



LEIA: Linguistic Embeddings for the Identification of Affect

Segun Taofeek Aroyehun^{1,2}, Lukas Malik^{3,4}, Hannah Metzler^{5,3,2}, Nikolas Haimerl⁶, Anna Di Natale^{5,3,2} and David Garcia^{1,2,3,5*} 

*Correspondence:

david.garcia@uni-konstanz.de

¹Department of Politics and Public Administration, University of Konstanz, Konstanz, Germany

²Graz University of Technology, Graz, Austria

Full list of author information is available at the end of the article

Abstract

The wealth of text data generated by social media has enabled new kinds of analysis of emotions with language models. These models are often trained on small and costly datasets of text annotations produced by readers who guess the emotions expressed by others in social media posts. This affects the quality of emotion identification methods due to training data size limitations and noise in the production of labels used in model development. We present LEIA, a model for emotion identification in text that has been trained on a dataset of more than 6 million posts with self-annotated emotion labels for happiness, affection, sadness, anger, and fear. LEIA is based on a word masking method that enhances the learning of emotion words during model pre-training. LEIA achieves macro-F1 values of approximately 73 on three in-domain test datasets, outperforming other supervised and unsupervised methods in a strong benchmark that shows that LEIA generalizes across posts, users, and time periods. We further perform an out-of-domain evaluation on five different datasets of social media and other sources, showing LEIA's robust performance across media, data collection methods, and annotation schemes. Our results show that LEIA generalizes its classification of anger, happiness, and sadness beyond the domain it was trained on. LEIA can be applied in future research to provide better identification of emotions in text from the perspective of the writer.

Keywords: Emotion detection; Natural language processing; Social media; Transfer learning

1 Introduction

Automatic identification of emotion in text is a valuable tool to study affect through social media and other text digital traces [1]. Word-based methods enabled the study of mood expressions on Twitter [2] in relation to daylight oscillations [3] and of collective emotions in social resilience [4]. Rule-based methods allowed the quantification of emotion contagion on Twitter [5] and the dynamics of emotions after affect labeling on social media [6]. More advanced classification methods trained on labeled data in various languages have been used to test the effect of air pollution on happiness in Weibo posts [7], to study the expression of emotions on Twitter about Black Lives Matter [8], and to validate social media emotion macroscopes against survey data [9, 10]. Beyond research, emotion detection

© The Author(s) 2023. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

from social media text has clinical potential to identify users at mental health risk [11] and can help platforms to detect abusive language [12].

Despite its potential, the use of emotion detection from social media text faces important challenges. Dictionary methods applied to social media text provide user-level metrics that are weakly correlated with answers to affective questionnaires [13]. Furthermore, dictionary-based emotion analysis methods have weak correlations with population-level emotion prevalence [14], but the same study shows that more advanced supervised methods bear promise to capture well-being. One of the sources of problems with the application of social media text to study emotions is the sensitivity of methods to particular domains. For example, [15] applied out-of-the-box sentiment analysis in a benchmark of different domains and found how methods are very sensitive to the medium and text source. This is part of a general problem in which language model performance degrades with distribution shifts [16], weakening the validity of emotion detection from text in out-of-domain (OOD) settings.

A source of error in emotion detection in social media is the way in which training labels are produced. While the target of applications is often to infer a subjective emotional state of the author of a social media post, the labels of training data are frequently produced by readers and not the authors of the post. The use of crowdsourcing can contribute to this problem, which can be alleviated by gathering several annotations per text but always carrying the potential noise source of readers not understanding the emotional state of writers. For example, a comparison between reader and writer annotations shows that they disagree 25% of the time [17]. To avoid this problem, experience sampling can be used to generate self-annotated emotion labels. For example, [18] gathered anxiety scores at the time when individuals posted tweets and compared self-reported anxiety with emotion text analysis. The results are correlations of at most 0.24, calling for studies that can leverage large datasets to identify emotional states more accurately. One must note, however, that we cannot assume that self-reported emotion labels are perfectly predictable from social media text, with only the natural language processing models as the missing piece. The upper limit on the performance of an emotion identification method is likely to be below 100%, as for example, well-being indicators correlate with each other correlations of approximately 0.84, which can be achieved with modern language models [19].

New platforms to share emotional experiences with other users offer the possibility to gather large-scale datasets with emotion self-annotations. Vent is an example that offers a particularly good source of self-annotated data, as the dataset available for researchers has millions of posts [20] and the design of the platform is precisely to share emotions rather than a smaller functionality as in other platforms. Recent research on Vent has shown the difficulty to predict Vent precise mood labels from text [21], but it is still left to explore how Vent can be used to infer more coarse emotion labels that can match discrete emotion classes from psychological research. In this work, we focus on a subset of Vent tags that can be mapped to standard emotional states, with the goal of training a better and more robust emotion detection model that can be applied to other text sources, especially from other social media. In the following, we present the design and development of LEIA, followed by an empirical analysis in a benchmark of in-domain and out-of-domain tests. We further analyze examples of classification errors and outputs of LEIA to understand its limitations and paths for improvement.

2 Related work

Emotion classification models mainly follow feature-based or neural approaches. Feature-based methods [22] employ handcrafted features built from resources such as emotion lexica. Neural approaches often rely on pre-trained representations such as word embeddings and contextual language models (LMs). The use of transformer-based LMs has been shown to yield state-of-the-art performance on natural language processing benchmarks. For emotion classification, recent research works have achieved better performance using pre-trained LMs [23–25].

Learning representations for affect A number of existing works learn representations for affective tasks. DeepMoji [26] is a neural network trained for predicting emoji in tweets using a large distant-labeled dataset considering 64 emojis as labels. Sentiment-specific word embeddings [27] encode sentiment information into the vector representation of words for sentiment analysis. Sentiment-aware language representation learning (SentiLARE) [28] incorporates part-of-speech and word polarity to enhance representation learning of a contextual language model for sentiment analysis tasks. Another effective strategy in several natural language processing tasks is to pre-train transformer models on a large collection of text and then fine-tune the model for other downstream tasks [29], including tasks in the social media domain [23, 25]. In this strategy, the adaptation step often relies on the masked language modeling objective where random tokens are masked and the model is trained to predict the masked tokens. Alternative masking strategies have been proposed to improve the pre-training task either by masking important words [30] or masking words relevant for a given downstream task. Recently, emotion masked language modeling (eMLM) was proposed in [31] to preferentially mask emotion words for contextual language representation learning. Similar to SentiLARE, eMLM also relied on existing lexical resources by masking emotional words more frequently when training a Bidirectional Encoder Representations from Transformers (BERT) model from scratch, yielding improvements in downstream affect-related tasks. Motivated by these results, we employ eMLM in the design of LEIA as we explain below.

Fine-tuning strategies and model generalization Supervised models can show a performance drop when faced with domain shifts, i.e. when they are applied to text from a domain that is not the same as the domain of their training data [16]. A recent result in computer vision [32] showed that this performance gap across domains can be mitigated with a fine-tuning strategy that first performs linear probing to align the features of the prediction head with the pre-trained base model and then fine-tuning all model parameters. This approach is similar to those proposed in [33] and provides a further theoretical basis as well as empirical validation. Linear probing is a non-destructive and computationally cheap approach that freezes the parameters of the base model and only updates the parameters of the prediction head during training. In this work, we consider this strategy in the context of text classification for the identification of emotion.

Emotion classification datasets Supervised models are trained and evaluated against emotion text datasets that are either constructed by manual labeling or automatically by using additional data sources and structures. Manually-labeled datasets are usually comparatively small while automatically-constructed datasets are built by identifying

emotion-bearing patterns of expression such as hashtags in the case of Twitter. The annotation of emotion datasets can also be divided into reader-labeled and writer-labeled datasets. Reader-labeled datasets are assigned labels by the annotators post-hoc based on their perception of the emotions expressed by a given content. On the other hand, writer-labeled datasets are usually self-annotated by the writer of the message to reflect their emotion.

Most of the existing work on emotion classification has drawn on manually annotated, automatically constructed, and reader-labeled datasets. Recently, large-scale writer-labeled datasets have been introduced [20, 34] and they are yet to become part of the benchmarks of emotion detection tasks. A notable example is the Vent dataset [20], which is produced by a specialized social media platform with the goal of encouraging people to write about their feelings and provide a tag. The quality of the self-annotated emotion data drawn from Vent was examined and led to the conclusion that the tagged emotional expressions are indicative of emotional content [35]. Furthermore, the distinction between reader-labeled and writer-labeled datasets was analyzed in [21] with the findings indicating that classifying the emotion labels of these datasets is a hard task when considering all available labels in the platform. As supervised methods tend to perform better than unsupervised ones and gathering manual annotations is time-consuming and expensive, this kind of self-annotated datasets offers a potential alternative beyond indirect self-annotations within the text as in Twitter hashtags.

3 Experimental setup

We illustrate our experimental setup in Fig. 1. Next, we describe this setup more in detail starting with the datasets for training and evaluating our models, followed by details on the implementation of our proposed models and baselines.

