

Introducing the NLP task of negative attitudinal function identification

Introduciendo la tarea de PLN de identificación de funciones actitudinales negativas

Nicolás José Fernández-Martínez

Departamento de Filología Inglesa, Universidad de Jaén
njfernan@ujaen.es

Abstract: On social media, users often express emotions, judgments, and evaluations on various social and private topics, detectable through automated methods. While NLP tasks like emotion detection and dialogue act classification focus on identifying emotions and intentions in texts, little attention has been paid to the attitudinal function of a text, such as expressing dislike, disagreement, pessimism, disapproval, etc. Our main contribution introduces the NLP task of negative attitudinal function identification, going beyond emotion detection and dialogue classification by focusing on users' intent and the expression of negative emotions, and negative ethical and aesthetic evaluations. We present a basic synthetic dataset for negative attitudinal functions built with foreign language teaching and learning resources. The dataset was used to develop negative attitudinal function models with supervised approaches, which were compared against other baseline models based on social media emotion detection datasets whose emotion categories were mapped to negative attitudinal functions. Our models, though not consistently outperforming baselines due to the qualitative differences of the tasks, use of out-of-domain data, and labeling issues found in the emotion detection datasets, exhibit promising capabilities with unseen data and in multilingual contexts.

Keywords: negative attitudinal function identification, NLP task, social media, synthetic dataset.

Resumen: En las redes sociales, los usuarios expresan con frecuencia sus emociones, juicios y evaluaciones sobre diversos temas sociales y privados, detectables mediante métodos automatizados. Mientras que tareas de PLN como la detección de emociones y la clasificación de actos de diálogo se centran en identificar emociones e intenciones en los textos, se ha prestado poca atención a la función actitudinal de un texto, como expresar desagrado, desacuerdo, pesimismo, desaprobación, etc. Nuestra principal contribución introduce la tarea de PLN de identificación de funciones actitudinales negativas, yendo más allá de la detección de emociones y la clasificación de diálogos al centrarse en la intención de los usuarios y la expresión de emociones negativas y evaluaciones éticas y estéticas negativas. Presentamos un dataset sintético básico para funciones actitudinales negativas construido con recursos obtenidos del campo de la enseñanza y aprendizaje de lenguas extranjeras. El conjunto de datos se utilizó para desarrollar modelos supervisados de funciones actitudinales negativas, que se comparó con otros modelos estándar basados en datasets de detección de emociones de redes sociales cuyas categorías de emociones fueron reasignadas a funciones actitudinales negativas. Nuestros modelos, aunque no superan sistemáticamente los modelos estándar debido a las diferencias cualitativas de las tareas, el uso de datos fuera de dominio y los problemas de etiquetado encontrados en los datasets de detección de emociones, muestran capacidades prometedoras con datos nunca antes vistos y en contextos multilingües.

Palabras clave: identificación de funciones actitudinales negativas, tarea de PLN, redes sociales, dataset sintético.

1 Introduction

On social media platforms such as Twitter (now called X), which has an ever-growing active user base of around 450 million monthly users, users post a staggering 500 million tweets a day on a wide range of public and private issues (Ruby, 2023).¹ This wealth of digital data can be leveraged with Natural Language Processing (NLP) tasks that focus on detecting users’ subjective emotions (i.e. emotion detection) (Cambria, 2016) or intent (i.e. dialog act classification) (Jurafsky et al., 1998). These NLP tasks provide manifold practical applications across diverse domains, from business and commerce to politics, disaster management, sociology, and digital humanities (Mohammad, 2021; Jurafsky and Martin, 2023).

Despite advances in understanding emotions and intent, a critical gap in existing NLP tasks remains unaddressed — the identification of the attitudinal function(s) of a given message. In linguistic terms, an attitudinal function reflects the speaker’s intention to communicate emotions, judgments, and appreciations about a topic, entity, or thing. This is particularly relevant in the context of smart cities, where predicting citizens’ dissatisfaction and concerns plays a pivotal role in enhancing their well-being and improving city services and infrastructure (Perián-Pascual, 2023). The focus on negative attitudinal functions in smart city contexts results from the research conducted in the ALLEGRO project (Perián-Pascual, 2023) where smart city problems are identified. Negative attitudinal functions take center stage in smart city contexts, reflecting citizens’ dissatisfaction with various aspects of their lives, from the condition of streets and parks to broader sociological issues such as economic inequality, racism, sexism, and political concerns. These functions encompass a spectrum of semantic subtleties, including expressions of dislike, disagreement, indifference, anger, threats, worries, distrust, pessimism, and more.

To bridge this gap, we reuse the concept of attitudinal function from linguistic theory. Drawing inspiration from functional theories of foreign language teaching and learning, systemic functional linguistics, and speech

act theory, we introduce the novel NLP task of negative attitudinal function identification. This task extends beyond emotion detection and dialog act classification, with a linguistically informed set of categories for dealing with the expression of intent with respect to emotions, judgments, and evaluations. We developed a basic synthetic dataset that captures negative attitudinal functions in smart city scenarios using existing linguistic resources. This dataset contains prototypical lexico-grammatical patterns that are formally realized by functions.² Leveraging this dataset, various supervised models were developed, including fine-tuned Transformers, contextualized sentence embeddings, and zero-shot classification with Natural Language Inference (NLI) using Transformers, together with traditional Machine Learning using a Naïve Bayes model.

Our study explores the models’ efficacy in identifying negative attitudinal functions in multilingual contexts, as well as the potential reuse of emotion detection datasets for this task. Our research seeks to contribute to the evolving landscape of NLP by addressing nuanced aspects of user expressions that were not captured in emotion detection and dialog act classification. This manuscript is organized as follows. Section 2 provides the background on emotion detection, dialog act classification, and negative attitudinal function identification. Section 3 describes the methodology used in our experiment, including the development of the synthetic dataset of negative attitudinal functions, the automatic mapping of social media emotion detection datasets, and the supervised approaches. Section 4 gives the results and discussion, including limitations, challenges, and future research directions. Section 5 presents the conclusion.

