

Emotion Stimulus Detection in German News Headlines

Bao Minh Doan Dang, Laura Oberländer, and Roman Klinger

Institut für Maschinelle Sprachverarbeitung, University of Stuttgart, Germany

st117194@stud.uni-stuttgart.de,

{laura.oberlaender, roman.klinger}@ims.uni-stuttgart.de

Abstract

Emotion stimulus extraction is a fine-grained subtask of emotion analysis that focuses on identifying the description of the cause behind an emotion expression from a text passage (e.g., in the sentence “I am happy that I passed my exam” the phrase “passed my exam” corresponds to the stimulus.). Previous work mainly focused on Mandarin and English, with no resources or models for German. We fill this research gap by developing a corpus of 2006 German news headlines annotated with emotions and 811 instances with annotations of stimulus phrases. Given that such corpus creation efforts are time-consuming and expensive, we additionally work on an approach for projecting the existing English GoodNewsEveryone (GNE) corpus to a machine-translated German version. We compare the performance of a conditional random field (CRF) model (trained monolingually on German and cross-lingually via projection) with a multilingual XLM-RoBERTa (XLM-R) model. Our results show that training with the German corpus achieves higher F1 scores than projection. Experiments with XLM-R outperform their respective CRF counterparts.

1 Introduction

Emotions are a complex phenomenon that play a central role in our experiences and daily communications. Understanding them cannot be accounted by any single area of study since they can be represented and expressed in different ways, e.g., via facial expressions, voice, language, or gestures. In natural language processing, most models build on top of one out of three approaches to study and understand emotions, namely basic emotions (Ekman, 1992; Strapparava and Mihalcea, 2007; Aman and Szpakowicz, 2007), the valence-arousal model (Russell, 1980; Buechel and Hahn, 2017) or cognitive appraisal theory (Scherer, 2005; Hof-

mann et al., 2020, 2021). Emotion classification in text has received abundant attention in natural language processing research in the past few years. Hence, many studies have been conducted to investigate emotions on social media (Stieglitz and Dang-Xuan, 2013; Brynielsson et al., 2014; Tromp and Pechenizkiy, 2015), in literary and poetry texts (Kim and Klinger, 2019; Haider et al., 2020) or for analysing song lyrics (Mihalcea and Strapparava, 2012; Hijra Ferdinan et al., 2018; Edmonds and Sedoc, 2021). However, previous work mostly focused on assigning emotions to sentences or text passages. These approaches do not allow to identify which event, object, or person caused the emotion (which we refer to as the *stimulus*).

Emotion stimulus detection is the subtask of emotion analysis which aims at extracting the stimulus of an expressed emotion. For instance, in the following example from FrameNet (Fillmore et al., 2003) “Holmes is happy having the freedom of the house when we are out” one could assume that *happiness* or *joy* is the emotion in the text. One could also highlight that the term “happy” indicates the emotion, “Holmes” is the experiencer and the phrase “having the freedom of the house when we are out” (underlined) is the stimulus for the perceived emotion. Detecting emotion stimuli provides additional information for a better understanding of the emotion structures (e.g., semantic frames associated with emotions). More than that, the fact that stimuli are essential in understanding the emotion evoked in a text is supported by research in psychology; Appraisal theorists of emotions seem to agree that emotions include a cognitive evaluative component of an event (Scherer, 2005). Therefore emotion stimulus detection brings the field of emotion analysis in NLP closer to the state of the art in psychology.

To the best of our knowledge, there are mostly corpora published for Mandarin (Lee et al., 2010b;

Gui et al., 2014, 2016; Gao et al., 2017) and English (Ghazi et al., 2015; Mohammad et al., 2014; Kim and Klinger, 2018; Bostan et al., 2020). We are not aware of any study that created resources or models for identifying emotion stimuli in German. We fill this gap and contribute the GERSTI (GERman STImulus) corpus with 2006 German news headlines. The headlines have been annotated for emotion categories, for the mention of an experiencer or a cue phrase, and for stimuli on the token level (on which we focus in this paper). News headlines have been selected as the domain because they concisely provide concrete information and are easy to obtain. Additionally, unlike social media texts, this genre avoids potential privacy issues (Bostan et al., 2020). Given that annotating such a corpus is time-consuming, we propose a heuristic method for projecting an annotated dataset from a source language to a target language. This helps to increase the amount of training data without manually annotating a huge dataset. Within this study, the GoodNewsEveryone corpus (GNE, Bostan et al., 2020) is selected as an English counterpart.

Our contributions are therefore: (1) the creation, publication, and linguistic analysis of the GERSTI dataset to understand the structure of German stimulus mentions;¹ (2), the evaluation of baseline models using different combinations of feature sets; and (3) comparison of this in-corpus training with cross-lingual training via projection and with a pre-trained cross-lingual language model with XLM-RoBERTa (Conneau et al., 2020).

2 Related Work

We now introduce previous work on emotion analysis and for detecting emotion stimuli.