3.1 Datasets

The *Vent dataset* consists of 33 Million posts from the Vent social media app [20]. Each post is annotated by its author with an emotion tag as a way to express their emotional state to others. While the dataset has 705 emotion tags, many are temporary tags about seasonal events that do not express a clear emotional state and the most frequent tags

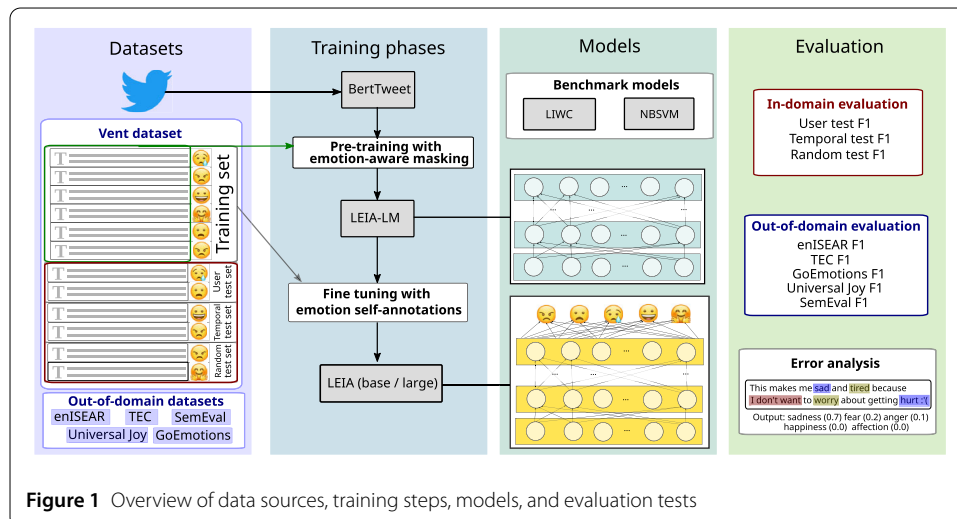


Figure 1 Overview of data sources, training steps, models, and evaluation tests

Table 1 Mapping of Vent categories to emotion labels

Label	Vent emotion tags
Sadness	Lonely, Sad, Miserable
Anger	Angry, Annoyed, Frustrated, Furious
Fear	Anxious, Stressed, Afraid, Nervous, Worried
Affection	Affectionate, Loving, Caring, Adoring, Cuddly, Supportive, Passionate, Infatuated
Happiness	Happy, Excited

are used on the vast majority of posts. Since Vent was designed to provide a nuanced expression of emotions rather than text classification, we mapped Vent emotion tags to a list of emotional states consistent with individual emotions from the affective science literature [36]. This way, we map emotion tags with words close in dimensional models of emotion [37] into the same label, for example, mapping the tags *angry* and *annoyed* into the same label of *Anger*. The precise mapping can be found in Table 1. Four of these emotion labels map to linguistic classes that have been consistently identified in emotional expression in text [38]: *Sadness*, *Anger*, *Fear*, and *Happiness*. We added a fifth category *Affection*, which occurs more frequently than *Happiness* and shows a social orientation of the expression of positive emotions on social media.

We pre-process the Vent dataset to generate a cleaner dataset of posts in English that were labeled by their authors with one of the tags of Table 1. We remove non-English posts using three language identification tools.^{1,2,3} For a post to be included in our analysis, at least two out of the three methods had to agree on detecting it as in English. After that, we remove duplicates and tag memes (invitations for a challenge to answer a question), following the approach in [35]. We remove posts with less than three words, excluding placeholders for links and user mentions in the word count. We also normalize the text by replacing multiple whitespaces with a single occurrence. We remove tab, new line and carriage return characters as well as Hypertext Markup Language codes. The resulting dataset contains more than nine million posts with metadata including the emotion labels, pseudonymized user ids, and timestamps when the post was written.

In-domain evaluation datasets An overview of this study can be seen in Fig. 1, including data sources and data splits for in-domain evaluation. We split the pre-processed Vent dataset into a training/development/test split with three disjoint test datasets to assess the capability of the model to generalize emotion identification. The random test set contains a uniformly random selection of 10% of all posts in the Vent dataset. The user test set consists of all posts written by a random sample of 10% of the users. This way, no post in the training set has been written by any of the users in the user test set. The temporal test set contains the last 10% of the posts according to their timestamp, thus allowing us to evaluate the model with future data with respect to its training set. We additionally extracted another 10% random set from the remaining posts as a development set to guide model design before the final run of all tests. All these subsets are disjoint and the three tests allow us to evaluate if and how the model generalizes across posts, users, and time. The resulting exact counts of posts and emotion labels in all splits can be found in Table 2.

¹<https://fasttext.cc/docs/en/language-identification.html>.

²<https://github.com/adbar/py3langid>.

³<https://github.com/google/cld3>.

Table 2 Frequency of occurrence of the labels on the data splits of the Vent dataset after pre-processing. The proportion of the total number of instances within the sample is in parenthesis

	Train	Development	User Test	Temporal test	Random test
Sadness	1,712,985 (27%)	199,890 (28%)	262,999 (27%)	293,993 (30%)	264,906 (27%)
Anger	1,517,282 (24%)	147,778(21%)	224,997 (23%)	205,598 (21%)	226,068 (23%)
Fear	1,341,624 (21%)	138,929 (20%)	198,264 (21%)	185,461 (19%)	201,563 (21%)
Affection	979,019 (15%)	144,175 (20%)	161,018 (17%)	191,022 (20%)	158,017 (16%)
Happiness	795,363 (13%)	74,369 (11%)	118,290 (12%)	91,127 (9%)	116,647 (12%)
<i>Total</i>	6,346,273	705,141	965,568	967,201	967,201

Out-of-domain evaluation datasets To evaluate if models learn about emotional expression beyond the domain of Vent as a social platform, we include five OOD datasets with emotion labels and texts associated with the emotions. The OOD datasets are the following:

- enISEAR [17] is a dataset of emotional event descriptions in English using the International Survey on Emotion Antecedents and Reactions (ISEAR) approach [39] via crowdsourcing. Annotators generated event-focused emotion descriptions using the template: “*I felt [emotion] when/because [situation]*”. While the study included annotations by readers, we only use the annotation of the author of the text to evaluate models. The dataset consists of 1001 instances for seven emotions, four of which match our emotion labels to provide an out-of-domain test. We design the task as a prediction of the text in which we have replaced the emotion word with the placeholder *mask*, which is a special token common in language models to denote a missing word. enISEAR is generated by asking participants to describe an emotion-inducing situation, a design that limits its external validity with respect to social media but that has the highest standard of internal validity with text annotations produced in a controlled setup. We consider enISEAR as the out-of-domain dataset most relevant to test the psychological validity of the emotion detection of models, while other datasets from social media are necessary to evaluate models in other domains once this psychological validity level is clear.
- GoEmotions [24] is a corpus of English comments extracted from Reddit with manual annotations for multiple emotions. It is a reader-labeled emotion dataset with labels assigned when at least three annotators gave the same label to a comment. For our out-of-domain test, we include the subset of the test split with a single label from among the Ekman category of the dataset, thus having *Sadness*, *Anger*, *Fear*, and *Joy* as a general positive emotion label.
- TEC [40] is a corpus of tweets posted between Nov. 15, 2011 and Dec. 6, 2011 with self-label for emotions using emotion-word hashtags. The hashtags serve as the emotion label for classification and are removed from the tweet texts. We sample 10% of the dataset at random as our out-of-domain test set. Since the hashtags are assigned by the authors of the tweets, the dataset can be considered labeled from the perspective of the writer.
- Universal Joy [34] is a collection of anonymized public Facebook posts in 18 languages labeled with five emotions: anger, anticipation, fear, joy, and sadness. The labels are derived from the Facebook “feelings tag” provided by the writers of the posts. We use the English subset of the test set for our analysis.

Table 3 Frequency of occurrence of the labels on the test sets of out-of-domain datasets

Dataset	Sadness	Anger	Fear	Happiness	Total
enISEAR	143	143	143	143	572
TEC	765	305	499	1627	3196
GoEmotions	259	520	77	1598	2454
Universal Joy	128	58	11	384	581
SemEval	312	511	165	706	1694

- SemEval [41] is a collection of tweets in three languages from 2016 and 2017 collected from Twitter using emotion keywords as queries. Subsequently, matching tweets were annotated by crowdworkers for emotion intensity, valence, and basic emotion classes. This dataset was the benchmark data for the competition about affect detection in SemEval. Here, we use the test data by including only instances with a single label that correspond to one of the labels in our model.

Note that for the OOD datasets (GoEmotions, TEC, Universal Joy, and SemEval), we use only the test sample for OOD evaluation and exclude other training or development samples. We do this to provide an evaluation that can be compared to previous and future supervised methods that use the training samples.

Based on our selection criteria, we find only 11 tweets with the *Affection* label in the SemEval dataset. So, we consider *Happiness* and *Affection* to be the *Happiness* emotion label, which limits the nuance in which we can assess classifications within positive emotions in out-of-domain settings but still enables a wider differentiation between general positive emotions and three negative emotions. Descriptive statistics of the counts and proportions of labels in the five datasets can be found in Table 3.

We use the in-domain and OOD datasets to evaluate the performance of models in our experimental setup. We calculate the macro-averaged F1 score over all emotion labels and report results with the F1 score of each of the emotion labels, as their frequencies greatly differ in several of the datasets we use for evaluation.