2 Background

2.1 Emotion detection

Emotions are integral to human life, influencing communication, interaction, and learning (Scherer, 2005). Studied historically in philosophy, particularly in Socratic schools like Aristotelianism and Stoicism (Sorabji, 2002),

¹<https://www.demandsage.com/twitter-statistics/>

²The dataset is available under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license at the following link: <https://github.com/njfm0001/Negative-Attitudinal-Function-Identification>

emotions are now being addressed in the psychological and computational literature (Picard, 1997; Scherer, 2005). Emotion detection in NLP is a sequence classification task (Cambria, 2016; Mohammad, 2021). Categorical models often rely on Ekman (1992)’s six universal emotions or Plutchik (1980)’s extended model. Applications span public health, safety, crisis response, advertising, and entertainment (Cambria, 2016; Mohammad, 2021). Users express emotions in text through emotionally loaded words, emojis, or implicit cultural cues (Mohammad, 2021). Important emotion markers include emotion verbs, adjectives, orthography, terms of address, adverbs, and exclamations (Mohammad and Alm, 2015; Alba-Juez, 2018). Emotions in text can refer to the writer, reader, or characters (Maia and Santos, 2018), although they typically correspond to the writer’s emotion (Buechel and Hahn, 2017). This task faces challenges like figurative language and cultural connotations (Mohammad, 2021). Efforts are now focusing on semantic role labeling in appraisal-based approaches (Campagnano, Conia, and Navigli, 2022; Wegge et al., 2022; Troiano, Oberländer, and Klinger, 2023).

Datasets for emotion detection come from various genres or domains, such as interviews (Scherer and Wallbott, 1994), news headlines (Strapparava and Mihalcea, 2007; Bostan, Kim, and Klinger, 2020), blog posts (Aman and Szpakowicz, 2007; Quan and Ren, 2009), tales (Alm, Roth, and Sproat, 2005), and, more recently, tweets (Mohammad, 2012; Roberts et al., 2012; Mohammad et al., 2015; Liew, Turtle, and Liddy, 2016; Mohammad et al., 2018; Plaza del Arco et al., 2020; Saravia et al., 2018) and Reddit comments (Demszky et al., 2020; Dwivedi-Yu and Halevy, 2022). Each domain has its own linguistic idiosyncrasies (Bostan and Klinger, 2018). Annotation methods include self-reported interviews (Scherer and Wallbott, 1994), distant supervision (Mohammad, 2012; Purver and Battersby, 2012; Wang et al., 2012; Dwivedi-Yu and Halevy, 2022), expert knowledge (Aman and Szpakowicz, 2007; Strapparava and Mihalcea, 2007), and crowd-sourcing (Mohammad and Alm, 2015; Liew, Turtle, and Liddy, 2016; Schuff et al., 2017; Demszky et al., 2020; Plaza del Arco et al., 2020). Approaches to emotion detection include symbolic models with lexica and rule-

based methods (Strapparava and Mihalcea, 2007; Dini and Bittar, 2016; Semeraro et al., 2023), and probabilistic models using machine learning or deep learning (Aman and Szpakowicz, 2007; Mohammad, 2012; Liew and Turtle, 2016). Probabilistic, deep learning models use networks like bidirectional LSTM or Transformers such as BERT (Devlin et al., 2019) or XLM-Twitter (Barbieri, Anke, and Camacho-Collados, 2022) via fine-tuning (Demszky et al., 2020; Vera, Araque, and Iglesias, 2021; Dwivedi-Yu and Halevy, 2022; Aroyehun et al., 2023) or through zero-shot classifiers (Basile, Pérez-Torró, and Franco-Salvador, 2021; Plaza del Arco, Martín-Valdivia, and Klinger, 2022; Yang et al., 2023).

2.2 Dialog act classification

Dialog act classification involves identifying users’ communicative intent in texts, typically dialogs (Searle, 1976; Jurafsky and Martin, 2023). Various annotation schemes like DAMSL (Core and Allen, 1997), SWBDD (Jurafsky, Shriberg, and Biasca, 1997), SPAAC (Leech and Weisser, 2003), DiAML (Bunt, 2009), DART (Weisser, 2018), and MIDAS (Yu and Yu, 2021) differ in dimensions, domains, and segmentation levels. DAMSL, the first scheme, is multidimensional, while SWBDD simplifies it for tagging dialog units. DiAML introduces DIT++ for annotation, encompassing clusters of general-purpose and dimension-specific functions. DART offers a fine-grained classification of speech acts. MIDAS is a multi-label scheme for human-machine interaction. These schemes apply to multi-turn texts but not tweets or other genres (Bunt, 2019).

Tagging pragmatic phenomena related to user intent faces challenges because of the complexity of the tagging schemes, requiring expert knowledge and labor-intensive efforts (Weisser, 2018; Yu et al., 2023). Datasets like Switchboard Dialog Act Corpus (Jurafsky, Shriberg, and Biasca, 1997), DailyDialog (Li et al., 2017), and DialogBank (Bunt et al., 2016) use various tagging schemes. Social media datasets, especially Twitter (now X), present difficulties due to their idiosyncratic nature (Saha et al., 2020; Baldwin et al., 2013). Symbolic approaches use lexico-syntactic cues like performative verbs and punctuation marks (Jurafsky et al., 1998).

Probabilistic models include Hidden Markov Models, Support Vector Machines, or Logistic Regression (Stolcke et al., 2000; Zhang, Gao, and Li, 2011; Vosoughi and Roy, 2016), and deep learning models include LSTM networks and Transformers (Khanpour, Gun-takandla, and Nielsen, 2016; Cerisara et al., 2018; Enayet and Sukthankar, 2020; Saha et al., 2020; Żelasko, Pappagari, and Dehak, 2021; Gung et al., 2023; Ostyakova et al., 2023a).

2.3 Negative attitudinal function identification

The appraisal framework in systemic functional linguistics explores three semantic domains related to the interpersonal metafunction: attitude, engagement, and graduation (Martin and White, 2005; Bednarek, 2008). Our focus is on attitude, which comprises affect (emotion) and opinion (judgment and appreciation). Affect pertains to emotional reactions, judgment involves ethical evaluations, and appreciation centers on aesthetic evaluations. In speech act theory (Searle, 1976), expressive and emotive speech acts are relevant to attitudes, dealing with psychological states (Sbisà, 1975; Guiraud et al., 2011; Zhabotynska and Slyvka, 2020). Com-misive and directive speech acts also carry attitudinal overtones (Sbisà, 1975). Existing speech act tagging schemes such as DiAML and DART offer some insights into expressing attitudes through speech acts. However, finer-grained distinctions are lacking.

To broaden the understanding of attitudinal functions, we also draw on the concept of communicative function from foreign language teaching and learning studies (Finocchiaro and Brumfit, 1983; Milanovic and Saville, 2012). Communicative functions serve specific purposes, including expressing attitudes. We developed a taxonomy of negative attitudinal functions for smart city scenarios, considering linguistic insights from all the previous approaches and function categorizations (Blundell, Higgins, and Middlemiss, 1982; Wilkins, 1976). The reason to focus on negative attitudinal functions is because, in smart city contexts, attitudes are often expressed in relation to negative events or entities, reflecting citizens’ feelings and ethical and aesthetic evaluations. For example, people might communicate their feelings towards inflation (affect), passing judgments

on economic policymakers for their decisions (judgment), or evaluating the effectiveness of economic policies independently of the policymakers’ actions (appreciation).

