved December 1, 2024]. Available at: http://nrs.harvard.edu/urn-3:hlnc.essay:VisvardiE.Collective_Emotion_in_Thucydides.2012<https://vimeo.com/57934923>

XIAOJIE, Li (2021): Research on the emotional tendency of web texts based on long short-term memory network. In: *Journal of Intelligent Systems*, 30(1), pp. 988–997. [online], [Retrieved December 1, 2024] Available at: <https://doi.org/10.1515/jisys-2021-0088>

ZIZIOULAS, J. (2007): *Communion and otherness*. London: T and T Clark.