3.2 Models

Model design and pre-training Pre-trained language models have shown state-of-the-art performance on many natural language processing tasks. We expect language models pre-trained on social media data to perform better on the Vent dataset. In preliminary experiments using performance on the development set, we test three pre-trained models based on the Robustly optimized BERT approach (RoBERTa) architecture and pre-training: Roberta-base [42], Twitter-RoBERTa [23], and BERTweet-base [25]. BERTweet-base had the best performance on the development set and thus we chose to continue our work with BERTweet-base and its large version, BERTweet-large, in all our experiments. BERTweet-base and BERTweet-large are transformers model pre-trained on 850M tweets with 12 and 24 layers, respectively. BERTweet-base has a maximum sequence length of 128 (sub)words while BERTweet-large has a maximum sequence length of 512 (sub)words [25]. Before training a classifier on the training set, we pre-train BERTweet-base (BERTweet-large) on the text of Vent posts in the training set ignoring all emotion labels. We perform task-adaptive pre-training [29] by preferentially masking emotion words using eMLM. We use the emotion terms in the emotion lexicon introduced in [43, 44] as it is one of the most extensive emotion lexicons available. We set the probability of masking emotion words to 0.5 following previous work [31]. We train with the eMLM objective for 100K steps using

the AdamW optimizer [45], a learning rate of $5 * 10^{-5}$, and a batch size of 128. We name the resulting models LEIA-LM-base and LEIA-LM-large, i.e. the result of our pre-training of BERTweet-base and BERTweet-large respectively. On an NVIDIA RTX8000 GPU, pre-training takes approximately a week for the base model and a month for the large model.

Model fine-tuning with labeled data We implement a multiclass classifier for the five emotion labels: *Anger*, *Fear*, *Sadness*, *Happiness*, and *Affection*. We train classifiers starting from LEIA-LM-base and LEIA-LM-large using a two-step approach. First, we perform linear probing to initialize the classifier head and then full fine-tuning of the model. For linear probing, only the classifier head is randomly initialized and trained on the training dataset while the remaining model parameters are fixed. This initial step can be seen as a way to align the features of the prediction head and the base model to minimize feature distortion [32]. In the subsequent full fine-tuning step, the prediction head is initialized from the parameters learned from the initial linear probing step. We also fine-tune a BERTweet-base and a BERTweet-large model without the eMLM step. To improve model generalization, we average model weights [46] of the two model variants (one with eMLM and one without eMLM) for each of the base and large architectures. The resulting models are respectively named LEIA-base and LEIA-large. We show the performance of the intermediate model variants on the in-domain and OOD test sets in Tables 9 and 10 in the Appendix. For the linear probing step, we use a learning rate of $5 * 10^{-4}$ and train only the classifier head while the other layers are frozen for 1000 steps. For fine-tuning, we set the learning rate to 10^{-5} with a constant learning rate schedule, embedding dropout of 0.1, weight decay factor of 0.01, and a label smoothing factor of 0.1. We train for 5 epochs using AdamW optimizer with an effective batch size of 256 and a maximum sequence length of 128. We jointly optimize a supervised contrastive loss and a cross-entropy loss [47]. The supervised contrastive loss ensures that the model captures the similarity between examples within a class while contrasting them with examples from other classes. This approach has been shown to aid model generalization. Following prior work [47], we set the weight of the contrastive loss to 0.9 and the temperature parameter to 0.3. The fine-tuning process takes approximately 24 hours for the base-sized model and 60 hours for the large-sized model on an Nvidia RTX8000 GPU with 48 GB memory.

Baselines As baselines, we use the popular Linguistic Inquiry and Word Count (LIWC) dictionary approach [48], the NRC emotion lexicon [43, 44], and a Naive Bayes Support Vector Machine (NBSVM) as a supervised baseline. For the LIWC approach, we map the score for the relevant LIWC categories to emotion labels as follows: *emo_anger* to Anger, *emo_anx* to Fear, *emo_sad* to Sadness, and *emo_pos* to Happiness. For NRC, we compute the frequency of emotion words corresponding to the emotion categories we consider normalized by the length of the text. We did not find a category that can be mapped to Affection in the LIWC and NRC categories, thus considering only 4 classes for the dictionary-based baselines. We convert the multiclass result of LIWC and NRC to a binary classification task for each emotion label using the “one-vs-rest” setting. For Sadness category as an example, we consider instances within the Sadness category as having a label of 1 if the frequency of occurrence of words in the Sadness category normalized by the length of the text is greater than 0. All other examples are assigned a label of 0.

We use NBSVM [49] as a supervised baseline. NBSVM is a strong baseline for text classification that uses Naive Bayes features for unigrams as input representation. We use the implementation in Ktrain [50] with a vocabulary size of 64K.

Polarity benchmark The discrete emotions that we consider can be grouped by valence into positive polarity (happiness and affection) and negative polarity (anger, fear, and sadness). We examine the effectiveness of post-processing the prediction from our models into a polarity classifier by taking the sum of the probability for happiness and affection as the positive polarity score. We then compare this strategy to two well-known approaches for polarity classification of online content: LIWC and VADER [51]. We use the Tone score from LIWC and the compound score from VADER as the polarity score, respectively. For each method, we compute the Area Under the Receiver Operating Characteristics Curve (AUC) for the prediction of positive/negative polarity.

Comparison with GPT models Large generative models provide the opportunity for the classification of emotions in social media text in emotion classes as our models, but rate limits and pricing hinder an exhaustive evaluation with our full Vent test datasets, which contain nearly 3 million posts. To provide a comparison against OpenAI's GPT-3.5 and GPT-4 models, we design an additional in-domain analysis based on a random sample of 1000 posts from each emotion category in the user test split of the Vent dataset. For an out-of-domain comparison, we ran both GPT models over the full test samples of the OOD datasets. We prompt GPT models with instructions to output only one of the emotion categories for each input text via the OpenAI *chat/completions* endpoint with a temperature of 0. For cases, where the model did not return the expected output we repeatedly prompt the model to output one of the emotion categories. In addition, we limit the number of output tokens to 3 (the maximum number of tokens required by the Generative Pretrained Transformer (GPT) tokenizer to cover all emotion categories).

4 Results and analysis

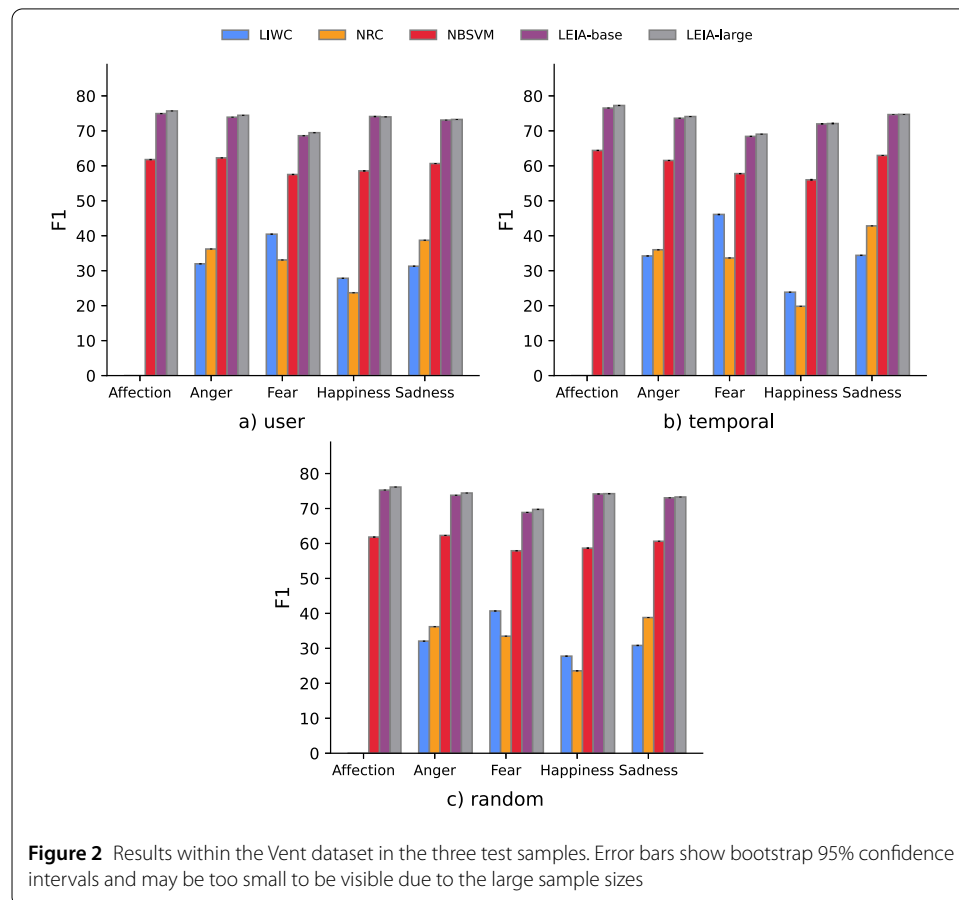
In this section, we report the performance of LEIA-base and LEIA-large in both in-domain and out-of-domain scenarios. We include the macro-F1 score and bootstrapping confidence intervals obtained from 10,000 bootstrap samples. We provide an error analysis on a sample of incorrect model predictions. We end by assessing the salient features on selected examples of model predictions.

