

Bridging Emotion Role Labeling and Appraisal-based Emotion Analysis

Roman Klinger

Institut für Maschinelle Sprachverarbeitung

University of Stuttgart, Germany

roman.klinger@ims.uni-stuttgart.de

Abstract

The term emotion analysis in text subsumes various natural language processing tasks which have in common the goal to enable computers to understand emotions. Most popular is emotion classification in which one or multiple emotions are assigned to a predefined textual unit. While such setting is appropriate to identify the reader’s or author’s emotion, emotion role labeling adds the perspective of mentioned entities and extracts text spans that correspond to the emotion cause. The underlying emotion theories agree on one important point; that an emotion is caused by some internal or external event and comprises several subcomponents, including the subjective feeling and a cognitive evaluation. We therefore argue that emotions and events are related in two ways. (1) Emotions are events; and this perspective is the fundament in NLP for emotion role labeling. (2) Emotions are caused by events; a perspective that is made explicit with research how to incorporate psychological appraisal theories in NLP models to interpret events. These two research directions, role labeling and (event-focused) emotion classification, have by and large been tackled separately. We contributed to both directions with the projects SEAT (Structured Multi-Domain Emotion Analysis from Text) and CEAT (Computational Event Evaluation based on Appraisal Theories for Emotion Analysis), both funded by the German Research Foundation. In this paper, we consolidate the findings and point out open research questions.

1 Introduction

“Communication is an exchange of facts, ideas, opinions, or emotions by two or more persons. The exchange is successful only when mutual understanding results.” (Newman et al., 1967, p. 219)

The development of computational models in natural language processing aims at supporting communication between computers and humans; with

language understanding research focusing on enabling the computer to comprehend the meaning of text. Sometimes, understanding facts is sufficient, for instance when scientific text is analyzed to automatically augment a database (Li et al., 2016; Trouillon et al., 2017). Factual statements can also comprise explicit reports of emotions or sentiments, such as “They were sad.”, and in such cases, the analysis of subjective language blends with information extraction (Wiebe et al., 2004).

Emotion analysis, however, goes beyond such analysis of propositional statements. To better understand what emotion analysis models are expected to do, it is worth reviewing emotion theories in psychology. There are many of them, with varying purposes and approaches, but most of them, if not all, agree on the aspect that *emotions are caused by some event* and come with a change of various subsystems, such as a change in motivation, a subjective perception, an expression, and bodily symptoms. Another component is the evaluation of the causing event, sometimes even considered to constitute the emotion (Scarantino, 2016).

The *emotion also corresponds to an event itself*, embedded in a context of other events, people, and objects. All components of such emotion events (cause, stances towards other involved people, opinions about objects) may be described along an explicit mention of an emotion name. Any subset of them may appear in text, and may or may not be sufficient to reliably assign an emotion representation to the text author, a mentioned entity, or to a reader (Casel et al., 2021; Cortal et al., 2023).

This complexity has led to a set of various emotion analysis tasks in NLP, which we exemplify in an integrated manner in Figure 1. The most popular task is emotion prediction, either representing the writer’s or the reader’s emotion as a category, as valence/arousal values, or as appraisal vector (at the bottom of Figure 1, we will describe the underlying psychological theories in Section 2.1). Adding the

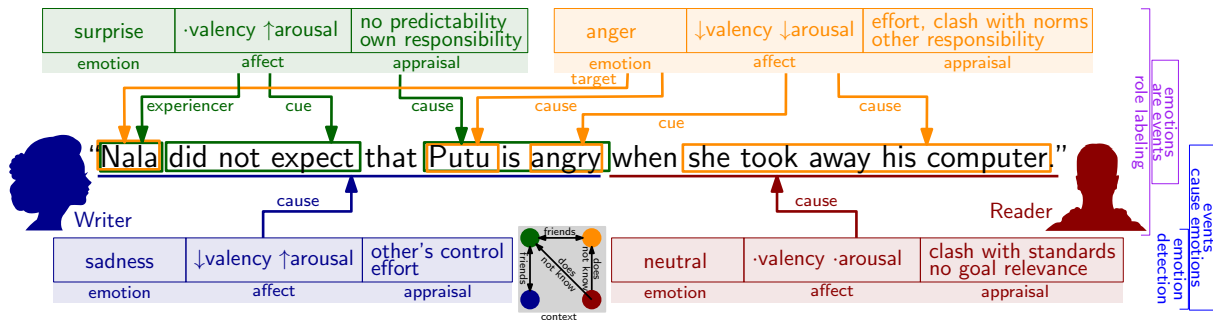


Figure 1: Integrated Visualization of Research Tasks in Emotion Analysis

task of cause detection bridges to the role labeling setup (visualized in more completeness at the top). Here, the emotion event is represented by the token span that represents the emotion experiencer, the cue, and the cause. *Emotion prediction focuses on understanding from text how events cause emotions, while role labeling focuses on understanding how emotions are represented as events themselves.*

We now introduce the background to emotion analysis, including psychological theories, related tasks, and use cases (§2). Based on that, we consolidate recent research on the interpretation of events to infer an emotion and on emotion role labeling (§3.1–3.2). We then point out existing efforts on bridging both fields (§3.3) and, based on this, develop a list of open research questions (§4).

2 Related Work

2.1 Emotion Theories in Psychology

Before we can discuss emotion analysis, we need to introduce what an emotion is. The term typically refers to some feeling, some sensation, that is defined following various perspectives. Scarantino (2016) provides an overview of various emotion theories and differentiates between a *motivation tradition*, a *feeling tradition*, and an *evaluative tradition*.

2.1.1 Categorical Models of Basic Emotions

The motivation tradition includes theories that are popular in NLP such as the basic emotions proposed by Ekman (1992) and Plutchik (2001). They differ in how they define what makes an emotion basic: Ekman proposes a list of properties, including an automatic appraisal, quick onset, brief duration, and distinctive universal signals. According to him, non-basic emotions do not exist but are rather emotional plots, moods, or personality traits. Plutchik defines basic emotions based on their function, and

non basic-emotions are gradations and mixtures. The set of basic emotions according to Ekman is commonly understood to correspond to joy, anger, disgust, fear, sadness, and surprise. However, in fact, the set is larger and there are even emotions for which it is not yet known if they could be considered basic (e.g., relief, guilt, or love, Ekman and Cordaro, 2011). The basic emotions according to Plutchik include anticipation and trust in addition. In NLP, such theories mostly serve as a source for label sets for which some evidence exists that they should be distinguishable, also in textual analysis. A study that uses a comparably large set of emotions is Demszky et al. (2020), while many other resource creation and modeling attempts focus on subsets (Alm et al., 2005; Strapparava and Mihalcea, 2007; Schuff et al., 2017; Li et al., 2017; Mohammad, 2012, i.a.).

2.1.2 Dimensional Models of Affect

An alternative to representing emotions as categorical labels is to place them in a (continuous) vector space, in which the dimensions correspond to some other meaning. The most popular one is the valence/arousal space, in which emotions are situated according to their subjective perception of a level of activation (arousal) and how positive the experience is (valence). This concept stems from the feeling tradition mentioned above and corresponds to affect (Posner et al., 2005). It also plays an important role in constructionist theories, which aim at explaining how the objectively measurable variables of valence and arousal may be linked by cognitive processes to emotion categorizations (Feldman Barrett, 2017). While we are not aware of any applications of the constructionist theories in NLP, emotion analysis has been formulated as valence/arousal regression (Buechel and Hahn, 2017; Preotiu-Pietro et al., 2016, i.a.).