2.1 Emotion Analysis

Emotion analysis is the task of understanding emotions in text, typically based on psychological theories of Ekman (1992), Plutchik (2001), Russell (1980) or Scherer (2005). Several corpora have been built for emotion classification such as Alm and Sproat (2005) with tales, Strapparava and Mihalcea (2007) with news headlines, Aman and Szpakowicz (2007) with blog posts, Buechel and Hahn (2017) with various domains or Li et al. (2017) with conversations. Some datasets were cre-

ated using crowdsourcing, for instance Mohammad et al. (2014), Mohammad and Kiritchenko (2015) or Bostan et al. (2020), that have been annotated with tweets, or news headlines, respectively. Some resources mix various annotation paradigms, for example Troiano et al. (2019) (self-reporting and crowd-sourcing) or Haider et al. (2020) (experts and crowdworkers).

Emotion analysis also includes other aspects such as emotion intensities and emotion roles (Aman and Szpakowicz, 2007; Mohammad and Bravo-Marquez, 2017; Bostan et al., 2020) including experiencers, targets, and stimuli (Mohammad et al., 2014; Kim and Klinger, 2018).

2.2 Stimulus Detection

Emotion stimulus detection received substantial attention for Chinese Mandarin (Lee et al., 2010b; Li and Xu, 2014; Gui et al., 2014, 2016; Cheng et al., 2017, i.a.). Only few corpora have been created for English (Neviarouskaya and Aono, 2013; Mohammad et al., 2014; Kim and Klinger, 2018; Bostan et al., 2020). Russo et al. (2011) worked on a dataset for Italian news texts and Yada et al. (2017) annotated Japanese sentences from news articles and question/answer websites.

Lee et al. (2010b,a) developed linguistic rules to extract emotion stimuli. A follow-up study developed a machine learning model that combines different sets of such rules (Chen et al., 2010). Gui et al. (2014) extended these rules and machine learning models on their Weibo corpus. Ghazi et al. (2015) formulated the task as structured learning.

Most methods for stimulus detection have been evaluated on Mandarin. Gui et al. (2016) propose a convolution kernel-based learning method and train a classifier to extract emotion stimulus events on the clause level. Gui et al. (2017) treat emotion stimulus extraction as a question answering task. Li et al. (2018) use a co-attention neural network. Chen et al. (2018) explore a joint method for emotion classification and emotion stimulus detection in order to capture mutual benefits across these two tasks. Similarly, Xia et al. (2019) evaluate a hierarchical recurrent neural network transformer model to classify multiple clauses. They show that solving these subtasks jointly is beneficial for the model’s performance.

Xia and Ding (2019) redefine the task as emotion/cause pair extraction and intend to detect potential emotions and corresponding causes in text.

¹The data is available at <https://www.ims.uni-stuttgart.de/data/emotion>.

Xu et al. (2019) tackle the emotion/cause pair extraction task by adopting a learning-to-rank method. Wei et al. (2020) also argue for the use of a ranking approach. They rank each possible emotion/cause pair instead of solely ranking stimulus phrases. Fan et al. (2020) do not subdivide the emotion/cause pair detection task into two subtasks but propose a framework to detect emotions and their associated causes simultaneously.

Oberländer and Klinger (2020) studied whether sequence labeling or clause classification is appropriate for extracting English stimuli. As we assume that these findings also hold for German, we follow their finding that token sequence labeling is more appropriate.

3 Corpus Creation

To tackle German emotion stimulus detection on the token-level, we select headlines from various online news portals, remove duplicates and irrelevant items, and further subselect relevant instances with an emotion dictionary. Two annotators then label the data. We describe this process in detail in the following.

3.1 Data Collection

We select various German news sources and their RSS feeds based on listings at a news overview website² and add some regional online newspapers.³ The collected corpus consists of headlines between September 30, 2020 and October 7, 2020 and between October 22 and October 23, 2020 with 9000 headlines, spread across several domains including *politics*, *sports*, *tech and business*, *science* and *travel*.

3.2 Data Preprocessing and Filtering

Short headlines, for instance “Verlobung!” or “Krasser After-Baby-Body” do not contain sufficient information for our annotation, therefore we omit sentences that have less than 5 words. Further, we remove generic parts of the headline, like “++ Transferticker ++”, “+++ LIVE +++” or “News-” and only keep the actual headline texts.

We also remove headlines that start with particular key words which denote a specific event which would not contribute to an understanding of

²<https://www.deutschland.de/de/topic/wissen/nachrichten>, accessed on April 27, 2021

³The list of RSS feeds is available in the supplemental material.

No.	Linguistics Rules
1.	Stimuli can be described by verbal or nominal phrases
2.	Subjunctions like “because of” belong to the sequence
3.	Conjunctions like “and”, “or” and “but” connect main clauses. They can therefore belong to a stimulus sequence.
4.	Antecedents, if present, are annotated as stimuli
5.	If antecedent is not present, an anaphora may be annotated instead
6.	Composites with “-” are considered a single word
7.	Stimuli can include one or multiple words
8.	Punctuation (e.g. ,;-;“”!?) should not be labeled as stimulus

Table 1: Linguistics rules for annotating stimuli.

emotions or stimuli, such as “Interview”, “Kommentare”, “Liveblog”, “Exklusive”, as well as visual content like “Video”, “TV” or “Pop”. Additionally, we discard instances which include dates, like “Lotto am Mittwoch, 30.09.2020” or “Corona-News am 05.10”.⁴

After filtering, we select instances that are likely to be associated with an emotion with the help of an emotion lexicon (Klinger et al., 2016). For this purpose, we accept headlines which include at least one entry from the dictionary.