In-domain results Table 4 shows that LEIA-base and LEIA-large outperform all models in all three Vent test samples, achieving a Macro-F1 of about 73 on random posts, text from unseen users and different time periods. Model performance is comparable across all three test sets, which indicates that its F1 score is not achieved by exploiting biases of user activity or high-volume time periods. The dictionary approaches have the lowest macro-F1 scores, being significantly outperformed by LEIA-base and LEIA-large. The supervised approach of NBSVM achieves macro-F1 scores of about 60 but is still substantially and significantly outperformed by LEIA-base and LEIA-large.

Figure 2 shows a breakdown of F1 per emotion class in the in-domain test samples. LEIA-base and LEIA-large show consistently high F1 score for all emotion classes. This shows that the general performance of LEIA-base and LEIA-large is not as a result of

Table 4 Macro-F1 scores on the Vent test sets. 95% Confidence interval in square brackets (computed over 10,000 bootstrap samples). For LIWC and NRC, we only consider 4 out of 5 labels and perform binary classification for each label using the “one-vs-rest” setting

	LIWC	NRC	NBSVM	LEIA-base	LEIA-large
User	32.88 [32.79, 32.96]	32.91 [32.85, 32.97]	60.15 [60.05, 60.25]	72.92 [72.82, 73.02]	73.37 [73.28, 73.46]
Temporal	34.64 [34.55, 34.72]	33.07 [33.01, 33.12]	60.52 [60.42, 60.63]	73.03 [72.95, 73.11]	73.43 [73.34, 73.53]
Random	32.84 [32.74, 32.94]	33.02 [32.95, 33.08]	60.26 [60.16, 60.36]	73.02 [72.94, 73.12]	73.57 [73.48, 73.66]



bias from higher performance on majority class. The only class that has a slightly lower F1 is Fear, but LEIA-base and LEIA-large still outperform all other methods on it. One observation is that NBSVM also performs slightly worse for Fear than for other emotions in contrast with LIWC, which obtains a comparatively better performance than NRC in the Fear category.

Out-of-domain results Our out-of-domain benchmark shows that LEIA can detect emotional states in other types of text and social media platforms beyond Vent. Table 5 shows the Macro-F1 scores for the five out-of-domain test sets. LEIA-base and LEIA-large have significantly higher F1 scores than all other methods when evaluated on 4 out of the 5 OOD datasets. The NBSVM has a comparable performance in the GoEmotions dataset, where the F1 of NBSVM and of LEIA-base are not significantly different. We also observe that a larger model does not necessarily lead to better performance on OOD datasets, as LEIA-large only shows a substantially different performance on the enISEAR dataset.

Table 5 Macro-F1 scores on out-of-domain datasets. 95% Confidence intervals in square brackets (computed over 10,000 bootstrap samples)

	LIWC	NRC	NBSVM	LEIA-base	LEIA-large
Universal Joy	23.45 [19.97, 27.29]	28.98 [26.15, 31.99]	41.70 [37.36, 46.08]	54.18 [48.79, 59.88]	54.17 [48.68, 59.84]
GoEmotions	45.81 [42.72, 48.63]	32.68 [31.18, 34.25]	48.23 [45.85, 50.59]	46.31 [43.98, 48.72]	45.75 [43.45, 48.09]
TEC	36.02 [34.02, 37.99]	33.92 [32.65, 35.27]	39.07 [37.28, 40.92]	43.87 [42.05, 45.61]	44.12 [42.34, 45.89]
SemEval	66.72 [64.61, 69.1]	49.86 [48.27, 51.4]	68.77 [66.29, 71.25]	71.68 [69.18, 74.19]	70.04 [67.48, 72.52]
enISEAR	23.51 [19.43, 26.89]	42.72 [40.26, 44.89]	55.33 [51.22, 59.41]	70.37 [66.63, 74.01]	79.94 [76.69, 83.14]

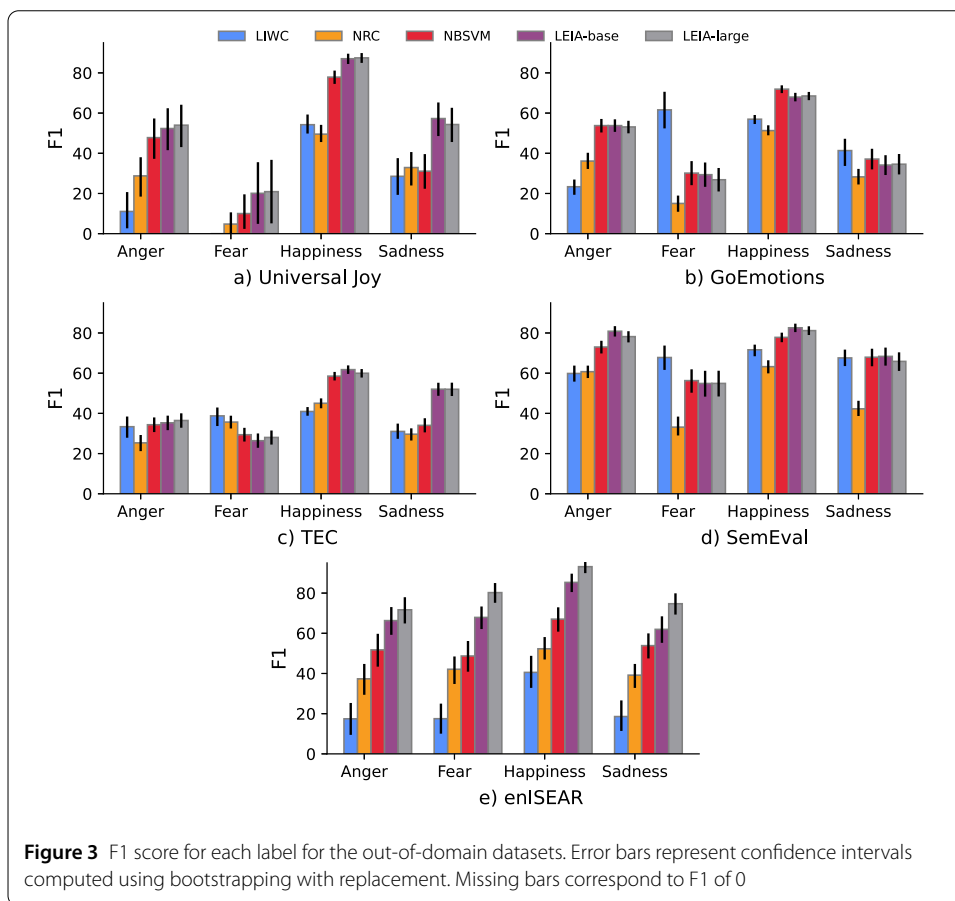


Figure 3 shows the F1 score for each class on the OOD datasets. In general, LEIA often outperforms baselines across labels. LEIA is significantly better than the baselines for Happiness and Sadness in the Universal Joy and TEC datasets, for all emotions in the enISEAR dataset, and for all emotions except Fear and Sadness in the SemEval dataset. On the GoEmotions dataset, LEIA is tied with NBSVM as the best method to detect Anger as F1 score is not significantly different. The Fear class evaluation poses some challenges in this OOD evaluation since evaluation samples for this class can be very small (e.g. 11 posts in Universal Joy and 77 in GoEmotions). In the case of Fear, LIWC performs significantly better than the supervised approaches on GoEmotions, SemEval, and TEC. Recall that the dictionary approach is based on a binary classification setting which is easier than a multiclass classification setting. Despite this, the performance of the dictionary approach is significantly lower for Happiness. This trend is similar to the performance observed on the in-domain test sets.

We can conclude that LEIA shows a good generalization beyond the domain it was trained on, first by achieving very high performance in enISEAR, the test closest to psychological methodology, but also achieving good performance for datasets that include posts from other social media such as Twitter and Facebook. The lower performance recorded for Fear on the out-of-domain test sets is not surprising as the model performance on this category tends to be lower on the in-domain test sets too. LEIA achieves a consistently high score for Happiness on the out-of-domain test sets despite the fact that it is one of the least frequent categories in the training set. This suggests that it constitutes an easier category for the model to recognize across domains than more nuanced negative emotions.