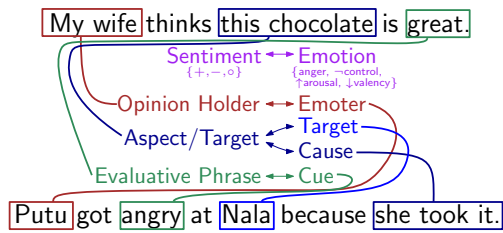


Figure 2: Comparison of structured sentiment analysis and emotion role labeling.

2.1.3 Appraisals

Affect is not the only so-called dimensional model to represent emotions. More recently, the concept of appraisals that represents the cognitive dimension of emotions, i.e., the cognitive evaluation of the event regarding the impact on the self, found attention in NLP. The set of appraisals that can explain emotions is not fixed and depends on the theory and the domain. It often includes variables that describe if an event can be expected to increase a required effort (likely to be high for anger or fear) or how much responsibility the experiencer of the emotion holds (high for feeling pride or guilt). Smith and Ellsworth (1985) showed that a comparably limited set of 6 appraisal variables can represent 15 emotion categories. Scherer et al. (2001) describes a multi-step process of appraisal evaluations as one part of the emotion – their emotion component process model also reflects on additional emotion components, namely the bodily reaction, the expression, the motivational aspect, and the subjective feeling. Appraisal theories led to a set of knowledge bases and models that link events to emotions (Balahur et al., 2012; Cambria et al., 2022; Shaikh et al., 2009; Udochukwu and He, 2015), but only recently, resources and models have been proposed which make appraisal variables explicit (Stranisci et al., 2022; Hofmann et al., 2020, 2021; Troiano et al., 2022, 2023; Wegge et al., 2022). This paper discusses work on appraisal theories to interpret events regarding the potentially resulting emotion in Section 3.1.

2.2 Tasks Related to Emotion Analysis

Emotion analysis is a task grounded in various previous research fields, from which we discuss sentiment analysis and personality profiling.

2.2.1 Sentiment Analysis

Sometimes, sentiment analysis is considered a simplified version of emotion analysis in which many emotion categories are conflated into two (positive

or negative). We would like to argue that the tasks differ in more than the number of labels. First of all, sentiment analysis is often equated to classifying the text into a more unspecific connotation of being positive or negative (Liu, 2012). Commonly, the sentiment of the text author is analyzed, which renders the task to be overlapping with opinion mining (Pang and Lee, 2008; Barnes et al., 2017). Emotion analysis is hardly ever about detecting the opinion regarding a product; while that is a common focus in sentiment analysis (Pontiki et al., 2014).

A more powerful approach to sentiment analysis is to not only detect if the author expresses something positive, but also to detect opinion holders, evaluated targets/aspects, and the phrase that describes the evaluation (Barnes et al., 2022; Pontiki et al., 2015, 2016; Klinger and Cimiano, 2013). The tasks of such “sentiment role labeling” and “emotion role labeling” do, however, barely match (see Figure 2):

- (1) The *opinion holder* in sentiment analysis is a person that expresses an opinion, regarding some object, service, or person. This commonly follows a cognitive evaluation, likely to be a conscious process rather than an unbidden reaction. We would therefore not call the person experiencing an emotion a “holder” but rather an *emotion experiencer*, or *feeler* (which is strangely ambiguous), or an *emoter* (to make the difference between an emotion and a feeling explicit).
- (2) The *aspect/target* in sentiment analysis might correspond to two things in emotion analysis. It can be a *target*, I can be angry *at* someone, who is not solely the *cause* of that emotion. I can be angry at a friend, because she did eat my emergency supply of chocolate. But I cannot be *sad at* somebody. In emotion analysis, we care more about the *stimulus* or *cause* of an emotion. Sometimes, targets and causes are conflated.
- (3) The *evaluative, subjective phrase* in sentiment analysis corresponds to emotion words (*cue* in Figure 1).

It is noteworthy that evaluative statements in sentiment also express an appraisal of something but the overlap with appraisal theories in emotion analysis is minimal – the evaluation of a product in sentiment analysis is often expressed explicitly. On the contrary, appraisal-based emotion analysis focuses on inferring the internal appraisal processes




Figure 3: Comparison of personality detection and emotion analysis.

of a person purely from an event description. We refer the interested reader to [Martin and White \(2005\)](#) for a comprehensive analysis of the language used to describe evaluations.

2.2.2 Personality Profiling

Sometimes the task of personality analysis is seen to be similar to emotion analysis, because both an emotion and the personality are based on a person. Personality is, however, a function that depends only on the person, while an emotion depends on the person in interaction with a situation (see Figure 3). Therefore, personality is a stable trait, while emotions are states that change more flexibly ([Geiser et al., 2017](#)). The most prominent model that found application in NLP is the OCEAN/Big-Five model ([Goldberg, 1999](#); [Roccas et al., 2002](#)), comprising openness to experience, conscientiousness, extraversion, agreeableness, and neuroticism ([Pizzolli and Strapparava, 2019](#); [Lynn et al., 2020](#); [Kreuter et al., 2022](#); [Golbeck et al., 2011](#)). An alternative is HEXACO, adding the dimension of honesty ([Lee and Ashton, 2018](#)), which did, however, lead to less attention in NLP ([Sinha et al., 2015](#)). The Myers–Briggs Type Indicator (MBTI, [Myers, 1998](#)) received attention in NLP, partially because of a straight-forward way to collect data with hash-tag-based self-supervision ([Plank and Hovy, 2015](#); [Verhoeven et al., 2016](#)). This model has weaknesses regarding reliability and validity ([Boyle, 1995](#); [Randall et al., 2017](#)) which affect the robustness of NLP models ([Stajner and Yenikent, 2021](#)).

2.3 Use-Cases of Emotion Analysis

Every kind of text in which an interpretation of the emotional connotation is of value constitutes a potential use case for emotion modeling. This includes the analysis of social media ([Mohammad et al., 2018](#); [Klinger et al., 2018](#); [Wang et al., 2012, i.a.](#)), of news articles ([Bostan et al., 2020, i.a.](#)), of figurative language ([Chauhan et al., 2020](#); [Dankers et al., 2019, i.a.](#)), of abusive language ([Rajaman-](#)

[ickam et al., 2020](#); [Plaza-del-Arco et al., 2021, i.a.](#)) of literature ([Kim and Klinger, 2018](#); [Alm and Sproat, 2005](#); [Dodds et al., 2011](#); [Kim et al., 2017, i.a.](#)), of clinically relevant disorders ([Islam et al., 2018](#); [Pestian et al., 2012, i.a.](#)), or the support of customer agents ([Labat et al., 2022](#)).

Each domain implicitly defines which tasks are relevant. For news headlines, the author’s emotion is least interesting while estimating the (intended) impact on the reader is important, for instance to understand reactions in the society and intentional use to manipulate readers ([Caiani and Di Cocco, 2023](#)). For hate speech detection or other social media analysis tasks, the author’s emotion is central. In literature, an interesting aspect is to understand which emotion is attributed to fictional characters ([Kim and Klinger, 2019b](#); [Hoorn and Konijn, 2003](#)).

Each domain also comes with particular challenges, stemming from varying task formulations: News headlines are short and highly contextualized in the outlet, the time of publication, and the reader’s stance towards topics ([Schaffer, 1995](#)). Social media comes in informal language ([Kern et al., 2016](#)). Literature often requires interpretations of longer text spans ([Kuhn, 2019](#)). Each of these applications therefore comes with design choices:

- What is the emotion perspective? (reader, writer, entities)
- What is the unit of analysis? (headline, tweet, paragraph, (set of) sentences)
- Is text classification of predefined units sufficient or does a model need to assign emotions to automatically detected segments in the text?
- What are the variables to be predicted and the possible value domain? (emotion categories, appraisals, affect, spans of different kind)

So far, models have mostly been developed for specific use-cases, where such constraints can be clearly identified. This has, however, an impact on the generalizability of models. We will now discuss the two perspectives of *events that cause emotions* as an interpretation of emotion analysis as text classification of predefined textual units (§3.1) and of *events as emotions*, the case of emotion role labeling (§3.2). After that, we explain the efforts to bring these two directions together (§3.3) and we build on top of this consolidation to point out important future research directions (§4).

