

RESEARCH

# Explainable depression symptom detection in social media



Eliseo Bao<sup>\*†</sup> , Anxo Pérez<sup>†</sup> and Javier Parapar<sup>†</sup>

## Abstract

Users of social platforms often perceive these sites as supportive spaces to post about their mental health issues. Those conversations contain important traces about individuals' health risks. Recently, researchers have exploited this online information to construct mental health detection models, which aim to identify users at risk on platforms like Twitter, Reddit or Facebook. Most of these models are focused on achieving good classification results, ignoring the explainability and interpretability of the decisions. Recent research has pointed out the importance of using clinical markers, such as the use of symptoms, to improve trust in the computational models by health professionals. In this paper, we introduce transformer-based architectures designed to detect and explain the appearance of depressive symptom markers in user-generated content from social media. We present two approaches: (i) train a model to classify, and another one to explain the classifier's decision separately and (ii) unify the two tasks simultaneously within a single model. Additionally, for this latter manner, we also investigated the performance of recent conversational Large Language Models (LLMs) utilizing both in-context learning and finetuning. Our models provide natural language explanations, aligning with validated symptoms, thus enabling clinicians to interpret the decisions more effectively. We evaluate our approaches using recent symptom-focused datasets, using both offline metrics and expert-in-the-loop evaluations to assess the quality of our models' explanations. Our findings demonstrate that it is possible to achieve good classification results while generating interpretable symptom-based explanations.

**Keywords:** Explainability, Interpretability, Depression detection, Social media

## Introduction

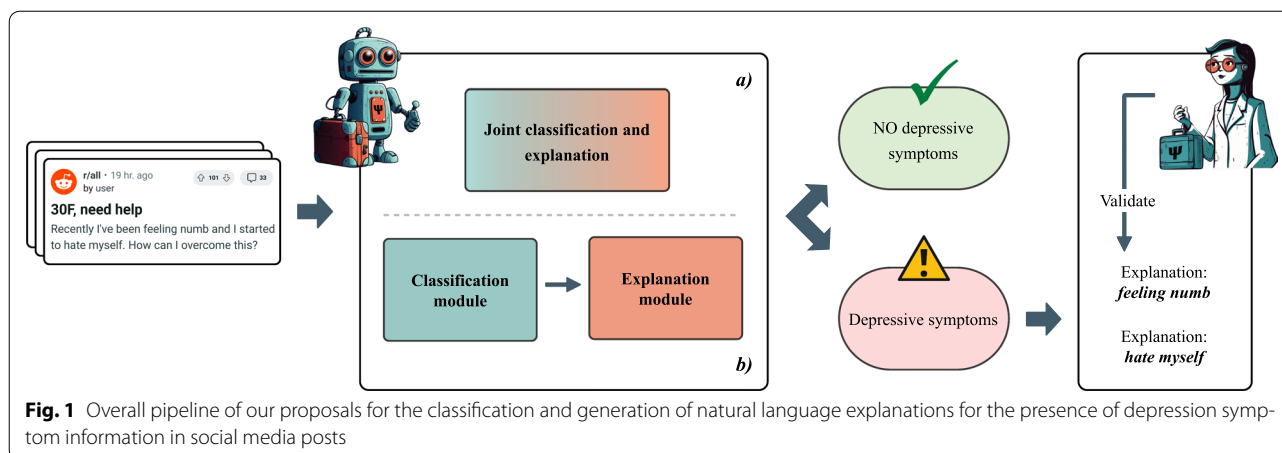
Mental health is a crucial component of overall health and well-being [1]. The most recent figures estimate mental disorders prevalence in adults over 20% [2]. The World Health Organization (WHO) estimates that approximately 332 million people globally are affected by depression [3]. Early intervention in mental disorders is crucial in mitigating their impact, especially for young individuals [4]. However, due to the stigma associated with depression, more than 60% of individuals with the condition do not seek professional support [5]. To address this problem, computational researchers have a

growing interest in assisting in the early detection and diagnosis of depression, thereby mitigating its societal impact [6].

In this scenario, researchers have found that the writings posted by individuals on social media platforms are valuable evidence for looking for early signs of depression [7–12]. Individuals experiencing depression find comfort in expressing their thoughts and emotions on these platforms, motivated by factors such as privacy or anonymity [13, 14]. Consequently, social media provides a complementary opportunity to access valuable information about individuals' state of mind beyond traditional professional therapy. The combined use of computational linguistic techniques and the vast amount of data from social networking has led to significant advancements in detecting signs of depression [15]. The field's critical nature motivated much effort in creating curated experimental benchmarks [16, 17], which allowed the development and evaluation of many new predictive models.

<sup>†</sup>Eliseo Bao, Anxo Pérez and Javier Parapar have contributed equally to this work.

\*Correspondence: [eliseo.bao@udc.es](mailto:eliseo.bao@udc.es)  
Information Retrieval Lab (IRLab), Centro de Investigación en Tecnoloxías da Información e da Comunicación (CITIC), Campus de Elviña, 15071 A Coruña, Galicia, Spain



**Fig. 1** Overall pipeline of our proposals for the classification and generation of natural language explanations for the presence of depression symptom information in social media posts

Traditional efforts used engineered features such as word counts, posting activity or emotion levels [7–9]. Recently, and due to the rise of transformer-based language models, many researchers used these deep learning models as classifiers to identify users at risk of depression and similar disorders in online environments [10–12].

However, researchers in this field do not plan to replace mental health professionals but rather offer support to their work. Licensed clinicians play a crucial role in validating the predictions made by computational models and taking appropriate actions with individuals when necessary. Computational models may only be used carefully to extend those professionals’ reach and facilitate their workflow. The existing models, however, present many limitations for achieving that goal [18]. One significant barrier is their limited ability to explain their predictions. Reliable interpretation of models’ decisions is mandatory for professionals to understand and trust these models and use them in their daily work [19]. One way to pursue that is by designing new models incorporating trustworthy and reliable explanations [20]. Following that path, recent research has explored the utilization of symptoms collected from validated clinical questionnaires. Most of these proposals, in the field of depression, used the markers from the Beck Depression Inventory-II (BDI-II) [21] or the 9-Question Patient Health Questionnaire (PHQ-9) [22], which encompass a range of depressive symptoms such as irritability, pessimism or sleep problems. The utilization of such symptom markers has been shown to improve the explainability, generalization and overall performance of depression detection models [23–27].

With that motivation, we aim to develop models that categorise whether or not social media posts exhibit markers of validated depressive symptoms. Accurately detecting symptom information along a user’s vast

amount of writing is the first step in developing explainable depression detection models. We go beyond classifying posts for depressive symptoms by providing an explanation for the decisions. It is important to clarify that our paper aims not to diagnose depressive disorder, but rather to identify markers indicative of potential symptoms. The diagnosis of depressive disorder, which is the ultimately aim and scope of psychologists in this domain, relies on clinical factors such as the presence of clinical symptoms and their temporality [28]. For this reason, this work seeks to aid in this diagnostic process by generating explanations of possible depressive symptomatology.

In this context, the terms ‘interpretability’ and ‘explainability’ can be challenging to define in our specific context, given the variations in their use across existing literature. Some authors use these terms interchangeably, referring to the ability to explain or present systems in a manner understandable to humans [29]. However, other authors consider them as distinct concepts. In this perspective, interpretability relates to the system’s capacity to be understood by humans, while explainability encompasses being true to reality [30]. In our work, we adopt a simple definition: interpretability/explainability refers to the extent to which humans can understand the reasons behind a decision [31].

In this paper, we introduce a text-to-text pipeline designed to achieve two main objectives: first, to classify the relevance of social media publications to depressive symptoms, and second, to provide explanations for these classifications. To implement this pipeline, we explore the effectiveness of state-of-the-art transformer-based models. Figure 1 illustrates the two approaches considered: Part a) employs text-to-text models that perform classification and explanation simultaneously. Part b) corresponds with two-step approaches, utilizing

separate models for classification and explanation. Additionally, we evaluate the capabilities of Large Language Models (LLMs) using recent general purpose conversational models such as GPT-3.5 Turbo [32], Vicuna-13B [33] and Mistral-7B-Instruct [34] within a few-shot in-context learning approach. We also investigate how fine-tuned conversational LLMs perform for this tasks using domain-specific models such as MentalLLaMA [35] and MentalMistral.