Comparison with GPT-3.5 and GPT-4 models Table 6 shows the performance comparison of LEIA-base, LEIA-large, GPT-3.5, and GPT-4 on a sample of 1000 examples for each emotion category drawn from the user test split of the Vent dataset. LEIA-base and LEIA-large perform better on all emotion classes of the Vent dataset than GPT-3.5 and GPT-4. We show the performance comparison on the out-of-domain datasets in Table 7. The F1 score per emotion category on each dataset is in Fig. 5 in the Appendix. GPT-3.5 and GPT-4 perform better than both LEIA-base and LEIA-large on the OOD datasets. Our results are consistent with recent findings showing that smaller models tailored for specific tasks perform better than large generative models such as GPT-3.5 and GPT-4, especially when evaluated against datasets that are unlikely to be part of the training data of GPT models [52, 53]. A visible trend from Table 7 is that GPT-3.5 and GPT-4 models show higher performance on datasets on which LEIA-base and LEIA-large show relatively higher performance and vice-versa. This may point to the level of difficulty of some of the OOD datasets. Although GPT-3.5 and GPT-4 models perform better than our models on the OOD datasets, we do not know whether this performance is clearly a capability of the model or due to data contamination as these models are trained on massive datasets which may include benchmark datasets [54]. Moreover, it has also been documented that it is challenging for large language models to infer mental state from textual data [55]. As noted by the authors of [53], an avenue for future work is to explore approaches that

Table 6 Comparison of LEIA-base, LEIA-large, GPT-3.5, and GPT-4 on a random sample of Vent user test split consisting of 1000 examples per emotion category

	LEIA-base	LEIA-large	GPT-3.5	GPT-4
Affection	74.48 [72.30, 76.57]	75.67 [73.56, 77.72]	41.38 [38.09, 44.69]	37.43 [34.02, 40.78]
Anger	72.92 [70.76, 75.01]	72.98 [70.87, 75.00]	61.79 [59.16, 64.26]	66.82 [64.42, 69.17]
Fear	69.01 [66.59, 71.35]	70.26 [67.89, 72.55]	51.55 [48.53, 54.55]	60.86 [58.17, 63.48]
Happiness	77.69 [75.52, 79.77]	77.58 [75.38, 79.64]	67.69 [65.60, 69.77]	68.70 [66.61, 70.77]
Sadness	67.28 [65.00, 69.47]	68.00 [65.73, 70.22]	59.94 [57.79, 62.11]	64.00 [61.82, 66.18]
Average	72.28 [71.04, 73.50]	72.90 [71.67, 74.11]	56.47 [55.07, 57.90]	59.56 [58.20, 60.92]

Table 7 Macro-F1 of LEIA-base, LEIA-large, GPT-3.5, and GPT-4 on the out-of-domain test sets

	LEIA-base	LEIA-large	GPT-3.5	GPT-4
Universal Joy	54.18 [48.79, 59.88]	54.17 [48.68, 59.84]	52.89 [47.8, 58.33]	56.43 [51.94, 60.9]
GoEmotions	46.31 [43.98, 48.72]	45.75 [43.45, 48.09]	59.06 [56.44, 61.67]	56.45 [53.97, 58.95]
TEC	43.87 [42.05, 45.61]	44.12 [42.34, 45.89]	52.66 [50.7, 54.59]	54.82 [52.96, 56.71]
SemEval	71.68 [69.18, 74.19]	70.04 [67.48, 72.52]	80.13 [77.9, 82.31]	81.72 [79.57, 83.81]
enISEAR	70.37 [66.63, 74.01]	79.94 [76.69, 83.14]	84.96 [81.96, 87.77]	89.97 [87.47, 92.32]

Table 8 Area Under the Receiver Operating Characteristic Curve (AUC ROC) for polarity classification on the out-of-domain datasets

	enISEAR	Universal Joy	GoEmotions	TEC	SemEval
LIWC	0.829 [0.810, 0.852]	0.680 [0.650, 0.721]	0.826 [0.814, 0.842]	0.677 [0.661, 0.692]	0.881 [0.868, 0.892]
VADER	0.852 [0.829, 0.875]	0.678 [0.637, 0.725]	0.880 [0.875, 0.885]	0.651 [0.633, 0.664]	0.905 [0.898, 0.911]
LEIA-base	0.983 [0.975, 0.988]	0.883 [0.856, 0.912]	0.820 [0.806, 0.832]	0.684 [0.674, 0.701]	0.920 [0.911, 0.931]
LEIA-large	0.989 [0.986, 0.992]	0.861 [0.834, 0.883]	0.817 [0.798, 0.835]	0.641 [0.626, 0.652]	0.894 [0.885, 0.904]

combine large generative models with smaller domain-specific models that can be applied efficiently and at scale.

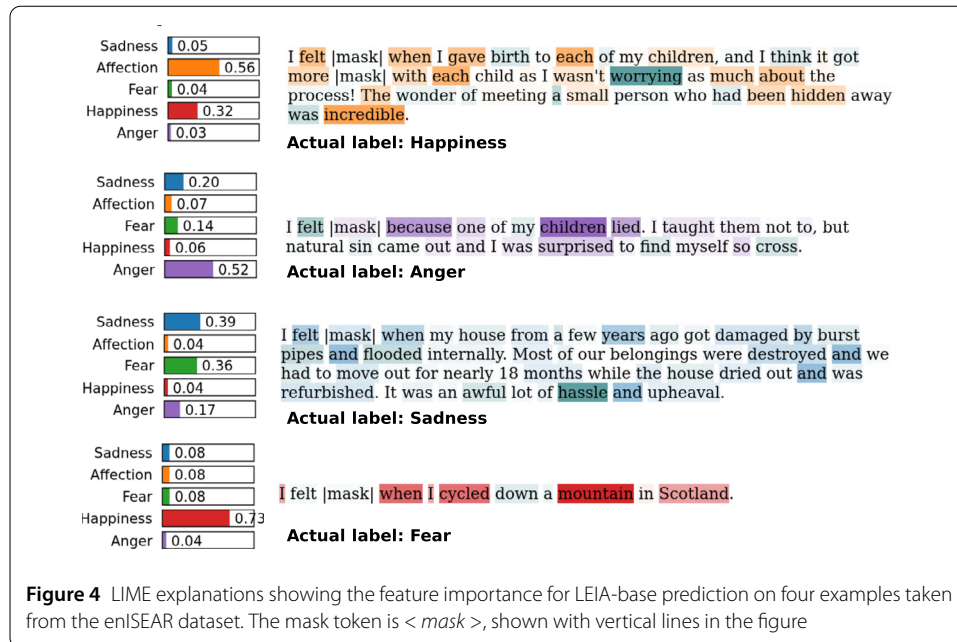
Polarity classification benchmark Results on the out-of-domain datasets can be found in Table 8. On 4 out of the 5 out-of-domain datasets LEIA-base and/or LEIA-large perform better than the dictionary baselines: LIWC and VADER. LIWC and VADER show better performance only on the GoEmotions dataset. Comparing LIWC and VADER, we find that VADER performs better than LIWC on enISEAR, GoEmotions, and SemEval while LIWC is superior on Universal Joy and TEC.

Error analysis We examine a random sample of 50 incorrect predictions from the user test split (10 per label) of the Vent dataset. We find that majority of errors in the sample can be categorized into the following cases:

1. Messages conveying an expectation of a positive outcome while the self-assigned label has negative valence (e.g., I need a good online game). These cases represent situations where the text is very similar to positive texts but subtle signals point toward negative states.
2. Expressions of both positive and negative emotions at the same time. These are assigned a single label by design but other labelling schemes could cope with mixed emotions.
3. Use of figurative expressions such as humor or sarcasm that the model does not recognize.
4. Very short posts that do not contain indications about the emotional state of the author (e.g., going for a coffee) where additional context is required.
5. Few instances where we find the model prediction more plausible than the assigned label.

As an additional form of error analysis, we assess whether our grouping of Vent tags into emotion categories shown in Table 1 is realistic. For this, we examine the variation of recall across tags within each emotion category. We compute the recall for each tag from the prediction of one of our models, LEIA-base, on the user test split of the Vent dataset. We consider the prediction a hit for a given tag when the model makes the correct emotion prediction and a miss otherwise. We find that the recall for each tag within each group is comparable ranging from about 0.6 to 0.8, which suggests that our grouping is reasonable. The details of the recall score for each tag are in Fig. 6 in the Appendix.

Feature attributions We examine the salient features that contribute to the predictions made by LEIA-base on a set of examples from the enISEAR dataset. We apply the Local Interpretable Model-agnostic Explanations (LIME) method for model interpretability [56],



an attribution method for identifying salient features as n-grams of the classified text. Figure 4 shows four examples, one for each class of emotions in the enISEAR test set. The first column shows the model confidence scores for each class supported by LEIA-base and the text is colored according to which words contribute to the prediction.

We observe that for the first example, the model incorrectly predicts Affection as the most likely label where the true label is Happiness, which is an error of a weaker kind since enISEAR does not have an Affection label and both emotions are close in terms of valence. The second highest class is Happiness and the prediction is positively based on words expressing high arousal and valence (e.g., “incredible”) and negatively based on the word “worrying”. In the second example, the model also seems to use relevant words linked to each other (e.g., “children” and “lied”) to make the correct prediction. The model correctly predicts Sadness for the third example building on negative words, including terms linked to property damage that caused an emotional loss. We observe that the scores for fear and sadness are very close and much higher than for other classes. This seems plausible as the first sentence in this example could be a fearful situation. The model prediction is Happiness in the fourth example instead of Fear, which was the true label. Even though the prediction relies on relevant features, the model seems to lack the commonsense knowledge that cycling down a mountain can be scary and not necessarily a pleasant experience.

The last two cases suggest that the emotion tag for some of the posts is used as the main medium to express the emotion, leaving the text to add other information. This is one of the limitations of using Vent as a training dataset, as labels are part of the communication and may sometimes be complementary or otherwise to the posts.

5 Discussion

We present LEIA, a language model in two sizes (LEIA-base and LEIA-large), that leverages approaches for adapting pre-trained language models for emotion identification. We show that using an emotion lexicon with task-adaptive pre-training, in this case focusing on emotion words, is effective for improving model performance using BERTweet-base

and BERTweet-large language models. LEIA generalizes beyond Vent posts as it shows better performance on texts written by users not included in its training data and future time periods. It achieves a balanced performance across emotion labels despite their imbalance in training data and this performance is also seen on out-of-domain texts for the considered emotions except for Fear. These results are in part possible thanks to focusing on a small set of emotions suggested by psychological research, as classifying the larger set of mood labels in Vent [21] is a substantially harder task we did not tackle here. Also, the Vent dataset, which despite being generated on a platform not as large as common ones in research, e.g. Twitter and Reddit, has a sufficiently large scale that enables the models to learn a broader range of emotional expressions.