Relevance	Implication	Coping	Normative Significance
Novelty •suddenness •familiarity •predictability •attention •att. removal Intrinsic Pleasantness •pleasant •unpleasant Goal Relevance •goal-related	Causality: agent •own responsib. •other's respons. •situational resp. Goal conduciveness •goal support Outcome probability •consequence anticipation Urgency •response urgency	Control •own control •others' control •chance control Adjustment •anticipated acceptance •effort	Internal standards compatibility •clash with own standards External stan- dards compatibility •clash with norms

Figure 4: Variables used by Troiano et al. (2023) to analyze text according to combined dimensions proposed by Scherer et al. (2001) and Smith and Ellsworth (1985).

3 The Link between Emotions and Events

3.1 Events cause Emotions: Appraisals

3.1.1 Traditional Emotion Analysis Systems

Most emotion analysis systems were, before the deep learning revolution in NLP, feature-based, and features often stemmed from manually created lexicons (Mohammad and Turney, 2013) and included manually designed features for the task (Štajner and Klinger, 2023; Aman and Szpakowicz, 2007). Since the state of the art for the development of text analysis systems is transfer learning by fine-tuning pretrained large language models (such as BERT, Devlin et al., 2019), the phenomenon-specific model development focuses on exploiting properties of the concept. One example is DeepMoji, which adapts transfer learning to the analysis of subjective language and identifies a particularly useful pretraining task, namely the prediction of emojis (Felbo et al., 2017). Another strain of research aims at developing models that aggregate multiple emotion theories (Buechel et al., 2021).

3.1.2 Event Interpretation

We focus on the aspect of emotions that they are caused by events. Interpreting events is challenging, because event descriptions often lack an explicit emotion mention. Such textual instances are considered “implicit” regarding their emotion (Udochukwu and He, 2015; Klinger et al., 2018): The challenge to be solved is to link “non-emotional” events to the emotion that they might cause. Balahur et al. (2012) tackled this by listing action units in an ontology, based on semantic parsing of large amounts of text. Cambria et al. (2022) developed a logics-based resource to associate events with their emotion interpretation.

3.1.3 Incorporating Appraisal Variables in Text Analysis Models

These attempts, however, do not model appraisal variables explicitly as a link between cognitive evaluations of events and emotions. There is also not only one appraisal theory, and depending on the theory, the computational modeling is realized in differing ways. Based on the OCC model (an appraisal theory that provides a decision tree of appraisal variables to characterize emotions, Steunebrink et al., 2009), both Shaikh et al. (2009) and Udochukwu and He (2015) develop methods to extract atomic variable values from text that are the building blocks for appraisal-based interpretations. An example appraisal variable is if an event is directed towards the self, for which they use semantic and syntactic parsers. Other such variables include the valence of events, the attitude towards objects, or the moral evaluation of people’s behaviours – all detected with polarity lexicons. These variables are then put together with logical rules, such as If Direction = ‘Self’ and Tense = ‘Future’ and Overall Polarity = ‘Positive’ and Event Polarity = ‘Positive’, then Emotion = ‘Hope’ (Udochukwu and He, 2015). The advantage of this approach is that it makes the appraisal-based interpretation explicit; however, it does not allow for reasoning under uncertainty, partially because these studies do not build on top of manually assigned appraisal variables to text.

3.1.4 Appraisal-Annotated Corpora

To understand the link better between appraisals in text and emotions, Hofmann et al. (2020) manually annotated autobiographical event reports (Troiano et al., 2019) for the appraisal dimensions identified by Smith and Ellsworth (1985): does the writer want to devote attention, were they certain about what was happening, did they have to expend mental or physical effort to deal with the situation, did they find the event pleasant, were they responsible for the situation, could they control the situation, and did they find that the situation could not be changed by anyone? They found that the annotation replicates the links to emotions as found in original studies (Hofmann et al., 2021, Fig. 1). Further, they showed that appraisals can reliably be detected, but they did not manage to develop a model that predicts emotions better with the help of appraisals than without. Hence, they proposed a new way of modeling emotions in text, but did not succeed to develop a multi-emotion-model ap-




Figure 5: The study design that led to the crowd-enVENT data set (Troiano et al., 2023).

proach.

3.1.5 Appraisal Annotations by Event Experiencers

To understand better if this inferiority of a joint model might be a result of an imperfect noisy appraisal annotation, and to create a larger corpus, Troiano et al. (2023) setup the experiment depicted in Figure 5 (replicating Troiano et al. (2019), but with appraisal variables). They asked crowdworkers to describe an event that caused a specific emotion and to then assign appraisal values (this time following the sequential approach by Scherer et al., 2001, with 21 variables, Figure 4) how they perceived the respective situation (Phase 1). They then asked other people to read the texts and reconstruct the emotion and appraisal (Phase 2). Unsurprisingly, the readers sometimes misinterpreted an event. For instance “I put together a funeral service for my Aunt” is mostly interpreted as something sad, while the original author was actually proud about it. These differences in interpretation can also be seen in the appraisal variables – Appraisals explain the differences in the event evaluation: The interpretation as being sad comes with evaluations as not being in control, while the interpretation to cause pride comes with being in control.

3.1.6 Emotion Modelling under Consideration of Appraisals

The modeling experiments of Troiano et al. (2023) confirm that also a larger set of variables can be reliably detected – similarly well as humans can reconstruct them. To further understand if such self-assigned appraisal labels enable an improvement also in the emotion categorization, they fine-tuned RoBERTa (Liu et al., 2019) and tested if adding appraisal values improves the result. They find that appraisals help the prediction of anger, fear, joy, pride, guilt, sadness, and anger. They showcase the event report “His toenails were massive.”, where the baseline model relies on something massive being associated to pride. With the appraisal information, it correctly assigns “disgust”.

3.1.7 Other research directions

More recently other research has been published with a focused on specific use-cases. Stranisci et al. (2022) who follow the appraisal model by Roseman (2013) postannotate Reddit posts which deal with situations that challenged the author to cope with an undesirable situation. Their APPReddit corpus is the first resource of appraisal-annotated texts from the wild. Cortal et al. (2023) follow a similar idea and acquire texts that describe how people regulate their emotions in specific situations. Next to their resource creation effort for French, they analyze which descriptions of cognitive processes allow to infer an emotion.

We conclude that appraisal-based emotion analysis research has the goal to better understand how emotions are implicitly communicated and to develop better emotion analysis systems.

3.2 Emotions are Events: Structured Analysis

The studies that we discussed so far put the aspect of emotion analysis on the spot that emotions are caused by events. As we argued before, emotions also constitute events. Similarly to the field of semantic role labeling (Gildea and Jurafsky, 2000) which models events in text following frame semantics, various efforts have been made to extract emotion event representations from text. The corpora that have been created come with differing modeling attempts, summarized in Figure 6.

3.2.1 Cue Phrase Detection

The early work by Aman and Szpakowicz (2007) focused on the emotion *cue* word, as an important part of role labeling. They annotated sentences from blogs, but did not propose an automatic cue identification system. A structurally similar resource with cue word annotations has been proposed by Liew et al. (2016).

3.2.2 Stimulus Detection

A few corpora have been developed focussing on stimuli: Ghazi et al. (2015) annotated sentences from FrameNet that are known to be associated to emotions and model the automatic prediction as sequence labeling. For German, Doan Dang et al. (2021) created a similar corpus based on news headlines. Gao et al. (2017) formulated stimulus detection as clause classification in Mandarin, which might, however, not be an appropriate approach for English (Oberländer et al., 2020).