3.3 Annotation

The annotation of the 2006 headlines which remain after preprocessing and filtering consists of two phases. In the first phase, emotion cues, experiencers and emotion classes are annotated, while stimuli are addressed in the second phase only for those instances which received an emotion label. Table 8 in the Appendix shows the questions to be answered during this annotation procedure. Each headline in the dataset is judged by two annotators. One of them is female (23 years old) while the other annotator is male (26 years old). The first annotator has a background in digital humanities and linguistics, while the second has a background in library and information management. After each phase, we combine overlapping stimulus annotations by choosing the parts annotated by both annotators, and discuss the cases where the annotations do not overlap until a consensus is reached.

Guidelines. We created an initial version of guidelines motivated by Lee et al. (2010b,a); Gui et al. (2014); Ghazi et al. (2015). Based on two batches of 25 headlines, and one with 50 headlines,

⁴Details in Supplementary Material.

Iteration	κ			F ₁		
				tok.	span	
	Cue	Exp.	Emo.	Stim.		
Prelim. 1	.22	.43	.25	—	—	—
Prelim. 2	.71	.49	.47	—	—	—
Prelim. 3	.46	.69	.44	—	.65	—
Final	.56	.57	.51	.68	.72	.56

Table 2: Inter-annotator agreement for the binary tasks of annotating the existence of cue mentions, experiencer mentions, the multi-label annotation of emotion labels, and the token-level annotation of stimulus spans. The F₁-span value for stimuli is an exact match value for the whole span.

we refined the guidelines in three iterations. After each iteration, we calculated inter-annotator agreement scores and discussed the annotator’s results. It should be noted that we only considered annotating emotions in the first two iterations. The sample annotation of emotion stimuli on the token-level has been performed in the third round, i.e., after two discussions and guideline refinements. During these discussions, we improved the formulation of the annotation task, provided more detailed descriptions for each predefined emotion and clarified the concept of sequence labeling using the IOB scheme. Additionally, we formulated several linguistic rules that help annotating stimuli (see Table 1).

Details. The goal of Phase 1 of the annotation procedure is to identify headlines with an emotional connotation. Those which do then receive stimulus annotations in Phase 2.

We annotated in a spread sheet application. In Phase 1a both annotators received 2006 headlines. They were instructed to annotate whether a headline expresses an emotion by judging if cue words or experiencers are mentioned in the text. Further, only one, the most dominant, emotion is to be annotated (*happiness, sadness, fear, disgust, anger, positive surprise, negative surprise, shame, hope, other* and *no emotion*). In Phase 1b we aggregated emotion annotations and jointly discussed non-overlapping labels to a consensus annotation.

In Phase 2a, annotators were instructed to label pretokenized headlines with the IOB alphabet for stimulus spans – namely those which received an emotion label in Phase 1 (811 instances). In Phase 2b, we aggregated the stimulus span annotations to a gold standard by accepting all overlapping tokens of both annotators in cases where they partially




Figure 1: κ for all emotion classes.

matched. For the other cases where the stimulus annotations did not overlap, we discussed the annotations to reach an agreement.

Agreement Results. Table 2 presents the inter-annotator agreement scores for the preliminary annotation rounds and for the final corpus. We observe that the results are moderate across classes. Figure 1 illustrates the agreement for each emotion class. The emotions *anger, fear, and happiness* show the highest agreement, while *surprise, other, and particularly disgust* show lower scores.

For the stimulus annotation, we evaluate the agreement via token-level Cohen’s κ , via token-level F₁, and via exact span-match F₁ (in the first two cases, B and I labels are considered to be different). The token-level result for the final corpus is substantial with $\kappa = .68$, F₁ = .72 and moderate for the exact span match, with F₁ = .56 (see Table 2).

4 Corpus Analysis

4.1 Quantitative Analysis

Our corpus consists of 2006 headlines with 20,544 tokens and 6,763 unique terms. From those, 811 instances were labeled with an emotion category and received stimulus annotations on the token-level. The shortest headline consists of five words, while the longest has 20 words. The headlines are on average short with nine words. The stimulus spans range from one to eleven tokens and have four words on average.

Table 3 summarizes the corpus statistics of GER-STI. For aggregating emotion cue and experiencer we accept instances for which the mention of these emotion roles has been annotated by one annotator. For all emotions, most instances include the mention of an emotion cue (likely biased by our sam-

Emotion	# inst.	w/ cue	w/ exp	w/ stimulus	avg. stimulus
Happiness	80	80	77	76	3.72
Sadness	65	65	54	59	4.07
Fear	177	117	138	167	3.83
Disgust	3	3	2	3	4.00
Anger	226	226	195	208	3.86
Pos. Surprise	51	51	45	44	4.11
Neg. Surprise	142	140	125	130	3.96
Shame	9	9	9	8	3.75
Hope	20	19	16	19	4.05
Other	38	37	26	34	3.71
No Emo.	1195	930	109	-	-
All	2006	1737	796	748	3.9

Table 3: Corpus statistics. Columns show the amount of annotated instances for emotion cue, experiencer, stimulus and the average length of all stimulus spans within each respective dominant emotion. For aggregating cue and experiencer, cases where one of the annotators annotated with a *yes* have been accepted.

pling procedure). Further, the number of headlines with mentions of a stimulus and an experiencer is also high for those instances which are labeled to be associated with an emotion.