The performance of LEIA-base is comparable to LEIA-large across tests in our benchmark with one notable exception: LEIA-large is substantially better for the enISEAR dataset. This dataset is especially important given the psychological methodology used to generate it, which allows us to compare the results of machine learning methods with self-reported labels in a controlled setup. LEIA's performance in enISEAR is especially high, reaching F1 of 70 for LEIA-base and 79 for LEIA-large, showing a high level of psychological validity, especially when compared to other methods in the benchmark that achieve at most 55. LIWC and NRC generally achieve low F1 in all tests except SemEval, which grants two notes. First, SemEval was generated by searching tweets with emotion-bearing terms, easing the task for dictionary approaches when classifying emotions based on similar word lists. Second, LIWC and NRC were not designed as an emotion classification method at the scope of a social media post. LIWC is a more general text analysis method that should be applied to longer texts and not necessarily for classification. We added LIWC and NRC to contrast with common methods applied in the field, but our comparison overstretches the applications for which these resources were designed.

On a sample of 1000 examples drawn from the user test split of Vent, LEIA-base and LEIA-large surpass the performance of GPT-3.5 and GPT-4. On the OOD datasets, GPT-3.5 and GPT-4 perform better than LEIA-base and LEIA-large. This finding is in line with existing findings that show that smaller domain-specific models perform better than larger general-purpose generative models. One issue with assessing the real capability of large language models is the possibility of data contamination where benchmark datasets can potentially be part of their training data. Large language models are often accessible via Application Programming Interfaces (APIs) which make it easy in practice to use with its attendant cost. However, rate limits and financial costs make it a less attractive option to apply at scale. This is even complicated by the need to make repeated calls to the API when the model does not follow the instructions provided in the prompt. LEIA-base and LEIA-large are openly available and can be run efficiently at scale when needed. Our models also provide an additional benefit, access to the confidence of the model predictions. This can benefit downstream analyses or can be useful to know when to rely on the model prediction. This is currently not possible by prompting existing large generative models.

Limitations While we show that our proposed models are effective, our experiments span two model sizes with the same architecture. Future research should conduct experiments on other pre-training approaches beyond masking as well as more efficient training techniques. In addition, we rely mostly on hyperparameter settings in the literature and optimizing them could lead to better performance. However, this is computationally

expensive and there might be unfavorable trade-offs between model performance and resources. Another limitation is our focus mainly on English posts, providing no evidence here of the potential of this approach for other languages. Furthermore, we study five emotion labels guided by psychological research, but several competing representations models for emotion are available. Humans are able to classify a larger number of basic emotions and can also quantify emotions in dimensional spaces, two open areas that can be explored with more nuanced labeling schemes. While self-annotated datasets have the potential to become the new gold standard beyond crowdworkers, the labeling scheme of the Vent dataset is designed as part of its interface rather than as a psychometric measure applied privately and not visible on the platform. This is still closer to general emotion expression than automatic labeling with emoji or hashtags, but models like LEIA-base or LEIA-large can be substantially improved with psychological methods like experience sampling [18] and with validated psychological scales to measure emotions in dimensional spaces [57, 58]. This would have the added value of being applicable to studying more nuanced emotion dynamics that need dimensional measurements and not just classification, for example using social media data [6, 59, 60].

Broader impact and ethical considerations This work shares the same ethical concerns with other emotion recognition systems as highlighted in [61]. Emotion detection models should be used responsibly and special care should be taken when they are applied in new scenarios, not only because of their possible lower performance but also due to possible different privacy expectations with respect to emotions. We must note that we have no way of estimating the demographic diversity of Vent users and it is very likely that the model misses idiosyncrasies of emotional expression in minority groups and in cultures not represented in the dataset. We acknowledge that we only consider one type of model evaluation focusing on accuracy while there are several aspects such as bias, fairness, and robustness that should be considered before a model is used in practice, especially when guiding any decision-making.

6 Conclusion

LEIA is an emotion detection method that achieves a balanced performance across emotions and generalizes across posts, users, and time. It shows satisfactory performance in out-of-domain tests, especially when compared to self-annotated texts produced with psychological methods. Beyond our validations, the language models within LEIA can be used as pre-training resources for future applications that employ annotated data in other domains, for example for tweets in particular contexts.

We named LEIA after Princess Leia from Star Wars, following the tradition of emotion method names set out by LIWC [62] (pronounced Luke, as in Luke Skywalker), and VADER [51] (as in Darth Vader). These three methods have a similar purpose but very different approaches that align with concurrent developments in text analysis. We published openly our models in HuggingFace (<https://huggingface.co/LEIA>) including both the classifier LEIA-base (LEIA-large) and the corresponding emotion-aware language model with the hope that they can be used in future work in emotion detection from text.

Appendix

Comparison of intermediate models Recall that we average the parameters of two model variants. A BERTweet-base model fine-tuned on the labeled Vent training set and a BERTweet-base with eMLM pre-training before fine-tuning to derive LEIA-base. We follow the same approach for BERTweet-large to derive LEIA-large. Here, we compare the performance of each of these models on in-domain and OOD test sets. Results in Table 9 show that LEIA-base and LEIA-large have the highest average rank across in-domain test sets. We also see a similar pattern in Table 10 for the OOD datasets. This suggests that taking an average of model parameters contribute to the generalization of LEIA-base and LEIA-large. In addition, the eMLM pre-training step is also effective given that the models with this pre-training step rank second across both in-domain and out-of-domain test sets.

Comparison of LEIA-base, LEIA-large, GPT-3.5, and GPT-4 on the OOD datasets Figure 5 provides the detailed F1 score per emotion category on each of the OOD datasets for LEIA-base, LEIA-large, GPT-3.5, and GPT-4.

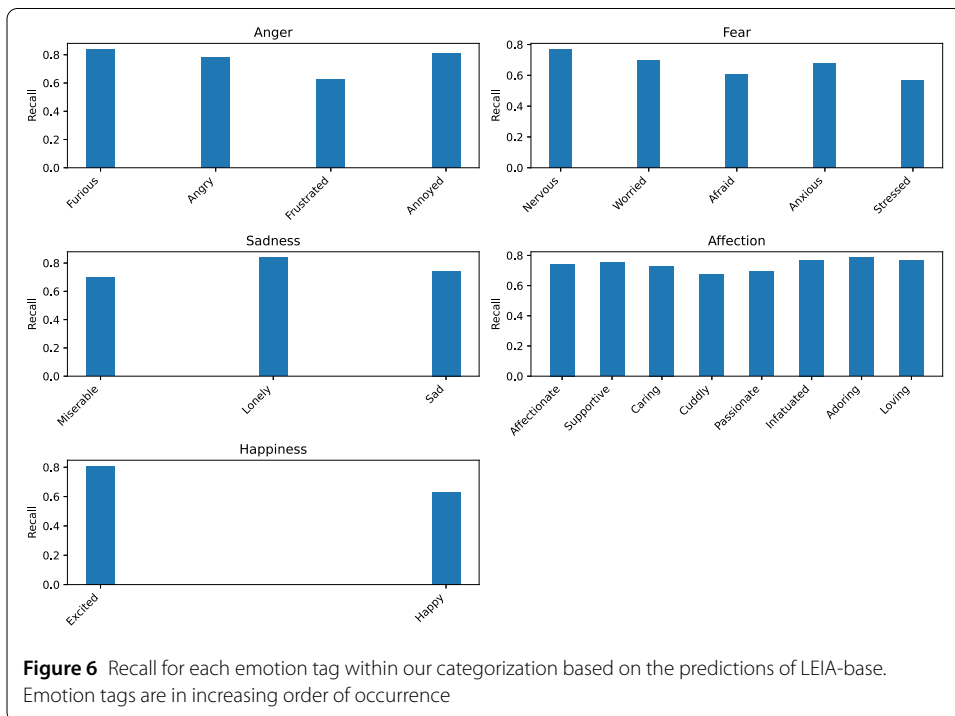
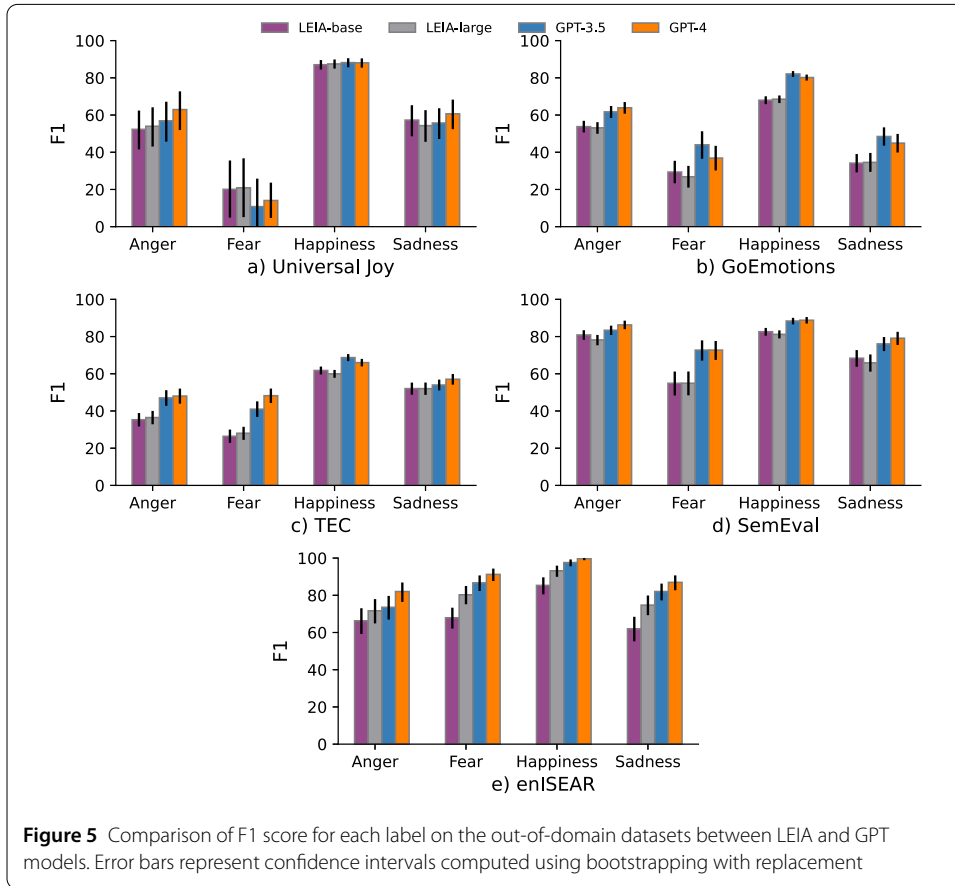
Recall per Vent tag within each emotion category In order to assess whether our grouping of Vent tags into emotion categories shown in Table 1 is realistic, we compute the recall for each tag from the prediction of one of our models, LEIA-base, on the user test split. Figure 6 depicts the recall for each group. Overall, we find that the recall for each tag within each group is comparable ranging from about 0.6 to 0.8. This suggests that our grouping is reasonable.