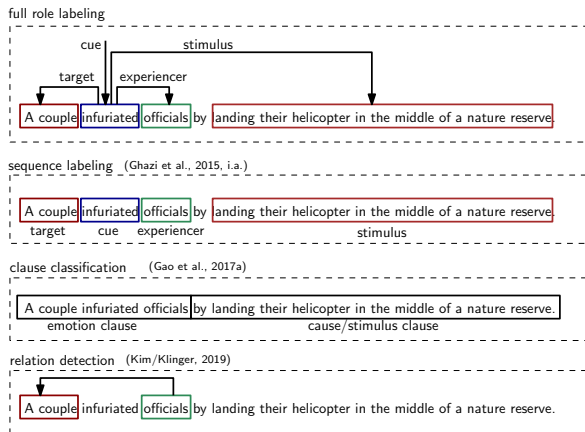


Figure 6: Emotion Role Modeling approaches (example from Bostan et al. (2020)). Full emotion role labeling has not been performed yet (top).

3.2.3 Role Labeling as Classification

An interesting attempt of emotion role labeling in texts from social media was the study on Tweets associated to a US election by Mohammad et al. (2014). The decision to focus on a narrow domain allowed them to frame the role identification task both in crowdsourced annotation and in modeling as a classification task; namely to decide if the emoter, the stimulus or the emotion target correspond to an entity from a predefined set (this modeling formulation is not shown in Figure 6).

3.2.4 Full Emotion Role Labeling Resources

Kim and Klinger (2018) and Bostan et al. (2020) aimed at creating corpora with full emotion role labeling information. The REMAN corpus (Kim and Klinger, 2019b) focused on literature from Project Gutenberg. Given the challenging domain, the authors decided to carefully train annotators instead of relying on crowdsourcing. Each instance corresponds to a sentence triple, in which the middle sentence contains the cue to which the roles of emoters, targets, and stimuli are to be associated. The sequence-labeling-based modeling revealed that cause and target detection are very challenging. The paper does not contain an effort to reconstruct the full emotion event graph structure.

Bostan et al. (2020) annotated news headlines, under the assumption that less context is required for interpretation (which turned out to not be true). To attribute for the subjective nature of emotion interpretations, they setup the annotation as a multistep crowdsourcing task. The modeling experiments on their GoodNewsEveryone corpus are limited to span prediction.




Figure 7: Example from the x-enVENT dataset

3.2.5 Role Labeling as Relation Detection

We are only aware of one work in the context of semantic role labeling that attempts to model the relational structure. Kim and Klinger (2019b) simplified role labeling to relation classification of emotional relations between entities. This allowed them to build on top of established methods for relation detection (Zhou et al., 2016) but they sacrificed explicit cue word detection and limited the analysis to emotion stimuli that have a corresponding entity.

3.2.6 Aggregated Corpora

There have been two efforts of data aggregation, by Oberländer et al. (2020) and Campagnano et al. (2022). The latter compared various models for role detection via span prediction. The prior we will discuss in the next section. To sum up, there have been some efforts to perform emotion role labeling, but in contrast to generic role labeling or to structured sentiment analysis, no models have yet been developed for full graph reconstruction. We visualize the differences in modeling attempts in Figure 6.

3.3 Bridging the Two Perspectives

We now discussed the two perspectives of *events causing emotions* (§3.1) and *emotions being events* (§3.2). The fact that these two analysis tasks have so far mostly been tackled separately leaves a lot of space for future research. However, some attempts to link the two areas already exist.

3.3.1 Do the tasks of emotion classification and role labeling benefit from each other?

Oberländer et al. (2020) aimed at understanding if knowledge of roles impacts the performance of emotion categorization. It turns out it does, either because the relevant part of the text is made more explicit (stimulus), or because of biases (emoter). Similarly, Xia and Ding (2019) setup the task of stimulus-clause and emotion-clause pair classification. Their corpora and a plethora of follow-up

work show that stimulus and emotion detection benefit from each other.

3.3.2 Descriptions of which emotion components enable emotion recognition?

A similar strain of research aims at understanding which components of emotions support emotion predictions. [Casel et al. \(2021\)](#) performed multi-task learning experiments with emotion categorization and emotion component prediction. [Kim and Klinger \(2019a\)](#) study how specific emotions are communicated, similarly to [Etienne et al. \(2022\)](#). [Cortal et al. \(2023\)](#) analyzed if particular ways of cognitively evaluating events support the emotion prediction more than others.

3.3.3 Linking Role Labeling and Appraisal-based Analysis

These works do, however, not link emotion roles explicitly to their cognitive evaluation dimensions. The only work that aimed at doing so is the corpus by [Troiano et al. \(2022\)](#), who label emoters for emotion categories and appraisals, the events that act as a stimulus on the token level, and the relation between them. [Figure 7](#) shows an example from their corpus. In their modeling efforts, however, they limited themselves to emoter-specific emotion/appraisal predictions and ignored, so far, the span-based stimulus annotations ([Wegge et al., 2022](#); [Wegge and Klinger, 2023](#)).

4 Open Research Tasks

We have now discussed previous work in emotion analysis, appraisal-based approaches and role labeling. In the following, we will make a set of aspects explicit that, from our perspective, need future work.

Full emotion role labeling. Several corpora exist now that have complex annotations of the emoter, their respective emotion stimuli, targets, and cue words; partially with sentence level annotations for the reader and writer in addition. Modeling, however, focused on sequence labeling for subsets of the roles or sentence level classification. There are no attempts of full emotion graph prediction, despite that role prediction subtasks might benefit from being modeled jointly. There is also only little work on exploiting role information for emotion categorization on the sentence level, a potentially valuable approach for joint modeling of a structured prediction task with text classification.

Role labeling/stimulus detection with appraisal information. The work that has been performed to understand the interaction between role prediction and emotion categorization focused on predicting discrete emotion classes. However, stimuli often correspond to event descriptions and therefore are a straight-forward choice for further analysis with appraisal variables. Also, understanding which event mentions in a text can function as an emotion stimulus could be supported with the help of appraisals. The detection of clauses or token sequences that correspond to emotion stimuli in context of appraisal-based interpretations therefore has potential to improve both subtasks.

Integration of other emotion models in role labeling. Emotion categorization is typically one variable to be predicted in stimulus detection and role labeling approaches, either for a writer or for entities. An additionally interesting approach would be to integrate other emotion representations with role labeling. An interesting choice would be to create a corpus of valence/arousal values, assigned to specific entities and linked to stimuli. Such approach comes with the general advantage of dimensional models, namely that emotion categories do not need to be predefined.

Robust cross-corpus modeling and zero-shot predictions. A similar motivation lead to recent work on zero-shot emotion prediction, in which emotion categories are to be predicted that are not available in the training data. [Plaza-del-Arco et al. \(2021\)](#) showed that the performance loss of natural language inference-based prompting in comparison to supervised learning leaves space for improvements. Such attempts might also bridge the gap between in-domain performance and cross-domain performance of emotion analysis systems ([Bostan and Klinger, 2018](#)). Zero-shot modeling or other approaches to find representations that are agnostic to the underlying emotion theory are essential for cross-corpus experiments, because the domains that are represented by different corpora require differing label sets.

Interpretation of event chains. Textual event descriptions can be interpreted with appraisal theories, but we rely on end-to-end learning to understand how sequences of events lead to specific emotions (for instance being afraid of a specific unconfirmed undesirable event $e \rightarrow e$ is disconfirmed \rightarrow relief). Dissecting events with semantic parsing, and com-

binning them with emotion role labeling leads to sequences of general and emotion events, which can be the input for a second-level emotion analysis. Such methods would be required to fully understand how emotions develop throughout longer sequences of stories, for instance in literature.