Our approach to attitudinal functions is a formal and discourse-pragmatic one, encompassing affect and evaluation through judgments and appreciations, and involves developing a dataset of attitudinal constructions that captures attitudinal meanings through lexico-grammatical patterns or constructions.

2.3.1 The task of negative attitudinal function identification

We define negative attitudinal function identification as a sequence classification task, where a given word sequence s is assigned one or more negative attitudinal function labels f from a set $f_1, f_2, f_3, \dots$, making it a single-label or multi-label task.

While linguistic constructions in the taxonomy serve as the primary cue for function identification, implicit expressions may also exist through other means (e.g. emojis, hashtags). For instance, the lexico-syntactic cue *disagree with something* clearly indicates the attitudinal function DISAGREE. However, in complex microtexts, multiple functions may be implied, requiring attitudinal function models to discern nuances. Defining functions by form poses challenges, as expressions like *I really hate* - may refer to multiple functions simultaneously (i.e. DISLIKE and ANGER). In social media microtexts, determining attitudinal functions becomes intricate due to potential implicit markers and polysemy.

Our synthetic dataset serves as a starting point for annotators or contextually rich models, like Transformers, to recognize the subtleties in expressing attitudinal functions, considering the complexities of context and implicit markers.

2.3.2 Differences and similarities with respect to emotion detection and dialog act classification

All these tasks involve sequence or text classification, assigning predefined categories to given texts, but they differ in aims, scope, and categories.

In emotion detection, the goal is to infer the writer’s, reader’s, and/or characters’ emotional states (Picard, 1997) using, for

instance, categorical models based on motivational theories (Ekman, 1992; Plutchik, 1980). The focus is thus on a set of universal emotions, employing psychological categorizations. Ekman (1992)’s categorization comprises the following emotion categories: anger, disgust, fear, joy, sadness, surprise. Plutchik (1980)’s tagset adds anticipation and trust.

Dialog act classification aims to identify users’ intent in utterances (usually in dialogs) using dialog acts (Jurafsky and Martin, 2023), loosely based on speech acts (Searle, 1976). Widely used dialog act tagsets include statements, questions, suggestions, comments, and miscellanea (Zhang, Gao, and Li, 2011). While most dialog acts lack attitudinal meanings, some taxonomies like DIT++ (Bunt, 2009) and DART (Weisser, 2018) include some attitudinal aspects that have not been explored in practice.

Negative attitudinal function identification detects users’ negative attitudinal intent, addressing emotional states, reactions, evaluative meanings (judgments and appreciations), and subjective states beyond emotional meanings. It bridges the intent and emotion focus of dialog act classification and emotion detection, encompassing linguistic nuances like distinctions between emotion, judgment, and appreciation. Unlike emotion detection, it adopts a linguistically grounded approach to cover a broader range of attitudinal meanings on the basis of categorizations found in theoretical and applied linguistic theory.

3 Methods

3.1 Development of the synthetic dataset of negative attitudinal functions

In our experimental setup, we built a synthetic dataset of linguistic constructions labeled with negative attitudinal functions, drawing from Blundell, Higgens, and Middlemiss (1982)’s comprehensive list. Blundell, Higgens, and Middlemiss (1982) is a practical textbook of functions for teaching and learning English as a Foreign Language that contains communicative functions of the following types: informational, attitudinal, and active. There are 12 informational functions, related to seeking or providing information; 48 attitudinal functions, in which

an attitude is expressed towards something (e.g. feelings, opinions, judgments); and 32 active functions, to establish courses of actions. Other function categories are given related to social formulas, communication strategies, and metalinguistic questions. This list was largely based on earlier categorizations in the functional-notional approach to foreign language teaching (Wilkins, 1976; Finocchiaro and Brumfit, 1983). Focusing on those attitudinal functions with negative overtones, we manually selected categories that were most relevant to smart cities and social media, where users express negative attitudes. Categories include PESSIMISTIC, WORRIED, ANGRY, DISAPPOINTED, BORED, DISLIKE, NOT_APPROVE, NOT_IMPORTANT, NOT_INTERESTED, DISAGREE, NOT_CORRECT, WARN, COMPLAIN, THREATEN, UNWILLING, REFUSE. Additionally, we introduced DISTRUST, using the Collins dictionary and the MacMillan dictionary thesauri, because this category was deemed crucial for smart city contexts, and the OTHER category, containing constructions from functions of different types found in Blundell, Higgens, and Middlemiss (1982). This results in 18 attitudinal functions with 362 constructions (Table 1).

Label	No.
OTHER	79
NOT_CORRECT	8
PESSIMISTIC	18
WORRIED	16
ANGRY	23
DISAPPOINTED	8
BORED	18
DISLIKE	22
NOT_APPROVE	19
NOT_IMPORTANT	20
NOT_INTERESTED	17
DISAGREE	23
WARN	10
COMPLAIN	21
THREATEN	13
UNWILLING	13
REFUSE	19
DISTRUST	15
Total	362

Table 1: Taxonomy of negative attitudinal functions.

Label	Examples
OTHER	My pleasure. Don't worry (about _).
NOT_CORRECT	You're/He's/She's/That's/We're/They're (all) wrong. (That's/It's) nonsense/rubbish/bullshit/bs/crap.
PESSIMISTIC	(I'm) not (too) happy (about _). There's no way.
WORRIED	I fear _. (I'm) (very) worried/uneasy (about _).
ANGRY	What an idiot/fool. _ (really) makes me mad.
DISAPPOINTED	That's/It's a real shame/pity/let-down. What a pity/disappointment.
BORED	_ is a (total) bore/drag. _ leaves me cold.
DISLIKE	(I) don't like _. How awful.
NOT_APPROVE	(I) don't think that's/it's (very) good. (I'm) dead against _.
NOT_IMPORTANT	(I) don't think that's (so) important. Does _ matter?
NOT_INTERESTED	(I'm) not (very) interested (in _). (I) couldn't care less (about _).
DISAGREE	(I) don't agree (with _). (I) can't go along (with _).
WARN	Watch out (for _). Make sure you don't do _.
COMPLAIN	_ really is the limit! (I'm) not at all satisfied (with _).
THREATEN	If I were you, I wouldn't do _. Don't do that or I'll do _.
UNWILLING	(I) don't (really) fancy doing _. I'd rather not (do _).
REFUSE	(I'm) sorry, I can't/couldn't (do _). Out of the question.
DISTRUST	Are you kidding? You must be joking.