Table 4 presents the most common sources, sorted by their frequencies, for each aggregated emotion during Phase 1b. Not surprisingly, *Bild-Zeitung* is to be found in the top three for almost all emotion classes, followed by *Stuttgarter-Zeitung* and *Welt*. In particular, in five out of ten of the emotions, *Bild-Zeitung* takes the first place. As Table 3 demonstrates, *disgust* is relatively rare, we therefore list all available sources for this emotion category. Furthermore, four in five most frequently annotated emotions are negative (*anger*, *fear*, *negative surprise*, *happiness*, *sadness*).

Note that this analysis does not necessarily reflect the actual quality of chosen news sources. The findings we report here might strongly be biased by the data collection time span.

4.2 Qualitative Analysis of Stimuli

To obtain a better understanding of stimuli in German, we analyse which words together with their preferred grammatical realizations are likely to indicate a stimulus phrase. For this purpose, we

Emotion	News Sources
Happiness	Bild, Welt, Stuttgarter Zeitung
Sadness	Bild, Spiegel, Stuttgarter Z.
Fear	Stuttgarter Z., Bild, Welt
Disgust	T-Online, Welt, Spiegel
Anger	Bild, Stuttgarter Z., Spiegel
Pos. Surprise	Welt, Focus, Bild
Neg. Surprise	Bild, Stuttgarter Z., Spiegel
Shame	Stuttgarter Z., Bild, Welt
Hope	T-Online, Bild, Stuttgarter Z.
Other	Bild, Stuttgarter Z., Welt

Table 4: Top three most observed media sources for each dominant emotion sorted by frequency.

examine the parts of speech⁵ of terms that are directly left positioned to stimulus phrases, inside the stimulus phrases and right after it (see Table 5). We further compare our findings with Mandarin (Lee et al., 2010a) and English (Bostan et al., 2020).

Our analysis shows that for GERSTI common nouns, proper nouns, punctuation, and verbs are most frequently located directly to the left of stimulus mentions (common nouns $\approx 26\%$, punctuation $\approx 28\%$, verbs $\approx 22\%$, proper nouns $\approx 0.09\%$). Often, these words are emotionally connotated, for instance as in the nouns “Streit”, “Angst”, “Hoffnung” or “Kritik” or the verbs “warnen”, “kritisieren”, “bedrohen”, “beklagen” or “kämpfen”.

There are discrepancies between German and Mandarin stimuli. Lee et al. (2010a,b) state that prepositions or conjunctions mostly indicate stimulus phrases in Mandarin, while this is not the case for German due to our predefined annotation rules (Rule 2 from Table 1). Furthermore, indicator words for Chinese stimulus events do not cover common nouns or proper nouns. However, verbs seem to emphasize emotion causes in both languages.

Compared to GNE, we also notice some differences: English stimuli do not begin with prepositions, but prepositions are most likely to be included in the stimulus span ((ADP) $\approx 0.14\%$ in GNE vs $\approx 0.03\%$ in GERSTI). Further, by looking at the part of speech tags that were relevant in indicating the stimuli for GERSTI we see that they are dominating for GNE as well. However, there are far more proper nouns than common nouns and quite fewer verbs that occur right before the stimulus phrase (common nouns $\approx 11\%$, punctuation $\approx 21\%$, verbs $\approx 0.09\%$, proper nouns $\approx 0.25\%$).

⁵We use spaCy, <https://spacy.io/usage/linguistic-features>, accessed on April 29, 2021

POS	GERSTI				GNE			
	All	Inside	Before@1	After@1	All	Inside	Before@1	After@1
NOUN	.28	.33 (1.17×)	.26 (0.93×)	.00 (0.01×)	.16	.17 (1.09×)	.11 (0.69×)	.17 (1.05×)
ADP	.15	.22 (1.48×)	.03 (0.19×)	.23 (1.54×)	.10	.12 (1.12×)	.14 (1.37×)	.20 (1.95×)
PROPN	.14	.09 (0.65×)	.09 (0.68×)	.01 (0.04×)	.30	.26 (0.89×)	.25 (0.86×)	.25 (0.83×)
PUNCT	.13	.02 (0.16×)	.28 (2.23×)	.49 (3.87×)	.09	.07 (0.82×)	.21 (2.40×)	.08 (0.91×)
VERB	.09	.09 (0.91×)	.22 (2.32×)	.16 (1.68×)	.11	.12 (1.06×)	.09 (0.80×)	.09 (0.85×)
DET	.05	.08 (1.47×)	.00 (0.09×)	.01 (0.16×)	.04	.05 (1.03×)	.04 (0.81×)	.03 (0.63×)
ADJ	.05	.07 (1.44×)	.00 (0.03×)	.01 (0.29×)	.05	.05 (1.09×)	.02 (0.42×)	.03 (0.53×)
ADV	.05	.05 (1.04×)	.04 (0.87×)	.04 (0.93×)	.02	.02 (1.07×)	.02 (0.80×)	.03 (1.47×)
AUX	.02	.01 (0.75×)	.04 (2.34×)	.03 (1.68×)	.03	.03 (1.01×)	.03 (1.16×)	.03 (1.11×)
PRON	.01	.01 (0.71×)	.02 (1.02×)	.00 (0.19×)	.03	.03 (1.14×)	.01 (0.45×)	.02 (0.63×)
NUM	.01	.02 (1.49×)	.00 (0.00×)	.00 (0.00×)	.02	.02 (1.15×)	.01 (0.27×)	.01 (0.34×)
CCONJ	.01	.01 (0.97×)	.01 (0.55×)	.01 (0.77×)	.01	.01 (1.21×)	.00 (0.64×)	.02 (3.82×)