Perspectivism. Appraisals do explain differences in the emotion assessment, based on differing interpretations of events (Troiano et al., 2023). We do, however, not know the role of underlying factors. A perspectivistic approach with the goal to uncover variables that lead to varying emotion constructions, e.g., based on demographic data of event participants or other data, might provide additional insight. This could also be applied to literature analysis, for instance by including personality information on fictional characters in the emotion prediction (Bamman et al., 2013). Such approach is well-motivated in psychology; we know that personality influences the interpretation of other’s emotions (Doellinger et al., 2021).

Integrate emotion models from psychology. Emotion analysis work so far focused on a comparably small set of emotion theories. The philosophical discussion by Scarantino (2016) offers itself as a guiding principle which other theories might be valuable to be explored. This does not only include entirely so-far-ignored theories (e.g., Feldman Barrett, 2017) but also knowledge from theories popular in NLP. For instance, Ekman (1992); Plutchik (2001) offer more information than lists of emotion categories. Integrating psychological knowledge in NLP models can improve the performance (Troiano et al., 2023). In a similar vein, there exist specific appraisal theories for particular domains, including, e.g., argument mining (Dillard and Seo, 2012).

Multimodal Modeling. We focused in our paper on analysis tasks from text, but there has already been work on multimodal emotion analysis (Busso et al., 2008, i.a.) and detecting emotion stimuli in images (Dellagiacomma et al., 2011; Fan et al., 2018, i.a.), also multimodally (Khlyzova et al., 2022; Cevher et al., 2019). However, we are not aware of any work in computer vision that interprets situations and the interactions of events with the help of appraisal theories. To fully grasp available information in everyday communication or (social) media, the presented approaches from this paper need to be extended multimodally.

Multilingual modeling. Most papers that we discuss in this paper focus on English – with very few exceptions, which we pointed out explicitly. We are not aware of any emotion role labeling corpus with full graph annotations in other languages, and there are only very few attempts to integrate appraisal theories in emotion detection on languages other than English. Such multilingual extension is not only relevant to achieve models that work across use-cases – the concept of emotion names might also differ between languages, and therefore comparing emotion concepts with the help of dimensional appraisal models between languages and cultures can provide interesting insights for both NLP and psychology.

5 Conclusion

With this paper, we discussed appraisal theory-based methods to interpret events, and how emotions can be represented as events with role labeling. We did that guided by our own two emotion analysis projects SEAT (Structured Multi-Domain Emotion Analysis from Text) and CEAT (Computational Event Evaluation based on Appraisal Theories for Emotion Analysis) which corresponded each to one of the two perspectives.

These two fields have been approached mostly separately so far and the main goal of this paper is to make the research narrative behind both transparent, and, based on this, point out open research questions. Such open tasks emerge from missing connections between the various goals in emotion analysis, but there are also other promising directions that we pointed out.

We do not believe that this list is comprehensive, but hope that the aggregation of previous work and pointing at missing research helps interested researchers to identify the gaps they want to fill. Emotion analysis is important to make computers aware of the concept, which is essential for natural communication.

In addition, research in these fields helps to better understand how humans communicate, beyond building impactful computational systems. Therefore, research in affective computing brings together psychology, linguistics, and NLP.

Limitations

This paper focused on appraisal theories and emotion role labeling mostly from a theoretical perspective. We aimed at pointing out open research

questions mostly based on conceptualizations of theories from semantics and psychology. To identify open research questions, a closer introspection of existing models need to be performed in addition. In our theoretical discussion, we assume that the open research questions have similar chances to succeed. In practical terms this is likely not the case and we therefore propose to first perform preliminary studies before definitely deciding to follow one of the research plans that we sketched.

Ethics Statement

The contributions in this paper do not directly pose any ethical issues: we did not publish data, models, or did perform experiments. However, the open topics that we identified might lead to resources and models that can in principle do harm to people. Following deontological ethics, we assume that no emotion analysis systems should be applied to data created by a person without their consent, if the results are used not only in aggregated form which would allow to identify the person who is associated with the analyzed data. We personally do not believe that a utilitarian approach may be acceptable in which reasons could exist that justify to use emotion analysis technology to identify individuals from a larger group. This is particularly important with methods discussed in this paper in comparison to more general emotion categorization methods, because we focus on implicit emotion expressions. The methods we discussed and future work we sketched would be able to identify emotions that are not explicitly expressed, and therefore humans that generate data might not be aware that their private emotional state could be reconstructed from the data they produce.

When creating data for emotion analysis, independent of its language, domain, or the task formulation as role labeling, classification, regression, using a dimensional model or a theory of basic emotions, fairness or developed system and bias in data and systems is typically an issue. While efforts exist to identify unwanted bias and confounders in automatic analysis systems, the possible existence of unidentified biases can never be excluded. Therefore, automatic systems always need to be applied with care while critically reflecting the automatically obtained results. This is particularly the case with systems that focus on interpreting implicit emotion communications that require reasoning under uncertainty. To enable such critical

reflection of a system's output, their decision must be transparently communicated to the users.

In general, the ability of automatic systems to interpret and aggregate emotions should not be used unaware of the people who created data, and decisions and actions following recognized emotions always need to remain in the responsibility of a human user.

We see our work mostly as a research contribution with the goal to better understand how humans communicate, not as an automatic enabling tool to provide insight in the private states of people.

Acknowledgements

We would like to thank all coauthors who contributed to our work on emotion analysis with the help of appraisal theories and in role labeling. These are (in alphabetical order) Amelie Heindl, Antje Schweitzer, Bao Minh Doan Dang, Enrica Troiano, Evgeny Kim, Felix Casel, Flor Miriam Arco Del Plaza, Jan Hofmann, Jeremy Barnes, Kai Sassenberg, Kevin Reich, Laura Oberländer née Bostan, Max Wegge, Sebastian Padó, Tornike Tsereteli, and Valentino Sabbatino. We further thank Alexandra Balahur, Orphée De Clercq, Saif Mohammad, Veronique Hoste, Valentin Barriere, and Sanja Štajner for discussions on the general topics of emotion analysis that helped us to develop this paper.

This work has been funded by two projects of the German Research Council (Deutsche Forschungsgemeinschaft), namely the project "Structured Multi-Domain Emotion Analysis from Text" (SEAT, KL 2869/1-1) and "Computational Event Evaluation based on Appraisal Theories for Emotion Analysis" (CEAT, KL 2869/1-2).¹

References

Cecilia Ovesdotter Alm, Dan Roth, and Richard Sproat. 2005. *Emotions from text: Machine learning for text-based emotion prediction*. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 579–586, Vancouver, British Columbia, Canada. Association for Computational Linguistics.