Table 2: Examples of negative attitudinal function constructions.

The dataset construction is linguistically and sociologically motivated: it is informed by expert linguistic knowledge and sociological insights into citizens' problems in smart city contexts (Periñán-Pascual, 2023). Many constructions were taken literally from Blundell, Higgins, and Middlemiss (1982), whereas others were adapted to fit the linguistic nature of the social media domain. The constructions follow methodological conventions involving parentheses, underscores, and slashes, for optional elements, unspecified topics, and alternative expressions, respectively. Rules apply to the 362 constructions to obtain the full synthetic dataset with 902 samples. For example, for the COMPLAINT category, the constructions (*I'm*) *not at all satisfied* (*with* _), after the application of the rules, generate the sam-

ples *not at all satisfied*, *not at all satisfied with it*, *I'm not at all satisfied*, and *I'm not at all satisfied with it*. Another example: for the NOT_CORRECT category, the construction (*That's/It's*) *nonsense/rubbish/bullshit/bs/crap*. generates the samples *nonsense*, *rubbish*, *bullshit*, *bs*, *crap*, *That's nonsense*, *That's rubbish*, *That's bullshit*, *That's bs*, *That's crap*, *It's nonsense*, *It's rubbish*, *It's bullshit*, *It's bs*, and *It's crap*. Some other examples can be found in Table 2.

3.2 Automatic mapping of social media emotion detection datasets

Our experiment utilized state-of-the-art datasets from emotion detection tasks, as some attitudinal functions involve explicit or

implicit emotion expressions. Specifically, we selected social media datasets due to their relevance in expressing attitudes of negative type. These include EmoEvent (Plaza del Arco et al., 2020), GoEmotions (Demszky et al., 2020), CARER (Saravia et al., 2018), and the AIT dataset (Mohammad et al., 2018). For multilingual datasets (e.g. EmoEvent with English and Spanish tweets or AIT with English, Spanish, and Arabic tweets), we merged them while respecting the original splits and without pre-processing. Each dataset is briefly described below:

- **EmoEvent:** it comprises 8,409 English and 7,303 Spanish tweets labeled with a single emotion category, topic, and offensiveness (Plaza del Arco et al., 2020). The emotion set aligns with Ekman (1992) (i.e. anger, disgust, fear, joy, sadness, surprise), including the 'other' category.
- **GoEmotions:** it consists of around 54,000 English Reddit comments annotated with 27 fine-grained emotions plus a neutral category (Demszky et al., 2020). The dataset is highly imbalanced.
- **CARER:** A dataset of tweets expressing emotions, loosely based on Ekman (1992)'s classification (i.e. anger, fear, joy, love, sadness, surprise) (Saravia et al., 2018). Built via distant supervision using emotion-related hashtags.
- **AIT dataset:** Used for SemEval-2018 Task 1, it includes English, Spanish, and Arabic tweets (Mohammad and Kiritchenko, 2018). Emotion categories are based on Plutchik (1980)'s classification: anger, anticipation, disgust, fear, joy, love, optimism, pessimism, sadness, surprise, and trust.

The mapping of emotions to negative attitudinal functions considered the different annotation schemes. The mapping process is detailed in Table 3. The negative attitudinal function identification task was evaluated as either single-label or multi-label depending on the tagging scheme of the dataset.

3.3 Experimental setup: supervised approaches

The experiment was conducted employing several supervised approaches, includ-

Dataset	Mapping
EmoEvent	ANGRY: anger DISLIKE: disgust WORRIED: fear PESSIMISTIC: sadness OTHER: joy, surprise, other
GoEmotions	OTHER: admiration, amusement, approval, caring, confusion, curiosity, desire, embarrassment, excitement, gratitude, joy, love, optimism, pride, realization, relief, surprise, neutral ANGRY: anger DISLIKE: annoyance, disgust DISAPPOINTED: disappointment NOT_APPROVE: disapproval WORRIED: fear PESSIMISTIC: grief, remorse, sadness
CARER	ANGRY: anger WORRIED: fear PESSIMISTIC: sadness OTHER: joy, love, surprise
AIT	ANGRY: anger DISLIKE: disgust WORRIED: fear PESSIMISTIC: pessimism, sadness OTHER: anticipation, joy, love, optimism, surprise, trust

Table 3: Mapped attitudinal function categories in the emotion detection datasets.

ing traditional machine learning, fine-tuned Transformers, negative attitudinal function metaembeddings, and zero-shot classification through Natural Language Inference (NLI). Baseline models were developed, using these approaches, with the train splits of each mapped emotion detection datasets. Negative attitudinal function models were developed, using these approaches, with the synthetic dataset. The evaluation was conducted with the test splits of each mapped emotion detection dataset. A brief explanation is given for each approach:

- **Traditional Machine Learning with Multinomial Naïve Bayes (NB):** We trained a Multinomial Naïve Bayes model using bag-of-words features for text classification tasks.
- **Fine-tuning Transformers (T-XLM-R):** We utilized XLM-T (Bambler, Anke, and Camacho-Collados,

2022), a variant of XLM-RoBERTa (Conneau et al., 2020) pre-trained with tweets in over 30 languages. For single-label tasks, we fine-tuned the model with a softmax classifier, while for multi-label tasks, we used binary cross-entropy loss or a softmax classifier with an adapted loss function.

- **Negative attitudinal function metaembeddings:** Contextualized embeddings were obtained using Paraphrase Multilingual MpNet Base v2 for semantic similarity tasks (Reimers and Gurevych, 2020). Mean pooling was applied to compute attitudinal function metaembeddings for each negative attitude function category on the basis of the samples of the synthetic dataset, and cosine similarity determined the semantic similarity between these metaembeddings and the sentence embeddings obtained from the tweets.
- **Zero-shot classification with NLI:** We employed mDeBERTa v3 fine-tuned on NLI tasks for zero-shot classification (He, Gao, and Chen, 2021). Prompt engineering involved 13 prompts focusing on emotional states, text emotions, functions, or intentions expressed in the text (Table 4). The candidate labels were the emotion categories, then mapped to functions.

Prompts
<i>This person feels ..</i>
<i>This person conveys ..</i>
<i>This person shows ..</i>
<i>This person expresses ..</i>
<i>This text is ..</i>
<i>This text is about ..</i>
<i>This text shows ..</i>
<i>This text expresses ..</i>
<i>This text conveys ..</i>
<i>The communicative function of this text is ..</i>
<i>The communicative intention of this text is ..</i>
<i>The emotion of this text is ..</i>
<i>The emotion expressed in this text is ..</i>

Table 4: Prompts used in our zero-shot classification approach with NLI.