Table 5: Relative frequencies of POS tags of all tokens in GERSTI and GNE datasets (All) vs relative frequencies of POS tags inside the stimuli spans (Inside), before and after the stimuli spans (Before@1, After@1). For all the columns that show frequencies of the spans related to the stimuli we show the factor (×) of how much it differs to the global frequencies in All.

Often, these indicator words of English stimuli do not as directly evoke an emotion. For instance, “say”, “make”, “woman”, “people” or “police” are often observed to be directly left located words of English stimuli. Nevertheless, similar to GERSTI, stimuli from GNE corpus are not indicated by conjunctions, numerals or pronouns.

The positioning of the stimuli is only similar to a limited degree in German and English: 53% of the instances in GERSTI end with the stimulus (86% in English GNE) and 13% begin with the stimulus (11% in GNE).

5 Experiments

In the following, we explain how we project annotation from an English stimulus corpus to a machine-translated counterpart. Based on this, we evaluate how well a linear-chain conditional random field (Lafferty et al., 2001) performs with the projected dataset in comparison to the monolingual setup. We compare that result to the use of the pre-trained language model *XLM-RoBERTa* (XLM-R) (Conneau et al., 2020).

5.1 Annotation Projection

We use the GNE dataset (Bostan et al., 2020) which is a large English annotated corpus of news headlines. Stimulus sequences in this dataset are comparatively longer with eight tokens on average.

We translate the GNE corpus via *DeepL*⁶ and perform the annotation projection as follows: We first translate the whole source instance t_{en} to the

⁶<https://www.deepl.com/en/translator>, accessed on May 20, 2021

translation t_{de} (from English to German). We further translate the stimulus token sequence stim_{en} to stim_{de} . We assume the stimulus annotation for t_{de} to correspond to all tokens in stim_{de} , heuristically corrected to be a consecutive sequence.

5.2 Experimental Setting

5.2.1 Models

CRF. We implement the linear-chain conditional random field model via the CRF-suite in Scikit-learn⁷ and extract different features. What we call *corpus-based features* contains the frequency of a current word in the whole corpus, position label for first (*begin*), last (*end*) and remaining (*middle*) words of the headline, if the current word is capitalized, or entirely in upper or lowercase, if the token is a number, a punctuation symbol, or in the list of 50 most frequent words in our corpus.

We further include *linguistic features*, namely the part-of-speech tag, the syntactic dependency between the current token and its head, if it is a stopword or if it has a named entity label (and which one it is).

We further add a feature which specifies whether the token is part of an emotion-word dictionary (Klinger et al., 2016). Additionally, we combine the feature vector of the preceding and succeeding token (we add the prefixes *prev* and *next* to each feature name) with the current token to get information about surrounding words. We mark the first and last token with additional features.

⁷<https://sklearn-crfsuite.readthedocs.io/en/latest/>, accessed on April 30, 2021

RoBERTa. We use the pre-trained *XLM-RoBERTa* base model with the *HuggingFace*⁸ library from Wolf et al. (2020). In addition to the pre-trained transformer, we add a linear layer which outputs a sequence of IOB tags for each input sentence. We fine-tune the language model in five epochs and use a batch size of 16 during training, a dropout rate of 0.5, and the Adam optimizer with weight decay (Loshchilov and Hutter, 2019), with a learning rate of 10^{-5} and a maximum gradient norm of 1.0.

Setup. For our experiments, we only use the 811 instances from the GERSTI dataset that received annotations for emotion stimuli. We split them into a train and validation subset (80%/20%) and perform experiments in three different settings. In the *in-corporus training*, we train with the GERSTI training data and test on the test corpus. In the *projection* setting, we train on the english GNE data and test on the German GERSTI test data (either with the CRF via projection or directly with the XLM-R model). In the *aggregation* setting, we use both the English train data and the German train data for training.

5.2.2 Evaluation Metrics

We evaluate the stimuli prediction as follows (following Ghazi et al. (2015) and Oberländer and Klinger (2020)): *Exact* match leads to a true positive for an exactly correct span prediction. *Partial* accepts a predicted stimulus as true positive if at least one token overlaps with a gold standard span. A variation is *Left/Right*, where the left/right boundary needs to perfectly match the gold standard.

5.3 Results

Table 6 reports the results for our experiments. The top four blocks compare the importance of the feature set choice for the CRF approach.