We utilize two datasets, PsySym [25] and BDI-Sen [36], for model training and evaluation. These datasets consist of sentences from social media posts associated with depressive symptoms. As we will discuss later, these sentences can be used as explanations for a post to explain its relevance to a symptom. To assess the quality of the generated explanations, we employ a combination of offline and expert-in-the-loop metrics. The experimental results demonstrate the effectiveness of our methods. By prioritizing explainability, our approach bridges the gap between automated predictions and human understanding, facilitating more informed clinical decision-making.<sup>1</sup>

Our study aims to address the following research questions:

- RQ1** Can we train transformer-based models to accurately classify the presence of depressive symptoms in social media posts, while providing explanations for their decisions?
- RQ2** How many training examples, i.e. hand-labelled posts by domain experts, are needed to generate good explanations?
- RQ3** Does a unified model, designed for both classification and explanation, significantly differ in performance from using two separated models trained for each task?
- RQ4** How do recent conversational LLMs perform in detecting and explaining the presence of depressive symptoms in social media content?

### Related work

Social media data has gained increasing attention for developing depression detection models in recent years [37–41]. The evaluation results of those methods show promising accuracy numbers [10–12, 15, 42]. Such platforms offer an opportunity to identify disorders at an early stage [43], a step that is crucial to reduce negative impacts and their associated costs [4, 44].

However, depression detection field still needs to overcome the cross-cutting problem of many machine learning applications: explainability. There are two common

approaches when explaining models' decisions. On the one hand, we may use intrinsically explainable models, such as decision trees or linear models [45]. On the other hand, we can opt for model-agnostic, typically post hoc, explainability methods [46, 47], which can be applied to any supervised machine learning model, regardless of its architecture. Moreover, the explanations can be local or global. Local explanations focus on individual decisions, allowing users to understand why the model produces a particular decision [48, 49]. These local explanations, such as LIME (Local Interpretable Model-Agnostic Explanations) [49] and SHAP (Shapley Additive exPlanations) [50], have already been used in the social psychology field [51, 52]. Alternatively, global explanations consider the whole machine learning model behaviour and its predictions altogether [53–55]. Most traditional detection methods used engineered features (e.g., word counts or sentiment analysis) and based their global explanations on feature importance metrics [9, 15]. These measures quantify the impact of each feature on the predictions, providing insights into the relative importance of each one.

With the rise of transformer-based models [56], researchers have leveraged these models for sentence classification tasks to detect signs of depression in social media users, achieving impressive results in experimental benchmarks [16, 17]. The encoder-decoder architecture of transformers effectively captures contextual information from input sequences and generates corresponding output sequences. However, maintaining the contextual integrity of each token in extensive texts presents a challenge. The attention mechanism tackles this by allocating weights to tokens relative to all others, allowing the model to focus on relevant parts of the input. Prior works have explored the attention mechanism's as an explainability tool [57–59]. Despite its utility, attention mechanisms also have limitations, particularly in interpreting the weights assigned to input features. There is an ongoing discussion in the community about the utility of attention weights for interpretability purposes, with some research advocating its benefits [60, 61], while others point out its constraints [62, 63].

More recently, an alternative strategy for building interpretable models has emerged: the use of generative natural language explanations. This technique has several benefits: (i) they are readily comprehensible to end users, (ii) human annotators can more easily work with natural language, simplifying data collection, and (iii) it may be feasible to extract natural language explanations from large datasets of domain-expert data, a promising prospect for future research [64].

Linked to the recent breakthrough in the field with the advent of modern LLMs, Yang et al. [65] explored

<sup>1</sup> Our implementation is available at: <https://gitlab.irlab.org/irlab/explainable-depression-symptom-detection-social-media>

their application in zero-shot/few-shot mental health analysis, assessing the influence of various emotion-enhanced prompts. This study investigates the potential of LLMs for explainable mental health assessment, elucidating the predictions through Chain-of-Thought (CoT) prompting [66]. While the study reveals that ChatGPT is capable of generating explanations at a human-level, it also notes that the approach has limitations, including unstable predictions and inaccurate reasoning. Following this study, Yang et al. [35] also introduced an interpretable mental health instruction dataset, constructed with ChatGPT-generated explanations. The authors used this data to train MentalLLaMA, a new open-source LLM that focuses on interpretable mental health analysis with conversational and reasoning skills. MentalLLaMA covers not only depression but also other related mental health conditions such as stress or suicide. In contrast, our study adopts a symptom-based approach, focusing on clinical markers for a nuanced analysis of depression. We include MentalLLaMA in our evaluations as a baseline, examining its performance against general-purpose models and its efficacy in symptom-level analysis. Our proposed method, MentalMistral, is another LLM fine-tuned for depression analysis but specifically at the symptom level. Notably, MentalMistral differs from MentalLLaMA in its data and explanation generation approach. While MentalLLaMA generates free-form explanations, MentalMistral is fine-tuned to generate explanations derived directly from segments of the input text, ensuring that its insights are firmly rooted in the original user-generated and human-annotated content. This approach ensures that our explanations are grounded in the text and provides a direct link between the user's words and the model's output.

Our goal is to cross the bridge by following prior efforts in training text-to-text models in order to produce natural language explanations for the depression detection problem by grounding all the explanations in clinical symptoms following the BDI-II questionnaire. As the explanations are formulated in human-readable text, our method makes the model's rationale behind the predictions transparent to clinicians. This level of interpretability is of critical importance to making informed decisions for depression detection [18], promoting both the efficacy and the trustworthiness of our models.

### Our proposal

In this section, we describe our proposal for generating explainable decisions in the context of detecting depressive symptoms in social media posts. The task is defined as follows: given a user post, the model classifies whether

it is indicative (positive) or non-indicative (negative) of any depressive symptom. A positive classification implies that the post reflects information for the user on one or more of the 21 symptoms described in the BDI-II questionnaire. Moreover, our objective extends to providing natural language explanations for any positive classification, showing the reasoning behind the model's decision. We undertake this task in two different manners: either as a single output (i.e. train together classification plus explanation), or in two separate steps (i.e. initial classification followed by post hoc explanation). These two scenarios are illustrated in Fig. 1. Part (a) is representative of the single-step approach, explained in "Single-step" section, while part (b) corresponds to the two-step approach, detailed in "Two-steps pipeline" section. If the pipeline identifies a post as having potential symptom risk, it provides corresponding explanations. Looking at the Figure, we can see that the explanations for the decision are 'feeling numb' and 'hate myself'. These generated explanations are designed to be reviewed and validated by health professionals.

### Explanations

Natural language explanation generation methods can be either extractive or abstractive [67]. In the *extractive* case, the model is asked to point out the parts of the original text that led to the decision. Conversely, *abstractive explanations* involve the model to explain the reasons for the decision in a free format. Our approach focuses on extractive explanations, where the models highlights relevant text spans from the whole input text. This decision was driven by data availability and the advantage of this method in allowing quantitative analysis of explanation quality through overlap statistics with ground truth.

Furthermore, our models are designed to explain only the positive classifications. This decision comes from two considerations. First, explanations in negative cases (where no depressive symptom markers are detected) are less clinically relevant, as the primary focus is to identify positive indications of depression. Secondly, for extractive explanations, negatives cases would typically be a generic responses, such as "there is no evidence in the post about any depression symptom". Therefore, to optimize the utility and specificity of our models, we concentrate on explaining only the positive cases.

### Finetuning text-to-text models

Raffel et al. [68] proposed the idea of transforming text-related tasks to follow the sequence-to-sequence (seq2seq) format [69], which is referred to as the text-to-text framework. This diverges from conventional methods like BERT-based approaches, which train models to yield a probability distribution over predefined output



classes [70]. In contrast, text-to-text models are trained to generate textual sequences. Consequently, these models may generate an unexpected output, which is considered a prediction error.

In this context, the input text sequence, which we will call  $\tilde{x}$ , is defined as `<task_prefix>: <input_text>`, where `<task_prefix>` identifies the task to be performed by the model. On the other hand, the output text sequence  $\tilde{y}$  will be of the form `<target>`, where `<target>` corresponds to the desired output. If we take a sentiment analysis task as example, the sequence would be:  $\tilde{x} = \text{'sentiment analysis: terrible product'}$  and  $\tilde{y} = \text{'negative'}$ . The Text-to-Text Transfer Transformer (T5) model is one of the most popular models under this paradigm [68], and it achieved state-of-the-art results in many NLP tasks [71]. While there exist other text-to-text models, such as BART [72], prior research indicates that it performs worse than T5 in terms of explanation generation [73]. These models are pre-trained, providing us with a robust base that we can further finetune under either the single-step or two-steps approaches discussed below for generating explanations.