4 Results and discussion

Table 5 offers the results of each experimental setup for each dataset. In bold, the best

scores achieved for each dataset are highlighted.

Baseline models often provided the best results due to the use of in-domain data from emotion detection datasets. On the other hand, the function models were expected to behave worse, due to their evaluation with out-of-domain data and the qualitative differences between the tasks of emotion detection and negative attitudinal identification. Despite that, we highlight the robust performance of the metaembedding function model with the CARER and AIT datasets, as they are loaded with many explicit emotion expressions, approaching the fine-tuned baseline. However, the function models displayed subpar performance with EmoEvent and GoEmotions, possibly due to dataset annotation issues impacting performance, as seen in the scores obtained by the baseline models and our error analysis stage. The best-performing prompt in the zero-shot NLI approach was *The emotion expressed in this text is {label}.*, except for GoEmotions, which was *This text is {label}.*

Notably, our function models, developed with the negative attitudinal constructions in English, showcased remarkable generalization to tweets in Spanish and Arabic, as shown by the results obtained in both the fine-tuning and metaembedding approaches of the function models in the AIT dataset. This may suggest the universal applicability of emotional expressions encoded in some function constructions, potentially obviating the need to build function datasets in other languages.

In summary, our models using the synthetic dataset of negative attitudinal functions exhibit promising capabilities in emotion detection for unseen data and multilingual contexts.

4.1 Limitations, challenges, and future research directions

Our evaluation of function datasets against emotion detection baselines faced challenges due to the qualitative differences of the tasks and annotation issues found in some emotion detection datasets. Mapping emotions to functions was not consistently equivalent, impacting fairness. Annotation issues in the EmoEvent and GoEmotions datasets affected the performance of all models. CARER and AIT, with higher quality annotations, led to

Dataset	Model	Evaluation metrics		
		Macro-F1	Micro-F1	Weighted-F1
EmoEvent en-es	NB-baseline	0.25	0.70	0.66
	NB-basic	0.16	0.38	0.46
	T-XLM-R-baseline	0.38	0.70	0.69
	T-XLM-R-basic	0.18	0.74	0.65
	Metaembedding-baseline	0.27	0.49	0.56
	Metaembedding-basic	0.25	0.56	0.59
	Zero-shot NLI	0.26	0.45	0.52
GoEmotions	NB-baseline	0.12	0.75	0.64
	NB-basic	0.12	0.38	0.42
	T-XLM-R-baseline	0.52	0.78	0.78
	T-XLM-R-basic	0.18	0.28	0.66
	Metaembedding-baseline	0.21	0.31	0.53
	Metaembedding-basic	0.23	0.36	0.49
	Zero-shot NLI	0.23	0.15	0.09
CARER	NB-baseline	0.77	0.84	0.82
	NB-basic	0.26	0.38	0.34
	T-XLM-R-baseline	0.95	0.96	0.96
	T-XLM-R-basic	0.23	0.48	0.33
	Metaembedding-baseline	0.62	0.67	0.67
	Metaembedding-basic	0.45	0.52	0.51
	Zero-shot NLI	0.58	0.64	0.65
AIT en-es-ar	NB-baseline	0.53	0.60	0.58
	NB-basic	0.16	0.19	0.18
	T-XLM-R-baseline	0.75	0.77	0.77
	T-XLM-R-basic	0.43	0.49	0.48
	Metaembedding-baseline	0.48	0.49	0.53
	Metaembedding-basic	0.50	0.54	0.54
	Zero-shot NLI	0.54	0.55	0.54

Table 5: Results of the experiments.

better model performance.

Improving emotion dataset annotation quality through manual supervision could enhance model performance. Another line of research could focus on manually tagging the emotion detection datasets with functions, using expert annotation and/or crowdsourcing. Future research could also focus on synthetic data generation (Dai et al., 2023) using large language models (LLMs) like ChatGPT or BARD, which can also be used as automatic annotators (Ostyakova et al., 2023b; Kaddour et al., 2023). We could also build tweet datasets with semi-automatic methods by leveraging the lexicogrammatical patterns found in our taxonomy and use expert annotation for revising the samples obtained. Human annotated and synthetically generated data could be compared and combined in future experiments. Future work could also employ semantic role labeling with frames (Fillmore, 2006; Baker,

Fillmore, and Lowe, 1998) to enhance attitudinal function conceptualization and granularity. Another line of research could focus on expanding and/or creating new function taxonomies for diverse purposes, contexts, and languages.

4.2 Ethical considerations

The increasing use of affect-related data by governments and corporations raises ethical concerns, particularly in the context of digital surveillance. Users willingly disclose extensive personal information on social media (Han, 2015), leading to potential misuse, such as predicting and manipulating users’ behavior for advertising and political purposes (McStay, 2020). This ‘psychological targeting’ involves building psychological profiles from digital footprints for manipulative ends (Matz, Appel, and Kosinski, 2020). The ethical implications include the infringement of freedom and harm to individuals’

interests. Proposed solutions involve limiting exposure, raising public awareness, and enforcing regulations (McStay, 2020). Ethical decision-making for affect-related systems should consider task design, data building, and annotation processes (Mohammad, 2022). The synthetic dataset was created to avoid privacy issues concerning the collection of user-generated data. Access to our dataset should prioritize social good and ethical considerations. Our negative attitudinal function identification task aims to enhance users' well-being in smart city contexts, not manipulate users' behavior.

5 Conclusion

We introduce the task of negative attitudinal function identification, which seeks to discern users' emotional reactions, judgments, and appreciations. It holds potential for smart city scenarios, where citizens' expressions of emotions, judgements, and appreciations can inform policies for social good, addressing concerns and enhancing well-being. Negative attitudinal functions enable citizens to voice complaints about various issues. Unlike emotion detection and dialog act classification, this task draws on resources from and insights into theoretical and applied linguistic research, offering a more nuanced understanding of attitudinal intent. Beyond blending the interests of both emotion detection and dialog act classification, our task widens its scope to include subjective states that do not necessarily carry emotional meaning, such as judgments and appreciations.

We constructed a linguistically informed synthetic dataset of negative attitudinal functions that contained lexico-grammatical patterns. This dataset was then used to develop different supervised approaches. Emotion detection datasets of tweets were reused for our task by mapping their emotion categories to the functions of our taxonomy. They were then used to develop different supervised approaches. This was the baseline. An evaluation stage was conducted to compare the performance of the function models against the baseline. Results revealed promising capabilities of our function models, despite challenges such as qualitative differences between the tasks, the use of out-of-domain data, and annotation noise. Despite these challenges, negative attitudinal function models demonstrate promising po-

tential, particularly in multilingual contexts, and with out-of-domain data. We hope that NLP practitioners and researchers can benefit from this new NLP task and the associated synthetic dataset of negative attitudinal functions.