Cecilia Ovesdotter Alm and Richard Sproat. 2005. *Emotional sequencing and development in fairy tales*. In *Affective Computing and Intelligent Interaction*,

¹<https://www.ims.uni-stuttgart.de/en/research/projects/seat/>, <https://www.ims.uni-stuttgart.de/en/research/projects/ceat/>

- pages 668–674, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Saima Aman and Stan Szpakowicz. 2007. [Identifying expressions of emotion in text](#). In *Text, Speech and Dialogue*, pages 196–205, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Alexandra Balahur, Jesus M. Hermida, and Andrew Montoyo. 2012. [Building and exploiting emotinet, a knowledge base for emotion detection based on the appraisal theory model](#). *IEEE Transactions on Affective Computing*, 3(1):88–101.
- David Bamman, Brendan O’Connor, and Noah A. Smith. 2013. [Learning latent personas of film characters](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 352–361, Sofia, Bulgaria. Association for Computational Linguistics.
- Jeremy Barnes, Roman Klinger, and Sabine Schulte im Walde. 2017. [Assessing state-of-the-art sentiment models on state-of-the-art sentiment datasets](#). In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, Copenhagen, Denmark. Workshop at Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics.
- Jeremy Barnes, Laura Oberlaender, Enrica Troiano, Andrey Kutuzov, Jan Buchmann, Rodrigo Agerri, Lilja Øvrelid, and Erik Velldal. 2022. [SemEval 2022 task 10: Structured sentiment analysis](#). In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1280–1295, Seattle, United States. Association for Computational Linguistics.
- Laura Ana Maria Bostan, Evgeny Kim, and Roman Klinger. 2020. [GoodNewsEveryone: A corpus of news headlines annotated with emotions, semantic roles, and reader perception](#). In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 1554–1566, Marseille, France. European Language Resources Association.
- Laura-Ana-Maria Bostan and Roman Klinger. 2018. [An analysis of annotated corpora for emotion classification in text](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2104–2119, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Gregory J. Boyle. 1995. [Myers-Briggs Type Indicator \(MBTI\): Some Psychometric Limitations](#). *Australian Psychologist*, 30(1):71–74.
- Sven Buechel and Udo Hahn. 2017. [EmoBank: Studying the impact of annotation perspective and representation format on dimensional emotion analysis](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 578–585, Valencia, Spain. Association for Computational Linguistics.
- Sven Buechel, Luise Modersohn, and Udo Hahn. 2021. [Towards label-agnostic emotion embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9231–9249, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeanette N. Chang, Sungbok Lee, and Shrikanth S. Narayanan. 2008. [Iemocap: interactive emotional dyadic motion capture database](#). *Language Resources and Evaluation*, 42(4):335.
- Manuela Caiani and Jessica Di Cocco. 2023. [Populism and emotions: a comparative study using machine learning](#). *Italian Political Science Review / Rivista Italiana di Scienza Politica*, page 1–16.
- Erik Cambria, Qian Liu, Sergio Decherchi, Frank Xing, and Kenneth Kwok. 2022. [SenticNet 7: A commonsense-based neurosymbolic AI framework for explainable sentiment analysis](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3829–3839, Marseille, France. European Language Resources Association.
- Cesare Campagnano, Simone Conia, and Roberto Navigli. 2022. [SRL4E – Semantic Role Labeling for Emotions: A unified evaluation framework](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4586–4601, Dublin, Ireland. Association for Computational Linguistics.
- Felix Casel, Amelie Heindl, and Roman Klinger. 2021. [Emotion recognition under consideration of the emotion component process model](#). In *KONVENS 2021*.
- Deniz Cevher, Sebastian Zepf, and Roman Klinger. 2019. [Towards multimodal emotion recognition in german speech events in cars using transfer learning](#). In *KONVENS*.
- Dushyant Singh Chauhan, Dhanush S R, Asif Ekbal, and Pushpak Bhattacharyya. 2020. [Sentiment and emotion help sarcasm? a multi-task learning framework for multi-modal sarcasm, sentiment and emotion analysis](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4351–4360, Online. Association for Computational Linguistics.
- Gustave Cortal, Alain Finkel, Patrick Paroubek, and Lina Ye. 2023. [Emotion recognition based on psychological components in guided narratives for emotion regulation](#). In *Proceedings of the 7th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 72–81, Dubrovnik, Croatia. Association for Computational Linguistics.

- Verna Dankers, Marek Rei, Martha Lewis, and Ekaterina Shutova. 2019. [Modelling the interplay of metaphor and emotion through multitask learning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2218–2229, Hong Kong, China. Association for Computational Linguistics.
- Michela Dellagiacoma, Pamela Zontone, Giulia Boato, and Liliana Albertazzi. 2011. [Emotion based classification of natural images](#). In *Proceedings of the 2011 International Workshop on DETecting and Exploiting Cultural Diversity on the Social Web, DETECT '11*, page 17–22, New York, NY, USA. Association for Computing Machinery.
- Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. 2020. [GoEmotions: A dataset of fine-grained emotions](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4040–4054, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- James Price Dillard and Kiwon Seo. 2012. [Affect and persuasion](#). In James Price Dillard and Lijiang Shen, editors, *The SAGE Handbook of Persuasion: Developments in Theory and Practice*, chapter 10. SAGE Publications.
- Bao Minh Doan Dang, Laura Oberländer, and Roman Klinger. 2021. [Emotion stimulus detection in German news headlines](#). In *Proceedings of the 17th Conference on Natural Language Processing (KONVENS 2021)*, pages 73–85, Düsseldorf, Germany. KONVENS 2021 Organizers.
- Peter Sheridan Dodds, Kameron Decker Harris, Isabel M. Kloumann, Catherine A. Bliss, and Christopher M. Danforth. 2011. [Temporal patterns of happiness and information in a global social network: Hedonometrics and twitter](#). *PLOS ONE*, 6(12):1–1.
- Lillian Doellinger, Petri Laukka, Lennart Björn Högman, Tanja Bänziger, Irena Makower, Håkan Fischer, and Stephan Hau. 2021. [Training emotion recognition accuracy: Results for multimodal expressions and facial micro expressions](#). *Frontiers in Psychology*, 12.
- Paul Ekman. 1992. [An argument for basic emotions](#). *Cognition & emotion*, 6(3-4):169–200.
- Paul Ekman and Daniel Cordaro. 2011. [What is meant by calling emotions basic](#). *Emotion Review*.
- Aline Etienne, Delphine Battistelli, and GwénoLé Lecorvé. 2022. [A \(psycho-\)linguistically motivated scheme for annotating and exploring emotions in a genre-diverse corpus](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 603–612, Marseille, France. European Language Resources Association.
- Shaojing Fan, Zhiqi Shen, Ming Jiang, Bryan L. Koenig, Juan Xu, Mohan S. Kankanhalli, and Qi Zhao. 2018. [Emotional attention: A study of image sentiment and visual attention](#). In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7521–7531.
- Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. 2017. [Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1615–1625, Copenhagen, Denmark. Association for Computational Linguistics.
- Lisa Feldman Barrett. 2017. [The theory of constructed emotion: an active inference account of interoception and categorization](#). *Social Cognitive and Affective Neuroscience*, 12(11):1833.
- Qinghong Gao, Jiannan Hu, Ruifeng Xu, Gui Lin, Yulan He, Qin Lu, and Kam-Fai Wong. 2017. [Overview of NTCIR-13 ECA task](#). In *Proceedings of the 13th NTCIR Conference on Evaluation of Information Access Technologies*, pages 361–366, Tokyo, Japan.
- Christian Geiser, Thomas Götz, Franzis Preckel, and Philipp Alexander Freund. 2017. [States and Traits: Theories, Models, and Assessment](#). *European Journal of Psychological Assessment*, 33(4):219–223.
- Diman Ghazi, Diana Inkpen, and Stan Szpakowicz. 2015. [Detecting emotion stimuli in emotion-bearing sentences](#). In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 152–165. Springer.
- Daniel Gildea and Daniel Jurafsky. 2000. [Automatic labeling of semantic roles](#). In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, pages 512–520, Hong Kong. Association for Computational Linguistics.
- Jennifer Golbeck, Cristina Robles, and Karen Turner. 2011. [Predicting personality with social media](#). In *CHI'11 extended abstracts on human factors in computing systems*, pages 253–262.
- Lewis R. Goldberg. 1999. [A broad-bandwidth, public domain, personality inventory measuring the lower-level facets of several five-factor models](#). *Personality psychology in Europe*, 7(1):7–28.