### Single-step

Narang et al. [74] explored how to teach text-to-text models to produce both classification and explanation. For that, they introduced an extension of T5 called WT5 ("Why T5?"). To use this method, the keyword "explain" is simply added to the input  $\tilde{x}$ , preceded by `<task_prefix>`. The target  $\tilde{y}$  is appended with the phrase `'explanation: <explanation>'`. Using this template, we adapted it to our task resulting in the new input/output format:

```

 $\tilde{x} = \text{'explain symptom post: <post_text>'}$ 
 $\tilde{y} = \text{'<target> [explanation: <explanation_1>]...}$ 
        $\text{... [explanation: <explanation_N>]'}$ 

```

The hard brackets denote potentially multiple explanation sentences. An example of input could be "explain symptom post: I absolutely hate myself (...) And I hate how I feel the need to burden other people with this. I am so whiny, so disgustingly insensitive (...)"<sup>2</sup> Its corresponding target output would be "positive explanation: I absolutely hate myself explanation: I am so whiny, so disgustingly insensitive". Thus, in one step, users' posts can be both classified and explained in an extractive manner according to whether they indicate depressive symptoms or not. Part a) in Fig. 1 illustrates

this approach, which we use to build two of our systems: **WT5** and **WBART**. Our WT5 system replicates the research of Narang et al. that we just briefly discussed, while our WBART system extends the same idea but to the BART (Bidirectional and Auto-Regressive Transformers) model [73].

While WT5 and WBART share the common objective of facilitating single-step classification and explanation, they diverge in their underlying architectures. WT5 builds upon the T5 model, whereas WBART utilizes the BART model. Both models adopt an encoder-decoder framework and are suitable for a wide range of sequence-to-sequence tasks. However, beyond considerations such as pretraining corpus, parameter initialization, and activation functions, the primary distinction between T5 and BART lies in their pretraining objectives. T5 implements a strategy where 15% of tokens in the input sequence are randomly dropped out. These dropped-out tokens are replaced by sentinel tokens, and each one has a unique token ID assigned. The model's objective is to predict these dropped-out spans, delimited by the sentinel tokens. In contrast, BART's training process involves corrupting documents and optimizing a reconstruction loss between the decoder's output and the original document. Unlike T5, BART allows for various document corruptions, including the complete loss of source information.

### Two-steps pipeline

Alternatively, rather than training a single model to execute both classification and explanation, we propose a two-steps pipeline. In this approach, one model is initially employed for classification, followed by the use of a different model for explanation. We develop three of our systems using this methodology: **T5 + T5**, **BERT + T5**, and **MBERT + T5**. This process is depicted in part b) of Fig. 1. Next, we define the stages of classification and explanation as they occur in this pipeline.

**Classification** In alignment with recent studies, we conduct experiments with various types of pre-trained language models to investigate their capabilities in terms of classification [24–26]. These models constitute the classification module and their sole responsibility is to determine whether a post is indicative of a depressive symptom. Firstly, we fine-tune a **T5** model with the aim of generating classification labels in a text format [68]. Secondly, we explore the utility of BERT-based architectures, utilizing the pre-trained **BERT** base uncased model [70]. Additionally, we finetuned MentalBERT (**MBERT**), a model purposely pre-trained for mental health applications using data compiled from various subreddits related to mental disorders [75].

<sup>2</sup> Full-length post is paraphrased and redacted for clarity and space reasons.

**Explanation** For the explanation module, we finetuned a T5 model responsible for explaining the evidence for those posts deemed positive by the classification models. As previously commented, we only trained this model to generate explanations for positive posts. Negative posts are discarded and do not enter in the second stage of the two-steps pipeline.

**In-context learning with LLMs**

In-context learning strategies allow guiding LLMs to specific behaviours within a given context without directly modifying their parameters [76]. In our study, we applied this technique to instruct recent conversational LLMs to perform our proposed single-step approach (*classify + explain*). Using specific guidelines<sup>3</sup>, we instructed the models to act as expert annotators responsible for detecting signs of depressive symptoms in user posts. Additionally, they must justify their decisions by quoting relevant spans from the original text. This approach aims to replicate the methodology discussed earlier, which outlined the process of finetuning text-to-text models for explanation.

We explored this strategy in two different ways: first, within a few-shot approach with general purpose models such as **Vicuna-13B** [33] **GPT-3.5 Turbo** [32] and **Mistral-7B-Instruct** [34]. Second, with domain-specific models that were previously finetuned such as **MentalLLaMA-chat-13B** [35] and **MentalMistral**. Both finetuned models can follow instructions to generate explanations for the predictions, but they differ in important ways. First, MentalMistral is a QLoRA [77]-finetuned version of the Mistral-7B-Instruct model that we implemented with the same symptom-based data as the rest of our models, whereas MentalLLaMA is based on LLaMA2-chat-13B [78]. Their strategies to identify depression also vary: MentalLLaMA focuses on assessing depression by classifying users into depressed and non-depressed categories, while MentalMistral uses clinical indicators to assess depression at the symptom level. Furthermore, the nature of their generated explanations are also different. While MentalLLaMA generates free-form explanations, all our methods, including MentalMistral, are specifically tuned to generate explanations from fragments of the original text.

**Experiments**

In this section, we describe the experiments conducted to evaluate the effectiveness of our approaches for symptom detection and explanation over social media content. We present the datasets utilized, followed by a description of the experimental configurations and training details.

<sup>3</sup> See B for details of the prompt strategy we adopted.

**Table 1 Statistics of the positive instances of our explainability dataset**

	BDI-Sen	PsySym
Number of posts	357	752
Number of explanations	546	764
Avg. number of expls. per post	1.53	1.02
Avg. length of post (in tokens)	127	515
Avg. length of explanation	13.76	13.44

Average length of posts/explanations are expressed as tokens

Subsequently, we explore the metrics employed to measure the performance of our methods, encompassing both offline evaluations and those involving human experts.

**Datasets**

In this work, we considered the symptoms of the BDI-II clinical questionnaire [21]. The BDI-II covers 21 recognized symptoms, such as pessimism, sleep problems or self-dislike. We have decided to align with the BDI-II symptoms by its widespread adoption in clinical practice [79] and its presence in prior literature concerning depression detection on the Internet [16, 23, 26, 41, 80]. We obtained the symptom evidence sentences from the BDI-Sen [36] and PsySym [25] resources. Both datasets consist of symptom-annotated sentences for depression sourced from the Reddit platform. BDI-Sen comprises relevant sentences related to the 21 symptoms covered in the BDI-II, while PsySym includes 14 main symptoms covered within the Diagnostic and Statistical Manual of Mental Disorders (DSM-5) [81]. Table 1 presents the statistics of the positive posts from the datasets. In addition to positive posts, we also included control instances in the experiments. We selected random control sentences from the datasets, totalling 1998 control posts. In "Ethical and environmental aspects" section, we comment on the anonymization and ethical use of the data.

The datasets comprise two different data sources. Thus, for our experiments, we can define five different settings for train and test sets by combining both resources. Initially, for both BDI-Sen and PsySym, we split their data 80-20 into training and test sets. Then, we fill these splits with control samples from PsySym with ratios 1:1 for training set and  $\approx 1:5$  for the test set.<sup>4</sup> Table 2 illustrates the five different settings we consider, where we test each dataset on itself, and also using as test the other one to evaluate if it can generalize correctly. In the last setting (*M-M*), we merge both datasets.

<sup>4</sup> We include more control sentences in the test set to replicate real scenarios, where there are many more negatives examples than positive ones.