Table 9 Comparison of performance of intermediate models on in-domain test sets

	User	Temporal	Random	Average rank
BERTweet-base	72.82 [72.72, 72.91]	72.85 [72.77, 72.95]	72.92 [72.81, 73.0]	3.00
BERTweet-base+eMLM	72.87 [72.79, 72.97]	72.91 [72.82, 73.00]	73.00 [72.92, 73.09]	2.00
LEIA-base	72.92 [72.82, 73.02]	73.03 [72.95, 73.11]	73.02 [72.94, 73.12]	1.00
BERTweet-large	73.03 [72.94, 73.11]	73.00 [72.91, 73.12]	73.19 [73.09, 73.28]	3.00
BERTweet-large+eMLM	73.19 [73.11, 73.28]	73.14 [73.05, 73.23]	73.37 [73.27, 73.46]	2.00
LEIA-large	73.37 [73.28, 73.46]	73.43 [73.34, 73.53]	73.57 [73.48, 73.66]	1.00

Table 10 Performance comparison of intermediate models on out-of-domain datasets

	Universal Joy	GoEmotions	TEC	SemEval	enlSEAR	Average rank
BERTweet-base	52.40 [47.01, 57.99]	47.16 [44.74, 49.61]	43.68 [41.90, 45.41]	69.70 [67.14, 72.25]	67.23 [63.30, 71.06]	2.60
BERTweet-base+eMLM	52.07 [46.85, 57.61]	46.26 [43.89, 48.66]	42.43 [40.64, 44.16]	70.45 [67.90, 72.97]	74.79 [71.24, 78.22]	2.00
LEIA-base	54.18 [48.79, 59.88]	46.31 [43.98, 48.72]	43.87 [42.05, 45.61]	71.68 [69.18, 74.19]	70.37 [66.63, 74.01]	1.40
BERTweet-large	51.87 [46.77, 57.32]	45.51 [43.14, 47.94]	44.08 [42.28, 45.85]	68.75 [66.20, 71.31]	80.49 [77.18, 83.72]	2.40
BERTweet-large+eMLM	52.28 [47.01, 57.79]	46.94 [44.58, 49.31]	43.78 [41.98, 45.57]	69.02 [66.47, 71.55]	80.09 [76.77, 83.37]	2.00
LEIA-large	54.17 [48.68, 59.84]	45.75 [43.45, 48.09]	44.12 [42.34, 45.89]	70.04 [67.48, 72.52]	79.94 [76.69, 83.14]	1.60



Funding

Open Access funding enabled and organized by Projekt DEAL. The research leading to these results received funding from the Vienna Science and Technology Fund (WWTF) [10.47379/VRG16005]. DG and SA acknowledge funding from the ERC Advanced Grant PRODEMINFO (101020961).

Abbreviations

LMs, Language Models; LEIA, Linguistic Embeddings for the Identification of Affect; SentiLARE, Sentiment-aware language representation learning; ISEAR, International Survey on Emotion Antecedents and Reactions; eMLM, Emotion Masked Language Modeling; BERT, Bidirectional Encoder Representations from Transformers; RoBERTa, Robustly optimized BERT approach; LIWC, Linguistic Inquiry and Word Count; NBSVM, Naive Bayes Support Vector Machine; OOD, out-of-domain; LIME, Local Interpretable Model-agnostic Explanations; LEIA-LM-base, LEIA language model with base-sized architecture; LEIA-LM-large, LEIA language model with large-sized architecture; LEIA-base, LEIA classifier with base-sized architecture; LEIA-large, LEIA classifier with large-sized architecture; TEC, Twitter corpus annotated with emotion-word hashtag; SemEval, Twitter dataset for identification of affect in SemEval competition; GoEmotions, Reddit dataset for emotion identification; enISEAR, English ISEAR corpus with writer-assigned emotion label; GPT, Generative Pretrained Transformer; GPT-3.5 and GPT-4, Generative Pretrained Transformer versions 3.5 and 4; API, Application Programming Interface; VADER, Valence Aware Dictionary and sEntiment Reasoner; NRC, Emotion lexicon developed by researchers at the National Research Council of Canada; AUC, Area Under the Receiver Operating Characteristics Curve.

Availability of data and materials

All data used in this article is of secondary use and is available for researchers in the citations provided in the main text. The models produced for this article are publicly available at <https://huggingface.co/LEIA>.

Declarations

Competing interests

The authors declare that they have no competing interests.

Author contributions

DG, HM, and ADN designed research. LM and NH contributed to model development. SA designed and trained the final model and performed all analyses. DG and SA wrote the first manuscript draft. All authors read and approved the final manuscript.

Author details

¹Department of Politics and Public Administration, University of Konstanz, Konstanz, Germany. ²Graz University of Technology, Graz, Austria. ³Complexity Science Hub, Vienna, Austria. ⁴Université Paris Saclay, Paris, France. ⁵Medical University of Vienna, Vienna, Austria. ⁶Vienna University of Technology, Vienna, Austria.

Received: 12 April 2023 Accepted: 30 October 2023 Published online: 16 November 2023

References

- Pellert M, Schweighofer S, Garcia D (2021) Social media data in affective science. In: Handbook of computational social science, vol 1, pp 240–255. Routledge, London. <https://doi.org/10.4324/9781003024583-18>
- De Choudhury M, Counts S, Gamon M (2012) Not all moods are created equal! Exploring human emotional states in social media. In: ICWSM, vol 6, pp 66–73
- Golder SA, Macy MW (2011) Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science* 333(6051):1878–1881
- García D, Rimé B (2019) Collective emotions and social resilience in the digital traces after a terrorist attack. *Psychol Sci* 30(4):617–628
- Ferrara E, Yang Z (2015) Measuring emotional contagion in social media. *PLoS ONE* 10(11):0142390
- Fan R, Varol O, Varamesh A, Barron A, van de Leemput IA, Scheffer M, Bollen J (2019) The minute-scale dynamics of online emotions reveal the effects of affect labeling. *Nat Hum Behav* 3(1):92–100
- Zheng S, Wang J, Sun C, Zhang X, Kahn ME (2019) Air pollution lowers Chinese urbanites' expressed happiness on social media. *Nat Hum Behav* 3(3):237–243
- Field A, Park CY, Theophilo A, Watson-Daniels J, Tsvetkov Y (2022) An analysis of emotions and the prominence of positivity in #blacklivesmatter tweets. *Proc Natl Acad Sci* 119(35):2205767119
- Pellert M, Metzler H, Matzenberger M, Garcia D (2022) Validating daily social media macroscopes of emotions. *Sci Rep* 12(1):1–8
- García D, Pellert M, Lasser J, Metzler H (2021) Social media emotion macroscopes reflect emotional experiences in society at large. *arXiv preprint. arXiv:2107.13236*
- Calvo RA, Milne DN, Hussain MS, Christensen H (2017) Natural language processing in mental health applications using non-clinical texts. *Nat Lang Eng* 23(5):649–685
- Rajamanickam S, Mishra P, Yannakoudakis H, Shutova E (2020) Joint modelling of emotion and abusive language detection. In: *ACL*, pp 4270–4279
- Beasley A, Mason W (2015) Emotional states vs. emotional words in social media. In: *WebSci*, pp 1–10
- Jaidka K, Giorgi S, Schwartz HA, Kern ML, Ungar LH, Eichstaedt JC (2020) Estimating geographic subjective well-being from Twitter: a comparison of dictionary and data-driven language methods. *Proc Natl Acad Sci* 117(19):10165–10171
- Ribeiro FN, Araújo M, Gonçalves P, Gonçalves MA, Benevenuto F (2016) Sentibench—a benchmark comparison of state-of-the-practice sentiment analysis methods. *EPJ Data Sci* 5(1):23