- Jan Hofmann, Enrica Troiano, and Roman Klinger. 2021. [Emotion-aware, emotion-agnostic, or automatic: Corpus creation strategies to obtain cognitive event appraisal annotations](#). In *Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 160–170, Online. Association for Computational Linguistics.
- Jan Hofmann, Enrica Troiano, Kai Sassenberg, and Roman Klinger. 2020. [Appraisal theories for emotion classification in text](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 125–138, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Johan F. Hoorn and Elly A. Konijn. 2003. [Perceiving and experiencing fictional characters: An integrative account: Perceiving and experiencing fictional characters](#). *Japanese Psychological Research*, 45(4):250–268.
- Md Rafiqul Islam, Muhammad Ashad Kabir, Ashir Ahmed, Abu Raihan M Kamal, Hua Wang, and Anwaar Ulhaq. 2018. [Depression detection from social network data using machine learning techniques](#). *Health Inf. Sci. Syst.*, 6(1):8.
- Margaret L. Kern, Gregory Park, Johannes C. Eichstaedt, H. Andrew Schwartz, Maarten Sap, Laura K. Smith, and Lyle H. Ungar. 2016. [Gaining insights from social media language: Methodologies and challenges](#). *Psychological Methods*, 21(4):507–525.
- Anna Khlyzova, Carina Silberer, and Roman Klinger. 2022. [On the complementarity of images and text for the expression of emotions in social media](#). In *Proceedings of the 12th Workshop on Computational Approaches to Subjectivity, Sentiment & Social Media Analysis*, pages 1–15, Dublin, Ireland. Association for Computational Linguistics.
- Evgeny Kim and Roman Klinger. 2018. [Who feels what and why? annotation of a literature corpus with semantic roles of emotions](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1345–1359, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Evgeny Kim and Roman Klinger. 2019a. [An analysis of emotion communication channels in fan-fiction: Towards emotional storytelling](#). In *Proceedings of the Second Workshop on Storytelling*, pages 56–64, Florence, Italy. Association for Computational Linguistics.
- Evgeny Kim and Roman Klinger. 2019b. [Frowning Frodo, wincing Leia, and a seriously great friendship: Learning to classify emotional relationships of fictional characters](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 647–653, Minneapolis, Minnesota. Association for Computational Linguistics.
- Evgeny Kim, Sebastian Padó, and Roman Klinger. 2017. [Investigating the relationship between literary genres and emotional plot development](#). In *Proceedings of the Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 17–26, Vancouver, Canada. Association for Computational Linguistics.
- Roman Klinger and Philipp Cimiano. 2013. [Bi-directional inter-dependencies of subjective expressions and targets and their value for a joint model](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 848–854, Sofia, Bulgaria. Association for Computational Linguistics.
- Roman Klinger, Orphée De Clercq, Saif Mohammad, and Alexandra Balahur. 2018. [IEST: WASSA-2018 implicit emotions shared task](#). In *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 31–42, Brussels, Belgium. Association for Computational Linguistics.
- Anne Kreuter, Kai Sassenberg, and Roman Klinger. 2022. [Items from psychometric tests as training data for personality profiling models of Twitter users](#). In *Proceedings of the 12th Workshop on Computational Approaches to Subjectivity, Sentiment & Social Media Analysis*, pages 315–323, Dublin, Ireland. Association for Computational Linguistics.
- Jonas Kuhn. 2019. [Computational text analysis within the Humanities: How to combine working practices from the contributing fields?](#) *Language Resources and Evaluation*, 53(4):565–602.
- Sofie Labat, Amir Hadifar, Thomas Demeester, and Veronique Hoste. 2022. [An emotional journey: Detecting emotion trajectories in Dutch customer service dialogues](#). In *Proceedings of the Eighth Workshop on Noisy User-generated Text (W-NUT 2022)*, pages 106–112, Gyeongju, Republic of Korea. Association for Computational Linguistics.
- Kibeom Lee and Michael C Ashton. 2018. [Psychometric properties of the hexaco-100](#). *Assessment*, 25(5):543–556.
- Xiang Li, Aynaz Taheri, Lifu Tu, and Kevin Gimpel. 2016. [Commonsense knowledge base completion](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1445–1455, Berlin, Germany. Association for Computational Linguistics.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. [DailyDialog: A manually labelled multi-turn dialogue dataset](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.

- Jasy Suet Yan Liew, Howard R. Turtle, and Elizabeth D. Liddy. 2016. *EmoTweet-28: A fine-grained emotion corpus for sentiment analysis*. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1149–1156, Portorož, Slovenia. European Language Resources Association (ELRA).
- Bing Liu. 2012. *Sentiment analysis and opinion mining*. Synthesis lectures on human language technologies. Springer Nature Switzerland.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. arXiv:1907.11692. <https://arxiv.org/abs/1907.11692>.
- Veronica Lynn, Niranjana Balasubramanian, and H. Andrew Schwartz. 2020. *Hierarchical modeling for user personality prediction: The role of message-level attention*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5306–5316, Online. Association for Computational Linguistics.
- J. R. Martin and P. R. R. White. 2005. *The Language of Evaluation*. Palgrave Macmillan UK, London.
- Saif Mohammad. 2012. *#emotional tweets*. In **SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 246–255, Montréal, Canada. Association for Computational Linguistics.
- Saif Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. *SemEval-2018 task 1: Affect in tweets*. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 1–17, New Orleans, Louisiana. Association for Computational Linguistics.
- Saif Mohammad and Peter D. Turney. 2013. *Crowdsourcing a word-emotion association lexicon*. *Computational Intelligence*, 29(3):436–465.
- Saif Mohammad, Xiaodan Zhu, and Joel Martin. 2014. *Semantic role labeling of emotions in tweets*. In *Proceedings of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 32–41, Baltimore, Maryland. Association for Computational Linguistics.
- Isabel Briggs Myers. 1998. *Introduction to Type: A Guide to Understanding Your Results on the MBTI Instrument*, 6th edition edition. Cpp. Inc.
- William Herman Newman, Charles Edgar Summer, and E. Kirby Warren. 1967. *The Process of Management: Concepts, Behaviour, and Practice*. Prentice-Hall.
- Laura Oberländer, Kevin Reich, and Roman Klinger. 2020. *Experiencers, stimuli, or targets: Which semantic roles enable machine learning to infer the emotions?* In *Proceedings of the Third Workshop on Computational Modeling of People’s Opinions, Personality, and Emotions in Social Media*, Barcelona, Spain. Association for Computational Linguistics.
- Bo Pang and Lillian Lee. 2008. *Opinion Mining and Sentiment Analysis*. *Foundations and Trends® in Information Retrieval*, 2(1–2):1–135.
- John P. Pestian, Pawel Matykiewicz, Michelle Linn-Gust, Brett South, Ozlem Uzuner, Jan Wiebe, K Bretonnel Cohen, John Hurdle, and Christopher Brew. 2012. *Sentiment analysis of suicide notes: A shared task*. *Biomed. Inform. Insights*, 5(Suppl 1):3–16.
- Daniele Pizzolli and Carlo Strapparava. 2019. *Personality traits recognition in literary texts*. In *Proceedings of the Second Workshop on Storytelling*, pages 107–111, Florence, Italy. Association for Computational Linguistics.
- Barbara Plank and Dirk Hovy. 2015. *Personality traits on Twitter—or—How to get 1,500 personality tests in a week*. In *Proceedings of the 6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 92–98, Lisboa, Portugal. Association for Computational Linguistics.
- Flor M. Plaza-del-Arco, Sercan Halat, Sebastian Padã³, and Roman Klinger. 2021. *Multi-task learning with sentiment, emotion, and target detection to recognize hate speech and offensive language*. In *FIRE 2021 Working Notes*, pages 297–318.
- Robert Plutchik. 2001. *The nature of emotions human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice*. *American Scientist*, 89(4):344–350.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, Mohammad AL-Smadi, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphée De Clercq, Véronique Hoste, Marianna Apidianaki, Xavier Tannier, Natalia Loukachevitch, Evgeniy Kotelnikov, Nuria Bel, Salud María Jiménez-Zafra, and Gülşen Eryiğit. 2016. *SemEval-2016 task 5: Aspect based sentiment analysis*. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 19–30, San Diego, California. Association for Computational Linguistics.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. *SemEval-2015 task 12: Aspect based sentiment analysis*. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 486–495, Denver, Colorado. Association for Computational Linguistics.
- Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Haris Papageorgiou, Ion Androutsopoulos, and Suresh