**Table 2 Settings considered in our experiments by combining the BDI-Sen and PsySym datasets**

Setting	Training			Test		
	Dataset	⊕/⊖	Avg. expls.	Dataset	⊕/⊖	Avg. expls.
B-B	BDI-Sen	285/285	1.48	BDI-Sen	72/359	1.72
B-P	BDI-Sen	285/285	1.48	PsySym	151/753	1.02
P-P	PsySym	601/601	1.01	PsySym	151/753	1.02
P-B	PsySym	601/601	1.01	BDI-Sen	72/359	1.72
M-M	Mix	886/886	1.16	Mix	223/1112	1.25

**Experimental settings**

**Single-step**

For our single-step strategy, we finetuned two distinct pretrained text-to-text models: **WT5** and **WBART**. For the WT5 model, we followed the training protocol from Narang et al. [74], applying 40 epochs of fine-tuning to the `t5-large` model. This adaptation was for our *classify + explain* task, utilizing the AdaFactor optimizer with a constant learning rate of 0.001. The input and output sequence lengths were set to 2048 and 512 tokens, respectively. In the case of WBART, we fine-tuned the `bart-large` model, adhering to the default learning rate and hyperparameters. Due to computational limitations, the maximum lengths for the source and target sequences were reduced to half of those used in WT5.

Additionally, we employed the conversational LLMs previously mentioned (**Vicuna-13B**, **GPT-3.5**, **Mistral-Instruct**, **MentalLLaMA**, and **MentalMistral**) within the single-step framework. All our models were trained and evaluated using NVIDIA A100-SXM4 80GB GPUs. For GPT-3.5, we relied on API calls. All conversational LLMs were directed following the in-context learning strategy described in "In-context Learning with LLMs" section, using a balanced set of 30 positive and 30 control samples randomly chosen.

**Two-steps**

In this approach, we first used two BERT-based models for text classification: BERT base, and MBERT base, following the existing implementations from the HuggingFace library without any additional hyperparameter tuning. The learning rate was  $2e^{-5}$  during 20 epochs and a batch size of 32. We also finetuned a T5 on its `t5-large` configuration for text classification by generating labels in textual form. For explanations, we used another `t5-large` to explain the symptoms presence. The parameters used to finetune the T5 models were the same as those used for WT5, with same sequence lengths. As a result, we constructed three different

variants within this approach: **T5 + T5**, **BERT + T5**, and **MBERT + T5**.

**Evaluation**

**Classification**

To evaluate the classification performance of our systems, we consider micro-averaged F1 due to the unbalanced nature of the datasets. By using micro-F1, we account for the importance of each instance, giving more weight to the minority class (positive). Additionally, we include the number of true positives (TPs) to gain insight into the actual number of correctly identified positive instances. This information is particularly significant since our systems only generate explanations for the positive samples.

**Explanation**

LLMs can generate text unsupported by the input and produce inaccurate responses, a phenomenon called "hallucination" [82]. This is a magnified risk considering the sensitive factor of the mental health domain [65]. For this reason, we force extractive explanations and instruct the models to reflect on a given text and provide a decision-supporting fragment. We ensure that our systems are less prone to hallucination than in a free-form explanation setup. In the clinical domain, where the validity and truthfulness of the explanations extracted are crucial [18], the generated texts' reliability is essential. To ensure the trustworthiness of the generated explanations, we complement classical offline metrics with *expert-in-the-loop* evaluation performed with three domain experts.

**Offline**

We used three offline metrics: ROUGE-L-F1 [83], Corpus BLEU [84], and Token F1 [85], each offering different evaluation perspectives. All these metrics compare the generated explanations with the text reference considered as golden truth. ROUGE emphasizes content overlap, ensuring that the generated output captures essential information from the references. On the other hand, BLEU focuses on fluency and adequacy, assessing the linguistic quality and alignment with references. Since these metrics effectively measure the quality of the

**Table 3 Results of our methods in the different proposed settings**

Setting	Metric	Finetuned PLMs					In-context convers. LLMs				
		Single-step		Two-steps			Few-shot			Finetuned	
		WT5	WBART	T5 + T5	BERT + T5	MBERT + T5	Vicuna-13B	GPT-3.5	Mistral-Instruct	MentalLLaMA	MentalMistral
<i>B-B</i>	F1	0.91	0.88	0.93	<b>0.95</b>	0.92	0.73	0.82	0.64	0.88	0.89
	TPs	<b>68</b>	<b>68</b>	67	58	65	41	62	68	51	51
	ROUGE	0.62	0.60	0.69	<b>0.75</b>	0.72	0.42	0.55	0.39	0.39	0.44
	BLEU	0.53	0.53	0.55	<b>0.61</b>	0.58	0.22	0.42	0.18	0.19	0.26
	TF1	0.48	0.48	0.54	<b>0.60</b>	0.58	0.34	0.45	0.36	0.37	0.31
<i>B-P</i>	F1	0.87	0.86	0.86	0.91	0.92	0.78	<b>0.87</b>	0.66	0.92	<b>0.93</b>
	TPs	125	134	79	79	144	140	<b>151</b>	150	131	141
	ROUGE	0.31	0.22	<b>0.35</b>	0.34	0.29	0.14	0.17	0.12	0.12	0.17
	BLEU	0.22	0.15	<b>0.23</b>	0.22	0.18	0.06	0.08	0.07	0.07	0.08
	TF1	0.10	0.10	<b>0.12</b>	<b>0.12</b>	0.10	0.11	0.10	0.12	0.17	0.09
<i>P-P</i>	F1	<b>0.98</b>	<b>0.98</b>	<b>0.98</b>	0.97	<b>0.98</b>	0.69	0.94	0.80	0.93	0.96
	TPs	143	149	145	150	149	150	<b>151</b>	150	147	144
	ROUGE	0.53	<b>0.56</b>	0.45	0.47	0.45	0.15	0.21	0.16	0.16	0.17
	BLEU	0.43	<b>0.53</b>	0.34	0.36	0.35	0.08	0.11	0.12	0.08	0.08
	TF1	0.47	<b>0.51</b>	0.38	0.38	0.38	0.09	0.11	0.08	0.09	0.08
<i>P-B</i>	F1	0.89	0.89	0.89	<b>0.91</b>	0.90	0.62	0.85	0.73	0.85	0.84
	TPs	35	35	40	48	42	<b>53</b>	49	47	51	29
	ROUGE	0.61	<b>0.71</b>	0.61	0.57	0.61	0.36	0.53	0.48	0.47	0.48
	BLEU	0.61	<b>0.68</b>	0.60	0.57	0.58	0.19	0.34	0.34	0.33	0.31
	TF1	0.46	<b>0.51</b>	0.45	0.42	0.45	0.23	0.36	0.32	0.35	0.30
<i>M-M</i>	F1	0.95	0.94	0.95	0.95	<b>0.96</b>	0.59	0.91	0.45	0.77	0.88
	TPs	209	211	206	211	211	195	200	<b>215</b>	191	191
	ROUGE	<b>0.57</b>	0.53	0.54	0.54	0.54	0.22	0.28	0.25	0.23	0.24
	BLEU	<b>0.54</b>	0.47	0.51	0.51	0.51	0.10	0.17	0.15	0.09	0.19
	TF1	<b>0.50</b>	0.48	0.46	0.46	0.45	0.14	0.19	0.17	0.17	0.14

Bold values represent the best performance for each setting and metric, with higher values indicating better results

F1, TPs, ROUGE, BLEU and TF1 denote micro-F1, number of true positives, ROUGE-L-F1, Corpus BLEU and Token F1. More details about these metrics can be found in "Evaluation" section. All metrics are in the range 0–1, except for TPs, which are an integer

generated text by comparing it against reference hypotheses, ROUGE and BLEU can evaluate the quality of our explanations while still allowing some paraphrasing. This flexibility can be beneficial when the explanation text slightly differs from the ground truth yet conveys the same meaning. Additionally, Token F1 measures which input tokens labelled as an explanation in the ground truth are present in the generated one. Thus, Token F1 score is computed through a token-by-token analysis, identifying if the tokens generated are the same as the references.

**Expert-in-the-loop**

We also performed an online evaluation using three domain experts to assess the practical utility and clinical relevance of the generated explanations. The evaluators consisted of two psychologists and a speech therapist. To ensure the consistency of judgments, we

organized training sessions with the evaluators during which we discussed the criteria for relevance and established a consensus on samples of both positive and negative explanations. An explanation was considered relevant if it successfully provided evidence for the presence of a depressive symptom in the post. For a detailed view of the instructions provided to the assessor please see "Appendix 1: Assessment guidelines" section. We provided the experts with explanations generated by our WT5 model variant under the *M-M* setting (see Table 3), leading to a total of 209 explanations being evaluated. We chose this setting because it includes both datasets and selected WT5 since it was the best performing model in this scenario.