In nearly all combinations of model and evaluation measure, the in-corporus evaluation leads to the best performance – adding data from the GNE corpus only slightly improves for the *Partially* evaluation setting when the CRF is limited to corpus features. The projection-based approach, where the model does not have access to the GERSTI training data consistently shows a lower performance, with approximately a drop by 50% in F_1 score.

⁸<https://huggingface.co/xlm-roberta-base>, accessed on April 30, 2021

Model	F_1	in-corp.	proj.	aggre.
CRF with corpus features	Exact	.38	.19	.33
	Partial	.49	.43	.52
	Left	.42	.22	.38
	Right	.51	.41	.51
CRF with linguistic features	Exact	.42	.16	.35
	Partial	.58	.41	.54
	Left	.52	.19	.43
	Right	.57	.40	.53
CRF with corp.+lingu. features	Exact	.45	.19	.35
	Partial	.57	.48	.53
	Left	.53	.24	.41
	Right	.56	.47	.52
CRF with all features	Exact	.42	.20	.36
	Partial	.56	.48	.55
	Left	.50	.25	.43
	Right	.55	.46	.53
RoBERTa XLM-R	Exact	.47	.25	.45
	Partial	.75	.61	.70
	Left	.68	.35	.58
	Right	.71	.59	.72

Table 6: Results for the CRF models with different feature sets and the XLM-R model. Highest F_1 -scores in each row printed with **bold face**, highest score in column/per evaluation measure is underlined, highest score in each column and per evaluation measure in the CRF is printed *italics*.

The linguistic features particularly help the CRF in the *Exact* evaluation setting, but all feature set choices are dominated by the results of the XLM-RoBERTa model. This deep learning approach shows the best results across all models, and is particularly better in the *Partial* evaluation setting, with 19pp, 13pp and 15pp improvement.

Both projection and aggregation models indicate that extracting the beginning of a stimulus span is challenging. We assume that both models have learned English stimulus structures and therefore could not generalize well on the German emotion stimuli (also see Section 4.2).

5.4 Error Analysis

We now discuss the model’s quality (see Table 7) based on various error types, namely *Early Start*, *Late Start*, *Early Stop*, *Late Stop*, *Surrounding* (*Early Start & Late stop*) and *Consecutive* error.

Both CRF and XLM-R with projection settings have largely generated *Early Start* and *Late Stop* errors. These models tend to detect longer stimulus segments than annotated in the gold data. This might be a consequence of English stimuli being longer than in German. Despite the fact that a CRF does not have an understanding of the length

Err. Type	Example	Setup
Early start	Hof in Bayern: 21-Jähriger [nach tödlichem Autounfall zu Court in Bavaria : 21-year-old [after deadly car-accident to Bewährungsstrafe verurteilt] probation convicted <i>Court in Bavaria: 21-year-old sentenced to probation after fatal car accident</i>	projection
Late start	Peter Madsen in Dänemark: Kim Walls Mörder [scheitert bei Peter Madsen in Denmark : Kim Wall's murderer fails by Fluchtversuch aus Gefängnis] escape-attempt from prison <i>Peter Madsen from Denmark: Kim Wall's killer fails in escape attempt from prison</i>	in-corporus
Early stop	Noch mehr Eltern erzählen [von den unheimlichen Dingen , die ihr Even more parents tell about the scary things , that their Kind mal gesagt hat] child once said has <i>More parents share creepy things their kid once said</i>	in-corporus
Late stop	In Paris: [Lauter Knall] schreckt Menschen auf - Ursache schnell In Paris : Loud bang scares people on - Cause quickly gefunden found <i>In Paris: Loud bang startles people - cause quickly found</i>	aggregation
Surrounding	EU-Gipfel: Streit [über Linie zur Türkei] - Erdogan reagiert mit Häme EU-summit : Dispute about line to Turkey - Erdogan reacts with gloat <i>EU-summit: Dispute over line on Turkey - Erdogan responds with gloating</i>	projection
Consecutive	Niederlage für Autohersteller: [Betriebsratswahl bei Daimler Defeat for car-manufacturer : work-council-election by Daimler ungültig] invalid <i>Defeat for car manufacturer: Daimler's work council election invalid</i>	aggregation

Table 7: Example headlines for examined error types. Gold annotations correspond to tokens between []. Predicted stimulus segments are highlighted as follows: red (B tag), blue (I tag). English translations for each sample are written in *italics*. All examples stem from the CRF models except the last one.

of span due to the Markov property, it has a bias weight for transitions between I labels. An example for such a case is the first instance from Table 7 the projection setting also extracted the token “21-Jähriger” as the start of the stimulus sequence. This explains the difference between partial and exact F₁ scores in Table 6.

The *Surrounding* exemplifies that the models tend to predict the beginning of a stimulus span directly after a colon. In contrast, in the in-corporus experiments (particularly with XLM-R), models tend to generate *Late Start* and *Early Stop* errors more often. For example the second headline from Table 7 shows a missing prediction of the verb “scheitert”. Instead, the preposition “bei” is found as the start of the stimulus phrase. Further, in the subse-

quent example, this model setting does not cover the phrase “die ihr Kind mal gesagt hat” in the stimulus segment. Both sample headlines demonstrate that in-corporus models tend to label prepositions as the start of stimulus sequences.