16. Elshahr H, Gallé M (2019) To annotate or not? Predicting performance drop under domain shift. In: EMNLP-IJCNLP, pp 2163–2173
17. Troiano E, Padó S, Klinger R (2019) Crowdsourcing and validating event-focused emotion corpora for German and English. In: ACL, pp 4005–4011
18. Elayan S, Sykora M, Shankardass K, Robertson C, Feick R, Shaughnessy K, Haydn L, Jackson T (2020) The stresscapes ontology system: detecting and measuring stress on social media. In: ECSM, pp 74–82
19. Kjell ON, Sikström S, Kjell K, Schwartz HA (2022) Natural language analyzed with AI-based transformers predict traditional subjective well-being measures approaching the theoretical upper limits in accuracy. *Sci Rep* 12(1):3918
20. Lykousas N, Patsakis C, Kaltenbrunner A, Gómez V (2019) Sharing emotions at scale: the Vent dataset. In: ICWSM, vol 13, pp 611–619
21. Alvarez-Gonzalez N, Kaltenbrunner A, Gómez V (2021) Uncovering the limits of text-based emotion detection. In: Findings in EMNLP, pp 2560–2583
22. Plaza-del-Arco FM, Martín-Valdivia MT, Ureña-López LA, Mitkov R (2020) Improved emotion recognition in Spanish social media through incorporation of lexical knowledge. *Future Gener Comput Syst* 110:1000–1008
23. Barbieri F, Camacho-Collados J, Espinosa Anke L, Neves L (2020) TweetEval: unified benchmark and comparative evaluation for tweet classification. In: Findings in EMNLP, pp 1644–1650
24. Demszky D, Movshovitz-Attias D, Ko J, Cowen A, Nemade G, Ravi S (2020) Goemotions: a dataset of fine-grained emotions. In: ACL, pp 4040–4054
25. Nguyen DQ, Vu T, Tuan Nguyen A (2020) BERTweet: a pre-trained language model for English tweets. In: EMNLP: system demonstrations, pp 9–14
26. Felbo B, Mislove A, Søgaard A, Rahwan I, Lehmann S (2017) Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. In: EMNLP, pp 1615–1625
27. Tang D, Wei F, Yang N, Zhou M, Liu T, Qin B (2014) Learning sentiment-specific word embedding for Twitter sentiment classification. In: ACL, pp 1555–1565
28. Ke P, Ji H, Liu S, Zhu X, Huang M (2020) SentiLARE: sentiment-aware language representation learning with linguistic knowledge. In: EMNLP, pp 6975–6988
29. Gururangan S, Marasović A, Swayamdipta S, Lo K, Beltagy I, Downey D, Smith NA (2020) Don't stop pretraining: adapt language models to domains and tasks. In: ACL, pp 8342–8360
30. Levine Y, Lenz B, Lieber O, Abend O, Leyton-Brown K, Tenenholz M, Shoham Y (2021) PMI-Masking: principled masking of correlated spans. In: ICLR
31. Sosea T, Caragea C (2021) eMLM: a new pre-training objective for emotion related tasks. In: ACL, pp 286–293
32. Kumar A, Raghunathan A, Jones RM, Ma T, Liang P (2022) Fine-tuning can distort pretrained features and underperform out-of-distribution. In: ICLR
33. Howard J, Ruder S (2018) Universal language model fine-tuning for text classification. In: ACL, pp 328–339
34. Lamprinis S, Bianchi F, Hardt D, Hovy D (2021) Universal joy a data set and results for classifying emotions across languages. In: Proceedings of the eleventh workshop on computational approaches to subjectivity, sentiment and social media analysis, pp 62–75
35. Malko A, Paris C, Duenser A, Kangas M, Molla D, Sparks R, Wan S (2021) Demonstrating the reliability of self-annotated emotion data. In: Proceedings of the seventh workshop on computational linguistics and clinical psychology: improving access, pp 45–54
36. Ekman P (1999) Basic emotions. In: Handbook of cognition and emotion, pp 45–60
37. Scherer KR (2005) What are emotions? And how can they be measured? *Soc Sci Inf* 44(4):695–729
38. Pennebaker JW, Boyd RL, Jordan K, Blackburn K (2015) The development and psychometric properties of liwc2015. Technical report
39. Scherer KR, Wallbott HG (1994) Evidence for universality and cultural variation of differential emotion response patterning. *J Pers Soc Psychol* 66(2):310
40. Mohammad S (2012) # emotional tweets. In: SemEval, pp 246–255
41. Mohammad S, Bravo-Marquez F, Salameh M, Kiritchenko S (2018) SemEval-2018 task 1: Affect in tweets. In: SemEval, pp 1–17
42. Liu Y, Ott M, Goyal N, Du J, Joshi M, Chen D, Levy O, Lewis M, Zettlemoyer L, Stoyanov V (2019) Roberta: a robustly optimized bert pretraining approach. arXiv preprint. [arXiv:1907.11692](https://arxiv.org/abs/1907.11692)
43. Mohammad S, Turney P (2010) Emotions evoked by common words and phrases: using Mechanical Turk to create an emotion lexicon. In: NAACL HLT workshop, pp 26–34
44. Mohammad SM, Turney PD (2013) Crowdsourcing a word-emotion association lexicon. *Comput Intell* 29(3):436–465
45. Loshchilov I, Hutter F (2019) Decoupled weight decay regularization. In: International conference on learning representations. <https://openreview.net/forum?id=Bkg6RiCqY7>
46. Wortsman M, Ilharco G, Gadre SY, Roelofs R, Gontijo-Lopes R, Morcos AS, Namkoong H, Farhadi A, Carmon Y, Kornblith S et al (2022) Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In: International conference on machine learning, pp 23965–23998. PMLR
47. Gunel B, Du J, Conneau A, Stoyanov V (2021) Supervised contrastive learning for pre-trained language model fine-tuning. In: International conference on learning representations. <https://openreview.net/forum?id=cu7lUjOhujH>
48. Boyd RL, Ashokkumar A, Seraj S, Pennebaker JW The development and psychometric properties of liwc-22
49. Wang S, Manning C (2012) Baselines and bigrams: simple, good sentiment and topic classification. In: ACL, pp 90–94
50. Maiya AS (2022) ktrain: a low-code library for augmented machine learning. *J Mach Learn Res* 23(158):1–6
51. Hutto C, Gilbert E (2014) Vader: a parsimonious rule-based model for sentiment analysis of social media text. In: Proceedings of the international AAAI conference on web and social media, vol 8, pp 216–225
52. Rathje S, Mirea D-M, Sucholutsky I, Marjeh R, Robertson C, Van Bavel JJ (2023) GPT is an effective tool for multilingual psychological text analysis. PsyArXiv preprint <https://osf.io/preprints/psyarxiv/sekf5/>
53. Ziems C, Held W, Shaikh O, Chen J, Zhang Z, Yang D (2023) Can large language models transform computational social science? arXiv preprint. [arXiv:2305.03514](https://arxiv.org/abs/2305.03514)
54. Jacovi A, Caciularu A, Goldman O, Goldberg Y (2023) Stop uploading test data in plain text: Practical strategies for mitigating data contamination by evaluation benchmarks. arXiv preprint. [arXiv:2305.10160](https://arxiv.org/abs/2305.10160)

55. Sap M, Le Bras R, Fried D, Choi Y (2022) Neural theory-of-mind? On the limits of social intelligence in large LMs. In: Proceedings of the 2022 conference on empirical methods in natural language processing, pp 3762–3780
56. Ribeiro MT, Singh S, Guestrin C (2016) “Why should I trust you?”: Explaining the predictions of any classifier. In: KDD, pp 1135–1144
57. Watson D (2000) *Mood and Temperament*. Guilford, New York
58. Posner J, Russell JA, Peterson BS (2005) The circumplex model of affect: an integrative approach to affective neuroscience, cognitive development, and psychopathology. *Dev Psychopathol* 17(3):715–734
59. Pellert M, Schweighofer S, Garcia D (2020) The individual dynamics of affective expression on social media. *EPJ Data Sci* 9(1):1
60. Eichstaedt JC, Weidman AC (2020) Tracking fluctuations in psychological states using social media language: a case study of weekly emotion. *Eur J Pers* 34(5):845–858
61. Mohammad SM (2022) Ethics sheet for automatic emotion recognition and sentiment analysis. *Comput Linguist* 48(2):239–278
62. Chung CK, Pennebaker JW (2012) Linguistic inquiry and word count (liwc): pronounced “luke,”... and other useful facts. In: *Applied natural language processing: identification, investigation and resolution*, pp 206–229. IGI Global. <https://doi.org/10.4018/978-1-60960-741-8.ch012>

Publisher’s Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

Submit your next manuscript at ▶ [springeropen.com](https://www.springeropen.com)