**Table 4 Metric averages per method and proposed setting**

<b>(a) Per method.</b>					
Method	F1	ROUGE	BLEU	TF1	
WTS	0.91	<b>0.53</b>	<b>0.47</b>	0.40	
WBART	0.91	0.52	<b>0.47</b>	<b>0.42</b>	
T5 + T5	0.92	<b>0.53</b>	0.45	0.39	
BERT + T5	<b>0.94</b>	<b>0.53</b>	0.45	0.40	
MBERT + T5	<b>0.94</b>	0.52	0.44	0.39	
Vicuna-13B	0.68	0.26	0.13	0.18	
GPT-3.5	0.88	0.25	0.22	0.24	
Mistral-Instruct	0.66	0.28	0.17	0.21	
MentalLLaMA	0.87	0.24	0.15	0.23	
MentalMistral	0.90	0.30	0.18	0.18	
<b>(b) Per setting.</b>					
Setting	F1	TPs	ROUGE	BLEU	TF1
<i>B-B</i>	0.85	60	<b>0.56</b>	0.41	<b>0.45</b>
<i>B-P</i>	0.86	127	0.22	0.14	0.11
<i>P-P</i>	<b>0.92</b>	148	0.33	0.25	0.26
<i>P-B</i>	0.84	43	0.54	<b>0.46</b>	0.39
<i>M-M</i>	0.84	204	0.38	0.32	0.32

Bold values represent the best performance for each setting and metric, with higher values indicating better results

F1, TPs, ROUGE, BLEU and TF1 denote micro-F1, number of true positives, ROUGE-L-F1, Corpus BLEU and Token F1. More details about these metrics can be found in "Evaluation" section. All metrics are in the range 0–1, except for TPs, which are an integer

## Results and discussion

In this section, we present the results of the experiments designed to address the four research questions posed in the introduction.

### RQ1: can we train transformers models to detect and explain depressive symptoms?

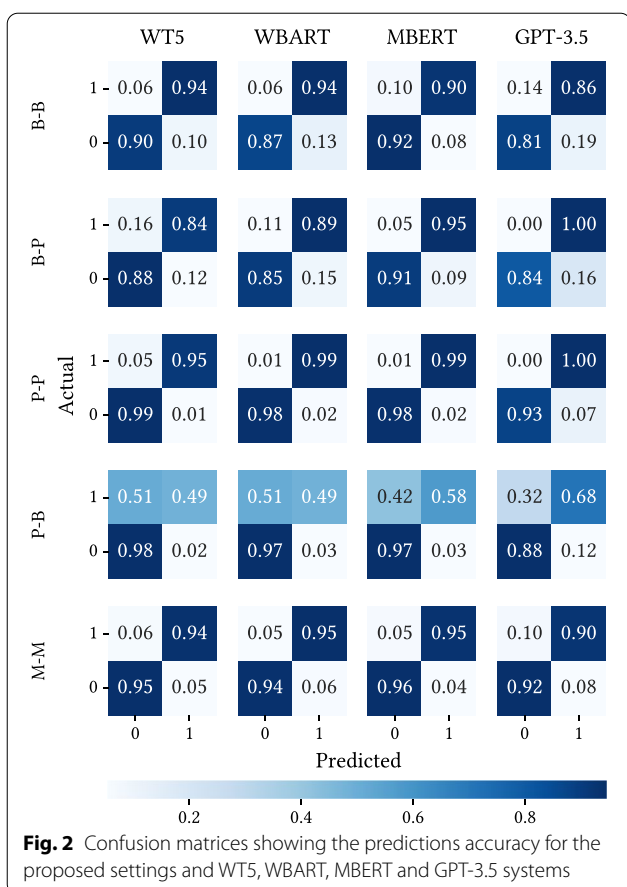
#### Symptoms classification

The results presented in Table 3 confirm that it is possible to train transformer architectures to detect traces of depression symptoms in social media posts. Across all the different settings, we observe that the worst performing finetuned PLM achieves a classification accuracy of 0.86, which is the case for the WBART and T5 + T5 methods for the *B-P* setting. On the other hand, in-context conversational LLMs have more difficulties in detecting the presence of depression symptoms within a few-shot approach, and in some cases their performance drops to an accuracy of 0.45, which is the case of the Mistral-Instruct method for *M-M* setting. However, conversational LLMs show the ability to learn this task easily when finetuned. Looking at the MentalMistral method, which is a finetuned version of Mistal-Instruct at the symptom level, we can see that it improves its classification accuracy from 0.45 to 0.88 and from 0.66 to 0.93 for the *M-M* and *B-P* settings, respectively. It is particularly

noteworthy that for the latter setting, MentalMistral is the method that best detects the symptoms of depression. Comparatively, MentalLLaMA, another fine-tuned LLM designed for broader mental health conditions beyond depression, shows slightly inferior performance compared to MentalMistral.

To better understand how well the different methods detect depression symptoms across settings, Table 4a contains a column with the F1 average across every setting. The values are consistent with the analysis just presented. On the one hand, all finetuned PLMs have a similar performance in terms of detection, which is slightly higher in the case of the two-step pipeline. It is also clear that the in-context conversational LLMs are a step behind, with the Mental-Instruct method being the worst performer (average F1 of 0.66), but that they improve significantly when finetuned, either at the symptom level (MentalMistral) or not (MentalLLaMA).

We also show in Table 4b a comparison of the performance of our methods across all the dataset settings. We observe the best classification figures in the *P-P* setting (training and testing on the PsySym dataset), with all finetuned PLMs achieving F1 values above 0.97 and all methods averaging an F1 value of 0.92. This numbers also provide insights into how the models generalise when tested on a different dataset (the *B-P* and *P-B*



settings). We observe bad generalisation in the case of the *P-B* setting. For instance, looking at specific models in Table 3 the T5 + MBERT model achieved an F1 of 0.98 when trained and tested on the PsySym dataset (*P-P* setting). Meanwhile, its performance dropped to 0.90 when tested on the BDI-Sen dataset (*P-B* setting). However, in the *M-M* setting, the trained models obtained an F1 greater than 0.94.

To further analyse this generalisation problem, we illustrate the distribution of true and false predictions of the systems in Fig. 2.

As observed in the confusion matrices,<sup>5</sup> the models have a high ratio of true positives, with the majority identifying over 90% of the positive instances. Furthermore, the ratio of false negatives is remarkably small. However, we can again observe significant deviation in these trends in the matrix for the *P-B* setting. Here we can see the reason for that problem: the PsySym dataset covers 14 symptoms of depression, while BDI-Sen includes all 21 symptoms of the BDI-II. In the *P-B* setting, not all symptoms of BDI-Sen have been seen during training,

leading to a substantially higher false positive number. This finding highlights the potential generalisation issues these models may encounter regarding the data used for training. False negatives are a crucial metric for clinical applications. In clinical practice, false negatives carry great significance. Failing to identify an individual at risk has far more severe implications than falsely identifying a healthy person (false positives).

**Symptoms explanation**

Regarding explanation quality, offline metrics and human evaluation results indicate promising results. Once again, finetuned PLMs outperform the most recent conversational instructed LLMs, which can be seen in columns ROUGE, BLEU and TF1 of Table 4a. It is interesting to note that, in contrast to symptom detection, modern LLMs do not significantly improve the quality of their explanations after finetuning. This idea is consistent with the results presented for RQ2, where it is shown that the quantity and quality of examples is much more critical for explanation than for classification supporting the idea that generating explanations is more challenging than the classification task. As previously stated ("*Explanations*" section), it is crucial to bear in mind that we generate explanations only for positive cases. Hence, the explanation quality numbers have to be jointly considered with the number of true positives (refer to the TPs rows in Table 3). When considering the offline metrics (ROUGE, BLEU, and TF1), the fairest comparison can be made in the *M-M* setting, since it is the largest setting and provides a similar number of TPs cases for all models. Here, WT5 emerges as the top performer, achieving a ROUGE-L F1-score of 0.57. The two-step models also perform closely to the single-step models (WT5 and WBART) in nearly all settings. In the *B-B* setting, where the numbers of true positives are also roughly similar, they outperform the single-step models.