In the XML-R experiments, we opted against the use of a Viterbi-decoded output layer (like a CRF output) – this leads to errors of the *Consecutive* type, as shown in the last example: start and end of the stimulus are correctly found, but tokens in between have been missed.

6 Conclusion and Future Work

We introduced the first annotated German corpus for identifying emotion stimuli and provided baseline model results for various CRF configurations

and an XLM-R model. We additionally proposed a data projection method.

Our results show training and testing the model in the same language outperforms cross-lingual models. Further, the XLM-R model that uses a multilingual distributional semantic space outperforms the projection. However, based on partial matches, we see that, when approximate matches are sufficient projection and multilingual methods show an acceptable result.

Previous work has shown that the task of stimulus detection can be formulated as token sequence labeling or as clause classification (Oberländer and Klinger, 2020). In this paper we limited our analysis and modeling on the sequence labeling approach. Thus, we leave to future work the comparison with the clause-classification approach. However, from the results obtained, we find sequence labeling an adequate formulation in German.

For further future work, we suggest experimenting with the other existing corpora in English to examine whether the cross-lingual approach would work well on other domains. Regarding this, one could also train and improve models not only for language change but also to extract stimuli across different domains. Subsequently, another aspect that should be investigated is the simultaneous recognition of emotion categories and stimuli.

Acknowledgements

This work was supported by Deutsche Forschungsgemeinschaft (project CEAT, KL 2869/1-2). Thanks to Pavlos Musenidis for fruitful discussions and feedback on this study.

References

- Cecilia Ovesdotter Alm and Richard Sproat. 2005. [Emotional sequencing and development in fairy tales](#). In *Affective Computing and Intelligent Interaction*, 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. Springer Berlin Heidelberg.
- 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.
- Joel Brynielsson, Fredrik Johansson, Carl Jonsson, and Anders Westling. 2014. [Emotion classification of social media posts for estimating people’s reactions to communicated alert messages during crises](#). *Security Informatics*, 3(1):1–11.
- 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.
- Ying Chen, Wenjun Hou, Xiyao Cheng, and Shoushan Li. 2018. [Joint learning for emotion classification and emotion cause detection](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 646–651, Brussels, Belgium. Association for Computational Linguistics.
- Ying Chen, Sophia Yat Mei Lee, Shoushan Li, and Churen Huang. 2010. [Emotion cause detection with linguistic constructions](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 179–187, Beijing, China. Coling 2010 Organizing Committee.
- Xiyao Cheng, Ying Chen, Bixiao Cheng, Shoushan Li, and Guodong Zhou. 2017. [An emotion cause corpus for chinese microblogs with multiple-user structures](#). *ACM Transactions on Asian and Low-Resource Language Information Processing*, 17(1).
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin 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.
- Darren Edmonds and João Sedoc. 2021. [Multi-emotion classification for song lyrics](#). In *Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 221–235, Online. Association for Computational Linguistics.
- Paul Ekman. 1992. [An argument for basic emotions](#). *Cognition and Emotion*, 6(3-4):169–200.
- Chuang Fan, Chaofa Yuan, Jiachen Du, Lin Gui, Min Yang, and Ruifeng Xu. 2020. [Transition-based directed graph construction for emotion-cause pair extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3707–3717, Online. Association for Computational Linguistics.
- Charles J. Fillmore, Miriam R. L. Petruck, Josef Ruppenhofer, and Abby Wright. 2003. [Framenet in action: The case of attaching](#). *International Journal of Lexicography*, 16:297–332.

- Qinghong Gao, Hu Jiannan, Xu Ruifeng, Gui Lin, Yulan He, Kam-Fai Wong, and Quin Lu. 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. National Institute of Informatics Test Collection for Information Resources.
- Diman Ghazi, Diana Inkpen, and Stan Szpakowicz. 2015. [Detecting emotion stimuli in emotion-bearing sentences](#). In *Computational Linguistics and Intelligent Text Processing*, pages 152–165, Cham. Springer.
- Lin Gui, Jiannan Hu, Yulan He, Ruifeng Xu, Qin Lu, and Jiachen Du. 2017. [A question answering approach for emotion cause extraction](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1593–1602, Copenhagen, Denmark. Association for Computational Linguistics.
- Lin Gui, Dongyin Wu, Ruifeng Xu, Qin Lu, and Yu Zhou. 2016. [Event-driven emotion cause extraction with corpus construction](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1639–1649, Austin, Texas. Association for Computational Linguistics.
- Lin Gui, Li Yuan, Ruifeng Xu, Bin Liu, Qin Lu, and Yu Zhou. 2014. [Emotion cause detection with linguistic construction in chinese weibo text](#). In *Natural Language Processing and Chinese Computing*, pages 457–464, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Thomas Haider, Steffen Eger, Evgeny Kim, Roman Klinger, and Winfried Menninghaus. 2020. [PO-EMO: Conceptualization, annotation, and modeling of aesthetic emotions in German and English poetry](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1652–1663, Marseille, France. European Language Resources Association.
- Afif Hijra Ferdinan, Andrew Brian Osmond, and Casi Setianingsih. 2018. [Emotion classification in song lyrics using k-nearest neighbor method](#). In *2018 International Conference on Control, Electronics, Renewable Energy and Communications (ICCEREC)*, pages 63–69.
- 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.
- 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. 2019. [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.
- Roman Klinger, Surayya Samat Suliya, and Nils Reiter. 2016. [Automatic Emotion Detection for Quantitative Literary Studies – A case study based on Franz Kafka’s “Das Schloss” and “Amerika”](#). In *Digital Humanities 2016: Conference Abstracts*, pages 826–828, Kraków, Poland. Jagiellonian University and Pedagogical University.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. [Conditional random fields: Probabilistic models for segmenting and labeling sequence data](#). In *Proceedings of the Eighteenth International Conference on Machine Learning*, page 282–289, San Francisco, CA. Morgan Kaufmann Publishers Inc.
- Sophia Yat Mei Lee, Ying Chen, and Chu-Ren Huang. 2010a. [A text-driven rule-based system for emotion cause detection](#). In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 45–53, Los Angeles, CA. Association for Computational Linguistics.
- Sophia Yat Mei Lee, Ying Chen, Shoushan Li, and Chu-Ren Huang. 2010b. [Emotion cause events: Corpus construction and analysis](#). In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC’10)*, Valletta, Malta. European Language Resources Association (ELRA).
- Weiyuan Li and Hua Xu. 2014. [Text-based emotion classification using emotion cause extraction](#). *Expert Systems with Applications*, 41(4):1742–1749.
- Xiangju Li, Kaisong Song, Shi Feng, Daling Wang, and Yifei Zhang. 2018. [A co-attention neural network model for emotion cause analysis with emotional context awareness](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language*