Acknowledgements

Financial support for this research has been provided under grant PID2020-112827GB-I00, funded by the Spanish Ministry of Science and Innovation MCIN/AEI/10.13039/501100011033, and grant number 101017861[project SMART-LAGOON], funded by the European Union's Horizon 2020 research and innovation program.

References

- Alba-Juez, L. 2018. Emotion and appraisal processes in language. In M. d. l. Á. Gómez González and J. L. Mackenzie, editors, *The Construction of Discourse as Verbal Interaction*. John Benjamins, pages 227–250.
- Alm, C. O., D. Roth, and R. Sproat. 2005. Emotions from text: Machine learning for text-based emotion prediction. In *HLT/EMNLP 2005 - Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, number October, pages 579–586.
- Aman, S. and S. Szpakowicz. 2007. Identifying expressions of emotion in text. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 4629 LNAI, pages 196–205. Springer Berlin Heidelberg.
- Aroyehun, S. T., L. Malik, H. Metzler, N. Haimerl, A. Di Natale, and D. Garcia. 2023. LEIA: Linguistic Embeddings for the Identification of Affect. *CoRR*, cs.CL/2304:1–15, apr.
- Baker, C. F., C. J. Fillmore, and J. B. Lowe. 1998. The Berkeley FrameNet Project. In *COLING 1998 Volume 1: The 17th International Conference on Computational Linguistics*.
- Baldwin, T., P. Cook, M. Lui, A. MacKinlay, and L. Wang. 2013. How Noisy Social Media Text, How Diffrent Social Me-

- dia Sources? *International Joint Conference on Natural Language Processing*, (October):356–364.
- Barbieri, F., L. E. Anke, and J. Camacho-Collados. 2022. XLM-T: Multilingual Language Models in Twitter for Sentiment Analysis and Beyond. In *2022 Language Resources and Evaluation Conference, LREC 2022*, number June, pages 258–266, Marseille, France. European Language Resources Association.
- Basile, A., G. Pérez-Torró, and M. Franco-Salvador. 2021. Probabilistic Ensembles of Zero- and Few-Shot Learning Models for Emotion Classification. *International Conference Recent Advances in Natural Language Processing, RANLP*, pages 128–137.
- Bednarek, M. 2008. *Emotion Talk across Corpora*. Palgrave Macmillan.
- Blundell, J., J. Higgins, and N. Middlemiss. 1982. *Function in English*. Oxford University Press.
- Bostan, L., E. Kim, and R. Klinger. 2020. Good news everyone: A corpus of news headlines annotated with emotions, semantic roles, and reader perception. In *LREC 2020 - 12th International Conference on Language Resources and Evaluation, Conference Proceedings*, number May, pages 1554–1566, Marseille, France. European Language Resources Association.
- Bostan, L.-A.-M. and R. 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.
- Buechel, S. and U. Hahn. 2017. Readers vs. writers vs. texts: Coping with different perspectives of text understanding in emotion annotation. In *LAW 2017 - 11th Linguistic Annotation Workshop, Proceedings of the Workshop*, pages 1–12. Association for Computational Linguistics.
- Bunt, H. 2009. The DIT++ taxonomy for functional dialogue markup. *Proceedings of EDAML/AAMAS Workshop “Towards a Standard Markup Language for Embodied Dialogue Acts”*, pages 13–24.
- Bunt, H. 2019. Guidelines for using ISO standard 24617-2. Technical report.
- Bunt, H., V. Petukhova, A. Malchanau, A. Fang, and K. Wijnhoven. 2016. The DialogBank. In *Proceedings of the 10th International Conference on Language Resources and Evaluation, LREC 2016*, pages 3151–3158, Portorož, Slovenia. European Language Resources Association (ELRA).
- Cambria, E. 2016. Affective Computing and Sentiment Analysis. *IEEE Intelligent Systems*, 31(2):102–107.
- Campagnano, C., S. Conia, and R. 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)*, volume 1, pages 4586–4601, Dublin, Ireland. Association for Computational Linguistics.
- Cerisara, C., S. Jafaritazehjani, A. Oluokun, and H. Le. 2018. Multi-task dialog act and sentiment recognition on Mastodon. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 745–754, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Conneau, A., K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, and V. Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Core, M. G. and J. F. Allen. 1997. Coding Dialogs with the DAMSL Annotation Scheme. In *Proceedings of the 1997 AAAI Fall Symposium: Communicative Action in Humans and Machines*, pages 28–35.
- Dai, H., Z. Liu, W. Liao, X. Huang, Z. Wu, L. Zhao, W. Liu, N. Liu, S. Li, D. Zhu, H. Cai, Q. Li, D. Shen, T. Liu, and X. Li. 2023. ChatAug: Leveraging ChatGPT for Text Data Augmentation. pages 1–12.
- Demszky, D., D. Movshovitz-Attias, J. Ko, A. Cowen, G. Nemade, and S. 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.
- Devlin, J., M. W. Chang, K. Lee, and K. Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, volume 1, pages 4171–4186.
- Dini, L. and A. Bittar. 2016. Emotion analysis on twitter: The hidden challenge. In *Proceedings of the 10th International Conference on Language Resources and Evaluation, LREC 2016*, pages 3953–3958, Portorož, Slovenia. European Language Resources Association (ELRA).
- Dwivedi-Yu, J. and A. Y. Halevy. 2022. "That's so cute!": The CARE Dataset for Affective Response Detection. In *Proceedings of the 26th Conference on Computational Natural Language Learning (CoNLL)*, pages 50–69. Association for Computational Linguistics.
- Ekman, P. 1992. An argument for basic emotions. *Cognition and Emotion*, 6(3-4):169–200, may.
- Enayet, A. and G. Sukthankar. 2020. A Transfer Learning Approach for Dialogue Act Classification of GitHub Issue Comments.
- Fillmore, C. J. 2006. Frame semantics. In D. Geeraerts, editor, *Cognitive linguistics: Basic readings*, volume 34. Mouton de Gruyter, Berlin, New York, pages 373–400.
- Finocchiaro, M. and C. Brumfit. 1983. *The functional-notional approach: from theory to practice*. Oxford University Press, New York.
- Guiraud, N., D. Longin, E. Lorini, S. Pesty, and J. Rivière. 2011. The face of emotions: A logical formalization of expressive speech acts. In Y. Tumer, P. Yolum, L. Sonenberg, and P. Stone, editors, *10th International Conference on Autonomous Agents and Multiagent Systems 2011 (AAMAS 2011)*, volume 2, pages 969–976, Taipei, Taiwan. International Foundation for Autonomous Agents and Multiagent Systems.
- Gung, J., R. Shu, E. Moeng, W. Rose, S. Romeo, A. Gupta, S. Mansour, and Y. Zhang. 2023. Intent Induction from Conversations for Task-Oriented Dialogue at DSTC 11. *CoRR*, cs.CL/2304.
- Han, B.-C. 2015. *The Transparency Society*. Stanford University Press, Stanford.
- He, P., J. Gao, and W. Chen. 2021. DeBERTaV3: Improving DeBERTa using ELECTRA-Style Pre-Training with Gradient-Disentangled Embedding Sharing. *CoRR*, cs.CL/2111:1–17.
- Jurafsky, D. and J. H. Martin. 2023. Chatbots Dialogue Systems. In *Speech and Language Processing*. Book in preparation, chapter 15.
- Jurafsky, D., E. Shriberg, and D. Biscia. 1997. Switchboard SWBD DAMSL Shallow-Discourse-Function Annotation. Technical report.
- Jurafsky, D., E. Shriberg, B. Fox, and T. Curl. 1998. Lexical, prosodic, and syntactic cues for dialog acts. *ACL/COLING Workshop on Discourse Relations and Discourse Markers*, pages 114–120.
- Kaddour, J., J. Harris, M. Mozes, H. Bradley, R. Raileanu, and R. McHardy. 2023. Challenges and Applications of Large Language Models. (1).
- Khanpour, H., N. Guntakandla, and R. Nielsen. 2016. Dialogue act classification in domain-independent conversations using a deep recurrent neural network. In *COLING 2016 - 26th International Conference on Computational Linguistics, Proceedings of COLING 2016: Technical Papers*, pages 2012–2021.
- Leech, G. and M. Weisser. 2003. Generic Speech Act Annotation for Task-Oriented Dialogue. In D. Archer, P. Rayson, A. Wilson, and A. McEnery, editors, *Proceedings of the Corpus Linguistics 2003 Conference*, pages 441–446, Lancaster. UCREL Technical Papers.
- Li, Y., H. Su, X. Shen, W. Li, Z. Cao, and S. Niu. 2017. DailyDialog: A Manually Labelled Multi-turn Dialogue Dataset. In *Proceedings of the The 8th International*