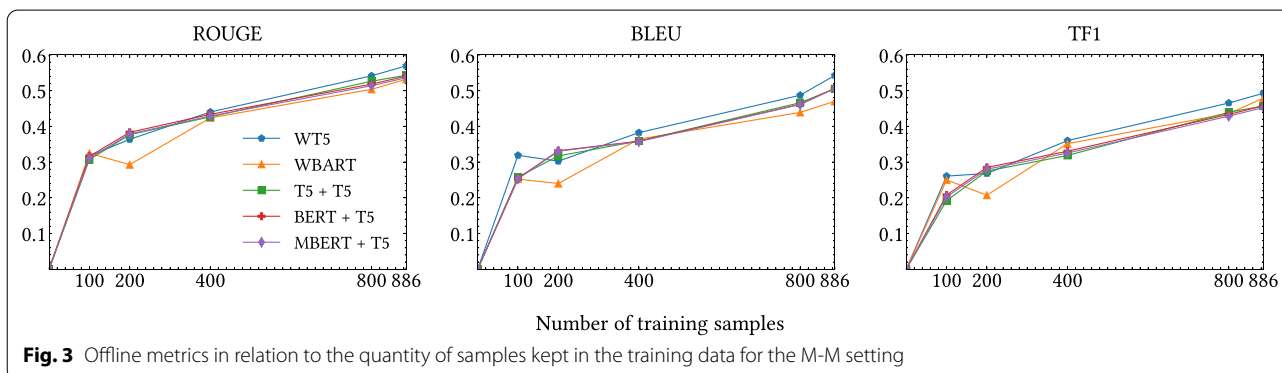
In terms of human evaluation, we presented the explanations generated by our WT5 model in the *M-M* setting to three domain experts, with 209 explanations provided to each of them. Table 5 provides four examples of positive predictions along with their generated explanations. The first three rows correspond to relevant explanations, since the three human assessors considered them as relevant. The last row shows a non-relevant explanation. In this case, while it appears to be topically related to a symptom of depression (*Suicidal ideas*), the experts classified it as non-relevant, as it does not provide any information about the symptom. Overall, the three assessors found 73%, 53%, and 77% of the explanations relevant and clinically useful, resulting in an average of 68%. A more granular inspection of the assessors' annotations indicates that, out of 209 explanations, 154 were deemed

<sup>5</sup> Due to clarity and space constraints, not all systems are shown.

**Table 5** Each row presents an example of positive predictions in the *M-M* setting using the WT5 model

Input post	Relevant
( . . . ) I want to start a business one moment, then pay out my IRA and travel throughout Europe. I do not comprehend who I am. <b>My short-term memory is terrible, and I can not concentrate.</b> I'm unsure of what to do. You guys are going to advise some really fantastic actions for me to pursue, but ultimately I lack the willpower or energy to carry out your advice.	1
Recently, I have been having a lot of difficulty with this. <b>I have been depressed, worried, or ill since I was a child.</b> Like my youth has been taken from me. As a last-ditch effort to feel like a person once more, I am actually considering seeing a naturopathic physician. Anyway, I hope you soon feel better.	1
I abhor myself to the core. Even reading back through Reddit postings I posted a few days ago makes me want to commit suicide. <b>I am such a disgusting waste of life - useless, unproductive, and with a future that is already in uncertainty.</b> And the fact that Im feeling this way on spring break is something I detest so fiercely. ( . . . )	1
( . . . ) how EMTs and first responders are looking after them, and how those individuals should persevere to witness another day. <b>None-theless, I find myself unable to avoid fall into thoughts about someone's death.</b> This triggers memories of my own experiences and how I might find myself in that person's position ( . . . )	0

Explanations are **highlighted** as judged by the assessors. The *Relevant* column displays the relevance labels after the human assessment. Samples have been paraphrased for privacy and shortened for clarity



clinically useful by at least two of the three assessors. In terms of total consensus, which is an important factor as it unambiguously confirms or rejects the practical utility of an explanation, we find that assessors agreed on the non-relevance of 25 explanations and the relevance of 86. Regarding the inter-annotator agreement among domain experts in our study, we computed a pairwise agreement to measure the percentage of cases where the three assessors reached a consensus in their annotations. The observed agreement values among the pairs were 76%, 66%, and 65%, yielding an average agreement of 69% regarding the clinical relevance of the explanations. This level of concordance aligns with similar studies in the field [25, 36, 86], considering the challenges associated with annotating in the complex domain of mental health [87].

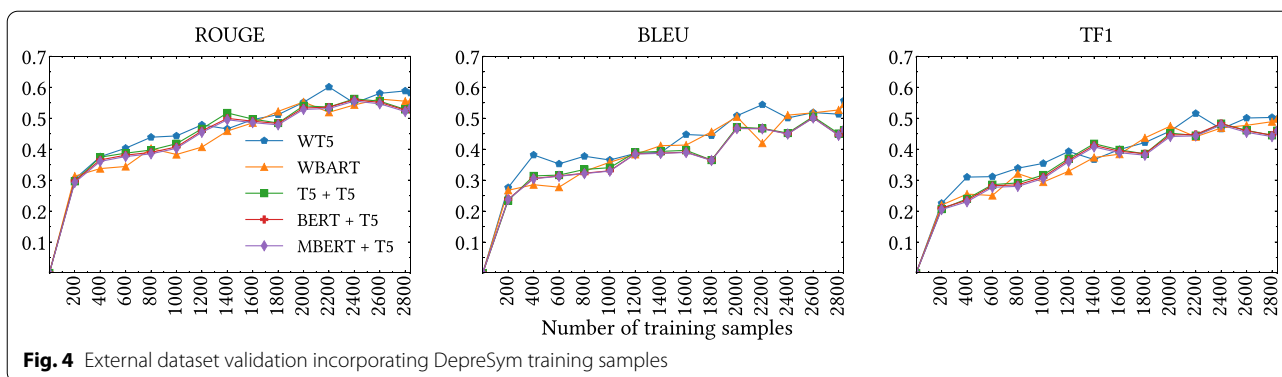
**RQ2: How many labels are needed to generate good explanations?**

Since collecting data labelled by domain experts in the mental health domain on social media is a costly process [86, 88], we investigate here how many training samples are needed to generate quality explanations. To analyze

this aspect, we designed two different experimental setups: First, in "Impact when training with limited explanations" section, we investigate the effect of the number of labelled examples on the quality of explanations for the experimental setup used in this paper. Then, in "Assessing explanation quality stabilization when varying number of training samples" section, we validate the results of this experiment adapting and including an additional external dataset.

**Impact when training with limited explanations**

To assess how the quantity of training samples would affect our experiments, we trained the PLMs with a gradually increasing number of examples using our larger setting, denoted as *M-M*. This setting includes 886 explanations in the training set, and we compared our models' performance using subsets of 100, 200, 400, and 800 training samples. Figure 3 presents the results across the three evaluation metrics (ROUGE, BLEU and TF1) that measure the quality of explanations generated. As shown in the Figure, there is a nearly linear improvement with respect to the number of training instances for all metrics and models. For instance, when our models are trained



with only 100 explanations, we achieve approximately half the performance compared to the entire training split. For completeness, we also investigated here the performance of the classification task, where we observed a more robust behaviour. For instance, using only 200 samples, all models consistently achieved a good F1 score, being higher than 0.9. In line with previous literature [74], we find that the number of training samples affects the quality of the explanations more than the classification results.

**Assessing explanation quality stabilization when varying number of training samples**

Following the above experimentation described in "Impact when training with limited explanations" section, we further investigated this aspect by conducting an additional experiment adapting the DepreSym [89] resource, an external dataset sourced from Reddit. This dataset also contains sentences that indicate the presence of markers for depression symptoms, allowing us to collect the user posts and then use these sentences as explanations. Using this resource, we shuffled these new explanations within the previous training set of the *M-M* setting and evaluated the explanation metrics every 200 training examples for the original *M-M* test set. Thus, we added 1956 new training examples, but kept the previous test set. Figure 4 shows our results. Looking at the performance of our models, we can see that the values of the metrics continue to improve after 800 training instances. However, around 2000 instances, we found a trend towards stabilisation of performance. Therefore, we can observe that after a certain point incorporate additional training samples does not guarantee an improvement on the performance.