- Processing*, pages 4752–4757, Brussels, Belgium. 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.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Rada Mihalcea and Carlo Strapparava. 2012. [Lyrics, music, and emotions](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 590–599, Jeju Island, Korea. Association for Computational Linguistics.
- Saif Mohammad and Felipe Bravo-Marquez. 2017. [Emotion intensities in tweets](#). In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM 2017)*, pages 65–77, Vancouver, Canada. Association for Computational Linguistics.
- Saif Mohammad and Svetlana Kiritchenko. 2015. [Using hashtags to capture fine emotion categories from tweets](#). *Computational Intelligence*, 31(2):301–326.
- 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.
- Alena Neviarouskaya and Masaki Aono. 2013. [Extracting causes of emotions from text](#). In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 932–936, Nagoya, Japan. Asian Federation of Natural Language Processing.
- Laura Ana Maria Oberländer and Roman Klinger. 2020. [Token sequence labeling vs. clause classification for English emotion stimulus detection](#). In *Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics*, pages 58–70, Barcelona, Spain (Online). Association for Computational Linguistics.
- 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.
- James A. Russell. 1980. [A circumplex model of affect](#). *Journal of personality and social psychology*, 39(6):1161–1178.
- Irene Russo, Tommaso Caselli, Francesco Rubino, Ester Boldrini, and Patricio Martínez-Barco. 2011. [EMOCause: An easy-adaptable approach to extract emotion cause contexts](#). In *Proceedings of the 2nd Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA 2.011)*, pages 153–160, Portland, Oregon. Association for Computational Linguistics.
- Klaus R. Scherer. 2005. [What are emotions? And how can they be measured?](#) *Social Science Information*, 44(4):695–729.
- Stefan Stieglitz and Linh Dang-Xuan. 2013. [Emotions and information diffusion in social media-sentiment of microblogs and sharing behavior](#). *Journal of management information systems*, 29(4):217–248.
- 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. Association for Computational Linguistics.
- 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.
- Erik Tromp and Mykola Pechenizkiy. 2015. [Pattern-based emotion classification on social media](#). In *Advances in social media analysis*, pages 1–20. Springer.
- Penghui Wei, Jiahao Zhao, and Wenji Mao. 2020. [Effective inter-clause modeling for end-to-end emotion-cause pair extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3171–3181, Online. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- 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.

- Rui Xia, Mengran Zhang, and Zixiang Ding. 2019. [Rthn: A rnn-transformer hierarchical network for emotion cause extraction](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI-19)*, pages 5285–5291, Macao. International Joint Conferences on Artificial Intelligence.
- Bo Xu, Hongfei Lin, Yuan Lin, Yufeng Diao, Liang Yang, and Kan Xu. 2019. [Extracting emotion causes using learning to rank methods from an information retrieval perspective](#). *IEEE Access*, 7:15573–15583.
- Shuntaro Yada, Kazushi Ikeda, Keiichiro Hoashi, and Kyo Kageura. 2017. [A bootstrap method for automatic rule acquisition on emotion cause extraction](#). In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 414–421, New Orleans, LA. Institute of Electrical and Electronics Engineers.

A Appendix

Question	Annotation	Labels
Phase 1: Emotion Annotation		
1. Are there terms in the headline which could indicate an emotion?	Cue word	0, 1
2. Does the text specify a person or entity experiencing an emotion?	Experiencer	0, 1
3. Which emotion is most provoked within the headline?	Emotion	Emotions
Phase 2: Stimuli		
4. Which token sequence describes the trigger event of an emotion?	Stimulus	BIO

Table 8: Questions for the annotation.