- Joint Conference on Natural Language Processing*, pages 986–995, Taipei, Taiwan. AFNLP.
- Liew, J. S. Y. and H. R. Turtle. 2016. Exploring Fine-Grained Emotion Detection in Tweets. In *Proceedings of NAACL-HLT*, pages 73–80, San Diego, California. Association for Computational Linguistics.
- Liew, J. S. Y., H. R. Turtle, and E. 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).
- Maia, B. and D. Santos. 2018. Language, emotion, and the emotions: The multidisciplinary and linguistic background. *Language and Linguistics Compass*, 12(6):1–17.
- Martin, J. R. and P. R. R. White. 2005. *The Language of Evaluation: Appraisal in English*. Palgrave Macmillan.
- Matz, S. C., R. E. Appel, and M. Kosinski. 2020. Privacy in the age of psychological targeting. *Current Opinion in Psychology*, 31:116–121.
- McStay, A. 2020. Emotional AI, soft biometrics and the surveillance of emotional life: An unusual consensus on privacy. *Big Data and Society*, 7(1).
- Milanovic, M. and N. Saville. 2012. The Theoretical Foundations for functions in the Council of Europe modern languages projects and the Common European Framework of Reference for languages. In *Language Functions Revisited*. Cambridge University Press.
- Mohammad, S. M. 2012. Emotional tweets. In **SEM 2012 - 1st Joint Conference on Lexical and Computational Semantics*, volume 1, pages 246–255, Montréal, Canada. Association for Computational Linguistics.
- Mohammad, S. M. 2021. Sentiment Analysis: Detecting Valence, Emotions, and Other Affectual States from Text. In *Emotion Measurement*. Elsevier, chapter 11, pages 201–237.
- Mohammad, S. M. 2022. Ethics Sheet for Automatic Emotion Recognition and Sentiment Analysis. *Computational Linguistics*, 48(2):239–278.
- Mohammad, S. M. and C. O. Alm. 2015. Computational Analysis of Affect and Emotion in Language. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing: Tutorial Abstracts*, Lisbon, Portugal. Association for Computational Linguistics.
- Mohammad, S. M., F. Bravo-Marquez, M. Salameh, and S. Kiritchenko. 2018. SemEval-2018 Task 1: Affect in Tweets. In *Proceedings of the 12th International Workshop on Semantic Evaluation (SemEval-2018)*, pages 1–17, New Orleans, Louisiana. Association for Computational Linguistics.
- Mohammad, S. M. and S. Kiritchenko. 2018. Understanding emotions: A dataset of tweets to study interactions between affect categories. In *LREC 2018 - 11th International Conference on Language Resources and Evaluation*, pages 198–209.
- Mohammad, S. M., X. Zhu, S. Kiritchenko, and J. Martin. 2015. Sentiment, emotion, purpose, and style in electoral tweets. *Information Processing and Management*, 51(4):480–499.
- Ostyakova, L., K. Petukhova, V. Smilga, and D. Zharikova. 2023a. Linguistic Annotation Generation with ChatGPT: a Synthetic Dataset of Speech Functions for Discourse Annotation of Casual Conversations ChatGPT : . In *Computational Linguistics and Intellectual Technologies: Proceedings of the International Conference “Dialogue 2023”*, pages 1–18.
- Ostyakova, L., V. Smilga, K. Petukhova, M. Molchanova, and D. Kornev. 2023b. ChatGPT vs. Crowdsourcing vs. Experts : Annotating Open-Domain Conversations with Speech Functions. In *Proceedings of the 24th Meeting of the Special Interest Group on Discourse and Dialogue*, pages 242–254. Association for Computational Linguistics.
- Periñán-Pascual, C. 2023. From Smart City to Smart Society: A quality-of-life ontological model for problem detection from user-generated content. *Applied Ontology*, pages 1–44.