**RQ3: single-step vs Two-steps**

Through our experimentation, we conducted a comparative analysis of finetuned PLMs to perform the *classify + explain* tasks in a single step (WT5 and WBART) against

those operating in two separate steps (T5 + T5, BERT + T5, MBERT + T5). Regarding the classification task, the two-step models exhibited a slight advantage, as reflected in column F1 from Table 4a. Specifically, the BERT-based models consistently outperformed the other models, with BERT and MBERT achieving the best classification performance in four of the five settings. Regarding explanation, to ensure a fair comparison, we focused on settings with a similar number of true positives (TPs), which led us to consider the *P-P* and *M-M* settings. In both cases, the models with the best explanations were WT5 and WBART.

In the context of our investigation, we can see that both the single-step and two-step approaches show promising potential for the classification and explanation tasks. However, we can see slight differences of the models in the two tasks. Looking at the results, we observe that, single-step systems often yield superior results in terms of explanation, even if their classification performance is marginally lower. This aligns with prior research suggesting that consolidating classification and explanation into a single step can enhance the quality of explanations [74, 90, 91]. Under the single-step paradigm, classification and explanation are treated as an interconnected whole rather than discrete operations. This integration may lead to more cohesive and informative explanations, underscoring the importance of considering the aggregated nature of the task. Regarding the slight better performance of the two-steps models in terms of classification, this also aligns with prior work, since BERT-based models are still predominant in classification for similar scenarios [65].

**RQ4: in-context Conversational instructed LLMs performance**

When evaluating the results achieved by conversational LLMs, it is important to distinguish between general purpose models used within a few-shot approach (Vicuna-13B, GPT-3.5 and Mistral-Instruct) and domain-specific

finetuned versions such as MentalLLaMA and MentalMistral. The overall results presented in Table 3 show that all the instructed LLMs achieve lower performance than our finetuned PLMs, especially for the explanation task. Table 4a reveals that it is GPT-3.5 method that comes closer to the best-performing systems among the general purpose LLMs in terms of classification but, remains far behind in terms of explanations quality. Moreover, our experiments also show that conversational LLMs benefit greatly from finetuning, whether they were finetuned at the symptom level (MentalMistral) or not (MentalLLaMA). This leap in performance can be demonstrated by comparing the results of MentalMistral with those of its base model, Mistral-Instruct. MentalMistral experiences a notorious boost and establishes itself as the model with the best classification results among the conversational LLMs. However, as noted in "Symptoms explanation" section, even these fine-tuned LLMs face challenges in generating good explanations, and can provide inappropriate and inaccurate explanations due to the hallucination effect [82]. In order to study this phenomenon in the context of our experiments, we calculated the percentage of cases where the generated explanations do not exactly match the original text. For the GPT-3.5 model and the *M-M* reference setting, this occurs in only 1.5% of cases. This low rate suggest a high degree of precision in the model's output, exemplifying its effectiveness in generating extractive explanations.

Despite these results, it is important to consider that we only used 30 positive and negative examples to guide these systems through in-context learning.<sup>6</sup> The limited exposure to examples for this new task might have contributed to their relatively lower performance. Moreover, recent studies have shown the variability in the performance of these models depending on the prompts quality [92, 93]. The prompt we defined for all of our experiments is listed on "Appendix 2: Prompt for conversational LLMs" section, and we plan to explore the impact of different prompt strategies in these particular scenarios.

### Conclusions and future work

In this paper, we presented the use of different text-to-text pipelines to classify and explain the presence of depressive symptoms on social media. Our models, in addition to detecting relevant social media posts related to depressive symptoms, also explain in natural language their decisions. Leveraging two datasets containing sentences indicative of depressive symptoms, we evaluated our model variants using a wide variety of offline metrics and expert-in-the-loop evaluations. Our

results are promising in both classification and explanation tasks. Within our methods, we compared single-step approaches (where classification and explanation are performed by the same model) against two-step approaches (where the two tasks are performed by separate models). Moreover, our analysis extended to examining the capabilities of recent conversational LLMs in detecting and explaining depressive symptoms. Our study enriches the expanding field of explainable AI in mental health applications. By offering healthcare systems models capable of explaining their decision-making process, we allow clinicians to understand the automated reasoning and build trust in the outputs of these models.

While we presented an extractive approach for generating explanations, our discussions with domain experts have revealed the potential to create abstract explanations, matching the way one expert would explain the existence of depressive symptoms to another. Moreover, it might be advantageous to identify the precise symptom present in positive posts, leaving behind the actual binary paradigm and defining a new multi-class approach. An important direction for future work involves incorporating temporal information into our models. We are planning to integrate the publication dates of Reddit posts, enabling our systems to process data in temporal batches. This approach would facilitate weekly or monthly analysis, thereby improving our decision-making processes. Such temporal integration is valuable for providing clinicians with insights and evidence of the evolution of depressive symptoms in individuals over time. Additionally, this incorporation would allow us to conduct population-level analyses, examining symptom prevalence during specific periods, such as the COVID-19 pandemic [94]. This temporal perspective may help to understand and respond to fluctuating mental health trends in the digital era. As a limitation, we are also aware that, although we use two different datasets, the data obtained comes from a single social media platform. In future research, extending our approach across diverse platforms, languages, and cultural contexts would be valuable to evaluate its cross-cultural applicability. Overall, we believe that, the creation of models capable of explaining the detection of clinical markers, is an important first step towards developing sophisticated natural language explanations for depression detection and related analysis on social media.

### Ethical and environmental aspects

The data used in our study were obtained from publicly accessible sources, adhering to the exempt status under title 45 CFR §46.104. The use of BDI-Sen, PsySym and DepreSym datasets was accomplished in full compliance with their respective data usage policies. To maintain

<sup>6</sup> Models have a maximum sequence length which implies a limit on the number of samples.



privacy, we implemented measures to ensure that any personal information was unidentifiable and all users remained anonymous. The data used were sourced from Reddit, and we strictly conformed to all terms specified by this platform. For the human evaluation, even though the assessors were domain experts, they were not subjected to any imposed time restrictions, and reported no adverse effects post-evaluation. Importantly, it must be emphasized that the systems described in this study are intended to complement the work of healthcare professionals, not replace them. The development of such technologies necessitates a cautious approach, with a continuous emphasis on their ethical use and a firm respect for patient privacy and autonomy.

The experiments for this study were conducted using our private infrastructure, with a carbon efficiency of 0.432 kgCO<sub>2</sub>eq/kWh, which reflects the OECD's 2014 yearly average. The resources utilised included 10 h of computation on an RTX A6000 device (with a TDP of 300W) and 50 h on an A100 PCIe 40/80GB device (with a TDP of 250W). Total emissions are estimated to be 8.43 kgCO<sub>2</sub>eq. To provide some perspective, this is equivalent to driving an average car for 34 kilometres. These figures were determined with the assistance of the Machine-Learning Impact Calculator [95].

### Appendix 1: Assessment guidelines

Considering different depressive symptoms based on the BDI questionnaire (e.g: Sleep problems, Suicidal Ideation, Sadness, Pessimism...), the original task is that, from a Reddit publication, the predicted explanations must give relevant information to some kind of depressive symptomatology.

#### Important Notes

A relevant explanation should provide some information about the state of the individual related to a depressive symptom. It is not necessary that the exact same words are used for the explanation, but they must correspond to what was described in the original post. Each post can have more than one explanation. In this case, if one of them is relevant, this is enough to assess it as relevant.

#### Examples (some original posts and the correct explanation)

- Relevant explanation
  - Post: “I’m very concerned about my future. I’m in uni atm, and I’m terrified of not getting a good job, not being able to have kids, not passing my exams, not having enough money to buy my own house,

take care of my kids (I’m staying hopeful) or do the things I want to do. I’m basically terrified of not being happy with my life”.

- Explanation: “I’m basically terrified of not being happy with my life”.

- Relevant explanation

- Post: “My problem was that i tried to do everything at the same time. Its important to take a breath few times a day. Like for me I’m a student and after school my plan was to, practice programming(pyhton), doing my website project, reading a book about programming and practicing art. I burnt out and I am still unmotivated and feeling like shit. Don’t do too many things at the same time”.

- Explanation: “I burnt out and i still am unmotivated and feeling like shit”.

- Relevant explanation

- Post: “Yeah. I feel like I don’t even get excited about things anymore”.

- Explanation: “I feel like I don’t even get excited about things anymore”.

- Non-Relevant explanation

- Post: “Even if it is, big fuckin’ woop. I liked the movie, and this event was fun. Who cares who organized it?”.