- Manandhar. 2014. **SemEval-2014 task 4: Aspect based sentiment analysis**. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 27–35, Dublin, Ireland. Association for Computational Linguistics.
- Jonathan Posner, James A. Russell, and Bradley S. Peterson. 2005. **The circumplex model of affect: an integrative approach to affective neuroscience, cognitive development, and psychopathology**. *Development and Psychopathology*, 17(3):715–734.
- Daniel Preoțiu-Pietro, H. Andrew Schwartz, Gregory Park, Johannes Eichstaedt, Margaret Kern, Lyle Ungar, and Elisabeth Shulman. 2016. **Modelling valence and arousal in Facebook posts**. In *Proceedings of the 7th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 9–15, San Diego, California. Association for Computational Linguistics.
- Santhosh Rajamanickam, Pushkar Mishra, Helen Yanakoudakis, and Ekaterina Shutova. 2020. **Joint modelling of emotion and abusive language detection**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4270–4279, Online. Association for Computational Linguistics.
- Ken Randall, Mary Isaacson, and Carrie Ciro. 2017. **Validity and reliability of the myers-briggs personality type indicator: A systematic review and meta-analysis**. *Journal of Best Practices in Health Professions Diversity*, 10(1):1–27.
- Sonia Roccas, Lilach Sagiv, Shalom H Schwartz, and Ariel Knafo. 2002. **The big five personality factors and personal values**. *Personality and social psychology bulletin*, 28(6):789–801.
- Ira J. Roseman. 2013. **Appraisal in the Emotion System: Coherence in Strategies for Coping**. *Emotion Review*, 5(2):141–149.
- Andrea Scarantino. 2016. **The philosophy of emotions and its impact on affective science**. In *Handbook of emotions*, chapter 4, pages 3–48. Guilford Press New York, NY.
- Deborah Schaffer. 1995. **SHOCKING SECRETS REVEALED! The Language of Tabloid Headlines**. *ETC: A Review of General Semantics*, 52(1):27–46. Publisher: Institute of General Semantics.
- Klaus R. Scherer, Angela Schorr, and Tom Johnstone. 2001. **Appraisal considered as a process of multi-level sequential checking**, volume 92. Oxford University Press.
- Hendrik Schuff, Jeremy Barnes, Julian Mohme, Sebastian Padó, and Roman Klinger. 2017. **Annotation, modelling and analysis of fine-grained emotions on a stance and sentiment detection corpus**. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 13–23, Copenhagen, Denmark. Association for Computational Linguistics.
- Mostafa Al Masum Shaikh, Helmut Prendinger, and Mitsuru Ishizuka. 2009. **A linguistic interpretation of the OCC emotion model for affect sensing from text**. *Affective Information Processing*.
- Priyanka Sinha, Lipika Dey, Pabitra Mitra, and Anupam Basu. 2015. **Mining HEXACO personality traits from enterprise social media**. In *Proceedings of the 6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 140–147, Lisboa, Portugal. Association for Computational Linguistics.
- Craig. A. Smith and Phoebe. C. Ellsworth. 1985. **Patterns of cognitive appraisal in emotion**. *Journal of Personality and Social Psychology*, 48(4).
- Sanja Stajner and Seren Yenikent. 2021. **Why is MBTI personality detection from texts a difficult task?** In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3580–3589, Online. Association for Computational Linguistics.
- Bas R. Steunebrink, Mehdi Dastani, and John-Jules Ch. Meyer. 2009. **The occ model revisited**. Online: https://people.idsia.ch/~steunebrink/Publications/KI09_OCC_revisited.pdf.
- Marco Antonio Stranisci, Simona Frenda, Eleonora Ceccaldi, Valerio Basile, Rossana Damiano, and Viviana Patti. 2022. **APPReddit: a corpus of Reddit posts annotated for appraisal**. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3809–3818, Marseille, France. European Language Resources Association.
- Carlo Strapparava and Rada Mihalcea. 2007. **SemEval-2007 task 14: Affective text**. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 70–74, Prague, Czech Republic. Association for Computational Linguistics.
- Enrica Troiano, Laura Ana Maria Oberlaender, Maximilian Wegge, and Roman Klinger. 2022. **x-enVENT: A corpus of event descriptions with experiencer-specific emotion and appraisal annotations**. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 1365–1375, Marseille, France. European Language Resources Association.
- Enrica Troiano, Laura Oberländer, and Roman Klinger. 2023. **Dimensional modeling of emotions in text with appraisal theories: Corpus creation, annotation reliability, and prediction**. *Computational Linguistics*, 49(1).
- Enrica Troiano, Sebastian Padó, and Roman Klinger. 2019. **Crowdsourcing and validating event-focused emotion corpora for German and English**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4005–4011, Florence, Italy. Association for Computational Linguistics.

- Théo Trouillon, Christopher R. Dance, Éric Gaussier, Johannes Welbl, Sebastian Riedel, and Guillaume Bouchard. 2017. [Knowledge graph completion via complex tensor factorization](#). *Journal Machine Learning Research*, 18(1):4735–4772.
- Orizu Udochukwu and Yulan He. 2015. [A rule-based approach to implicit emotion detection in text](#). In *Natural Language Processing and Information Systems*.
- Ben Verhoeven, Walter Daelemans, and Barbara Plank. 2016. [TwiSty: A multilingual Twitter stylometry corpus for gender and personality profiling](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1632–1637, Portorož, Slovenia. European Language Resources Association (ELRA).
- Sanja Štajner and Roman Klinger. 2023. [Emotion analysis from texts](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics: Tutorial Abstracts*, pages 7–12, Dubrovnik, Croatia. Association for Computational Linguistics.
- Wenbo Wang, Lu Chen, Krishnaprasad Thirunarayan, and Amit P. Sheth. 2012. [Harnessing Twitter ‘Big Data’ for Automatic Emotion Identification](#). In *2012 ASE/IEEE International Conference on Social Computing and 2012 ASE/IEEE International Conference on Privacy, Security, Risk and Trust*, pages 587–592. IEEE.
- Maximilian Wegge and Roman Klinger. 2023. [Automatic emotion experiencer recognition](#). In *3rd Workshop on Computational Linguistics for the Political and Social Sciences (CPSS)*.
- Maximilian Wegge, Enrica Troiano, Laura Ana Maria Oberlaender, and Roman Klinger. 2022. [Experiencer-specific emotion and appraisal prediction](#). In *Proceedings of the Fifth Workshop on Natural Language Processing and Computational Social Science (NLP+CSS)*, pages 25–32, Abu Dhabi, UAE. Association for Computational Linguistics.
- Janyce Wiebe, Theresa Wilson, Rebecca Bruce, Matthew Bell, and Melanie Martin. 2004. [Learning subjective language](#). *Computational Linguistics*, 30(3):277–308.
- Rui Xia and Zixiang Ding. 2019. [Emotion-cause pair extraction: A new task to emotion analysis in texts](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1003–1012, Florence, Italy. Association for Computational Linguistics.
- Peng Zhou, Wei Shi, Jun Tian, Zhenyu Qi, Bingchen Li, Hongwei Hao, and Bo Xu. 2016. [Attention-based bidirectional long short-term memory networks for relation classification](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 207–212, Berlin, Germany. Association for Computational Linguistics.