- Picard, R. W. 1997. *Affective computing*. MIT Press, Cambridge, MA, US.
- Plaza del Arco, F. M., M.-T. Martín-Valdivia, and R. Klinger. 2022. Natural Language Inference Prompts for Zero-shot Emotion Classification in Text across Corpora. *CoRR*, cs.CL/2209.
- Plaza del Arco, F. M., C. Strapparava, L. Alfonso Ureña-López, and M. Teresa Martín-Valdivia. 2020. EmoEvent: A multilingual emotion corpus based on different events. In *LREC 2020 - 12th International Conference on Language Resources and Evaluation, Conference Proceedings*, number May, pages 1492–1498, Marseille, France. European Language Resources Association.
- Plutchik, R. 1980. A general psychoevolutionary theory of emotion. In R. Plutchik and H. B. T. T. o. E. Kellerman, editors, *Emotion: Theory, research, and experience: Vol. 1. Theories of emotion*. Academic Press, chapter A GENERAL, pages 3–33.
- Purver, M. and S. Battersby. 2012. Experimenting with distant supervision for emotion classification. In *EACL 2012 - 13th Conference of the European Chapter of the Association for Computational Linguistics, Proceedings*, pages 482–491, Avignon, France. Association for Computational Linguistics.
- Quan, C. and F. Ren. 2009. Construction of a blog emotion corpus for Chinese emotional expression analysis. In *EMNLP 2009 - Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: A Meeting of SIGDAT, a Special Interest Group of ACL, Held in Conjunction with ACL-IJCNLP 2009*, number August, pages 1446–1454.
- Reimers, N. and I. Gurevych. 2020. Sentence-BERT: Sentence embeddings using siamese BERT-networks. In *EMNLP-IJCNLP 2019 - 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing, Proceedings of the Conference*, pages 3982–3992.
- Roberts, K., M. A. Roach, J. Johnson, J. Guthrie, and S. M. Harabagiu. 2012. EmpaTweet: Annotating and detecting emotions on twitter. In *Proceedings of the 8th International Conference on Language Resources and Evaluation, LREC 2012*, pages 3806–3813, Istanbul, Turkey. European Language Resources Association (ELRA).
- Ruby, D. 2023. 58+ Twitter Statistics For Marketers In 2023 (Users Trends).
- Saha, T., S. Ramesh Jayashree, S. Saha, and P. Bhattacharyya. 2020. BERT-Caps: A Transformer-Based Capsule Network for Tweet Act Classification. *IEEE Transactions on Computational Social Systems*, 7(5):1168–1179.
- Saravia, E., H. C. Toby Liu, Y. H. Huang, J. Wu, and Y. S. Chen. 2018. Carer: Contextualized affect representations for emotion recognition. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018*, pages 3687–3697.
- Sbisà, M. 1975. Speech acts and the expression of affect.
- Scherer, K. R. 2005. What are emotions? And how can they be measured? *Social Science Information*, 44(4):695–729, dec.
- Scherer, K. R. and H. G. Wallbott. 1994. Evidence for universality and cultural variation. *Journal of Personality and Social Psychology*, 66(2):310–328.
- Schuff, H., J. Barnes, J. Mohme, S. Padó, and R. 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.
- Searle, J. R. 1976. A classification of illocutionary Acts. *Language in Society*, 5(1):1–24.
- Semeraro, A., M. Stella, A. Semeraro, S. Vilella, S. M. Mohammad, G. Ruffo, and M. Stella. 2023. EmoAtlas: An emotional profiling tool merging psychological lexicons, artificial intelligence and network science.

- Sorabji, R. 2002. *Emotion and Peace of Mind: From Stoic Agitation to Christian Temptation*. Oxford University Press.
- Stolcke, A., K. Ries, N. Coccaro, E. Shriberg, R. Bates, D. Jurafsky, P. Taylor, R. Martin, C. Van Ess-Dykema, and M. Meteer. 2000. Dialogue Act Modeling for Automatic Tagging and Recognition of Conversational Speech. *Computational Linguistics*, 26(3):35.
- Strapparava, C. and R. Mihalcea. 2007. Learning to identify emotions in text. In *Proceedings of the ACM Symposium on Applied Computing*, pages 1556–1560.
- Troiano, E., L. Oberländer, and R. Klinger. 2023. Dimensional Modeling of Emotions in Text with Appraisal Theories: Corpus Creation, Annotation Reliability, and Prediction. *Computational Linguistics*, 49(1):1–72.
- Vera, D., O. Araque, and C. A. Iglesias. 2021. GSI-UPM at IberLEF2021: Emotion analysis of Spanish tweets by fine-tuning the XLM-RoBERTa language model. In *CEUR Workshop Proceedings*, volume 2943, pages 16–26.
- Vosoughi, S. and D. Roy. 2016. Tweet acts: A speech act classifier for Twitter. In *Proceedings of the 10th International Conference on Web and Social Media, ICWSM 2016*, pages 711–714.
- Wang, W., L. Chen, K. Thirunarayan, and A. P. Sheth. 2012. Harnessing twitter ‘big data’ for automatic emotion identification. In *Proceedings - 2012 ASE/IEEE International Conference on Privacy, Security, Risk and Trust and 2012 ASE/IEEE International Conference on Social Computing, SocialCom/PASSAT 2012*, pages 587–592.
- Wegge, M., E. Troiano, L. Oberländer, and R. Klinger. 2022. Experienter-Specific Emotion and Appraisal Prediction. In *Proceedings of the Fifth Workshop on Natural Language Processing and Computational Social Science (NLP+CSS)*, pages 25–32. Association for Computational Linguistics.
- Weisser, M. 2018. *How to Do Corpus Pragmatics on Pragmatically Annotated Data*. John Benjamins Publishing Company, Amsterdam/Philadelphia.
- Wilkins, D. A. 1976. *Notional Syllabuses*. Oxford University Press.
- Yang, K., S. Ji, T. Zhang, Q. Xie, and S. Ananiadou. 2023. On the Evaluations of ChatGPT and Emotion-enhanced Prompting for Mental Health Analysis. *CoRR*, cs.CL/2304.
- Yu, D., L. Li, H. Su, and M. Fuoli. 2023. Using LLM-assisted Annotation for Corpus Linguistics A Case Study of Local Grammar Analysis. *CoRR*, cs.CL/2305.
- Yu, D. and Z. Yu. 2021. MIDAS: A dialog act annotation scheme for open domain human machine spoken conversations. In *EACL 2021 - 16th Conference of the European Chapter of the Association for Computational Linguistics, Proceedings of the Conference*, pages 1103–1120. Association for Computational Linguistics.
- Żelasko, P., R. Pappagari, and N. Dehak. 2021. What helps transformers recognize conversational structure? Importance of context, punctuation, and labels in dialog act recognition. *Transactions of the Association for Computational Linguistics*, 9:1179–1195.
- Zhabotynska, S. and N. Slyvka. 2020. Emotive speech acts and their discourse modifications in the literary text. *Discourse and Interaction*, 13(1):113–136.
- Zhang, R., D. Gao, and W. Li. 2011. What are tweeters doing: Recognizing speech acts in twitter. In *AAAI Workshop - Technical Report*, volume WS-11-05, pages 86–91.