- Explanation: “I liked the movie, and this event was fun.”.

- Non-Relevant explanation

- Post: “Nope. We are from Germany”.

- Explanation: “We are from Germany”.

Your job is to assess if the explanations given to a specific post are informative about some kind of depressive symptomatology. The relevance grades are:

- Relevant (1): A relevant explanation should be relevant to a depressive symptom, and it must correspond to the content written in the original publication.
- Non-relevant (0): A non-relevant sentence does not address any topic related to a depressive symptom, or the explanation is not informative about the original publication.

The Excel sheet contains three columns (Post, Explanation and Relevant Explanation). The column "Post" contains the original post, the "Explanation" corresponds to the explanation based on that specific post. The last column, "Relevant Explanation", is where you have to fill in 0 or 1.

To measure the assessment effort, we ask you to record the time spent on fully evaluating the explanations presented.

## Appendix 2: Prompt for conversational LLMs

See Fig. 5.

```
[System]
You are an expert annotator that finds whether there is evidence or not about depression symptoms in
user posts, and you can justify your decision selecting post's excerpts generating the answer in XML
format.

[Instructions]
As an expert annotator, you determine whether a user's writing contains evidence of depression symptoms
(relevant) or not (non-relevant). You need to assess whether the writing shows any information, on the
positive or negative side, about the writer (not about anybody else) related to any of the following
symptoms of depression: sadness, pessimism, sense of failure, loss of pleasure, guilt, punishment,
self-dislike, self-incrimination, suicidal ideas, crying, agitation, social issues, indecision,
feelings of worthlessness, low energy, sleep issues, irritability, appetite issues, problems with
concentration, fatigue, and low libido.

In addition to identifying if there is information about one or more symptoms in the user's posts, when
there is evidence (relevant post), you must support your decision by selecting one excerpt from the
original post that provides the information supporting your judgment.

This is the expected input and output XML format:

Post: "I haven't even had the energy to care because I have simply been sleeping for the past three
days"
<output>
  <decision>relevant</decision>
  <excerpt>I haven't even had the energy to care</excerpt>
</output>

[In-context learning]
Given the provided instructions judge the following post for relevance and select the excerpt from the
post on relevant ones.

Post: "... "
<output>
  <decision>{label}</decision>
  <excerpt>{explanation}</excerpt>
</output>

...

[Question]
Post: "{post}"
```

**Fig. 5** Prompt used for our experiments with conversational LLMs

**Acknowledgements**

This work has received support from projects: PLEC2021-007662 (MCIN/AEI/10.13039/501100011033, Ministerio de Ciencia e Innovación, European Union NextGenerationEU/PRTR) and PID2022-137061OB-C21 (MCIN/AEI/10.13039/501100011033/, Ministerio de Ciencia e Innovación, ERDF A way of making Europe, by the European Union); Consellería de Educación, Universidade e Formación Profesional, Spain (grant number ED481A-2024-079 and accreditations 2019–2022 ED431G/01 and GPC ED431B 2022/33) and the European Regional Development Fund, which acknowledges the CITIC Research Center. Funding for open access charge: Universidade da Coruña/CISUG. We would also like to thank Desireé Pombar, Silvia López-Larrosa and Laura Hermo for their efforts in evaluating the results. Their work allowed us to assess the practical utility and clinical relevance of the explanations generated.

**Data availability**

The datasets used in this work can be obtained upon request to their original authors.

Received: 29 January 2024 Accepted: 17 August 2024

Published: 6 September 2024

**References**

- Prince M, Patel V, Saxena S, Maj M, Masello J, Phillips MR, Rahman A. No health without mental health. *Lancet*. 2007;370(9590):859–77.
- Mental Health TNS. Wellbeing: mental health: prevalence and impact. Technical report, Australian Institute of Health and Welfare 2022
- Organization WH, et al. Depression and other common mental disorders: global health estimates. Technical report, World Health Organization; 2017
- Picardi A, Lega I, Tarsitani L, Caredda M, Matteucci G, Zerella MP, Miglio R, Gigantesco A, Cerbo M, Gaddini A, Spandonaro F, Biondi M. SET-DEP Group: a randomised controlled trial of the effectiveness of a program for early detection and treatment of depression in primary care. *J Affect Disord*. 2016;198:96–101.
- Gulliver A, Griffiths KM, Christensen H. Perceived barriers and facilitators to mental health help-seeking in young people: a systematic review. *BMC Psychiatry*. 2010;10:113.
- De Choudhury M, Gamon M, Counts S, Horvitz E. Predicting depression via social media. In: Proceedings of the international AAAI conference on web and social media. 2013. p. 128–37
- Sadeque F, Xu D, Bethard S. Measuring the latency of depression detection in social media. In: Proceedings of the eleventh ACM international conference on web search and data mining. WSDM '18. New York: Association for Computing Machinery; 2018. p. 495–503
- Yates A, Cohan A, Goharian N. Depression and self-harm risk assessment in online forums. In: Proceedings of the 2017 conference on empirical methods in natural language processing. Copenhagen: Association for Computational Linguistics; 2017. p. 2968–78
- Cacheda F, Fernandez D, Novoa FJ, Carneiro V. Early detection of depression: social network analysis and random forest techniques. *J Med Internet Res*. 2019;21(6):12554.
- Beheshti A, Moraveji-Hashemi V, Yakhchi S, Motahari-Nezhad H.R, Ghafari S.M, Yang J. Personality2vec: enabling the analysis of behavioral disorders in social networks. In: Proceedings of the 13th international conference on web search and data mining. WSDM '20. New York: Association for Computing Machinery; 2020. p. 825–28
- Bucur A.-M, Cosma A, Rosso P, Dinu LP. It's just a matter of time: detecting depression with time-enriched multimodal transformers. In: European conference on information retrieval; 2023. Springer. p. 200–15
- Aragon M, Lopez Monroy A.P, Gonzalez L, Losada D.E, Montes M. Dis-orBERT: a double domain adaptation model for detecting signs of mental disorders in social media. In: Proceedings of the 61st annual meeting of the association for computational linguistics (vol. 1: Long Papers). Toronto: Association for Computational Linguistics; 2023. p. 15305–18
- Callahan A, Inckle K. Cybertherapy or psychobabble? A mixed methods study of online emotional support. *Br J Guidance Counsell*. 2012;40(3):261–78.
- Ferraro G, Loo Gee B, Ji S, Salvador-Carulla L. Lightme: analysing language in internet support groups for mental health. *Health Inf Sci Syst*. 2020;8(1):34.
- Rissola EA, Losada DE, Crestani F. A survey of computational methods for online mental state assessment on social media. *ACM Trans. Comput. Healthc*. 2021;2(2):1–31.
- Parapar J, Martín-Rodilla P, Losada D.E, Crestani F. eRisk 2023: depression, pathological gambling, and eating disorder challenges. In: Advances in information retrieval: 45th European conference on information retrieval, ECIR 2023, Dublin, Ireland, April 2–6, 2023, Proceedings, Part III. Springer; 2023. p. 585–92
- Zirikly A, Atzil-Slonim D, Liakata M, Bedrick S, Desmet B, Ireland M, Lee A, MacAvaney S, Purver M, Resnik R, Yates A, editors. Proceedings of the eighth workshop on computational linguistics and clinical psychology. Seattle: Association for Computational Linguistics; 2022
- Walsh CG, Chaudhry B, Dua P, Goodman KW, Kaplan B, Kavuluru R, Solomonides A, Subbian V. Stigma, biomarkers, and algorithmic bias: recommendations for precision behavioral health with artificial intelligence. *JAMIA Open*. 2020;3(1):9–15.
- Hauser TU, Skvortsova V, De Choudhury M, Koutsouleris N. The promise of a model-based psychiatry: building computational models of mental ill health. *Lancet Digit. Health*. 2022;4(11):816–28.
- Ernala SK, Birnbaum ML, Candan KA, Rizvi AF, Sterling WA, Kane JM, De Choudhury M. Methodological gaps in predicting mental health states from social media: triangulating diagnostic signals. In: Proceedings of the 2019 CHI conference on human factors in computing systems. CHI '19. New York: Association for Computing Machinery; 2019. p. 1–16
- Beck AT, Steer RA, Brown G. Beck Depression Inventory–II. American Psychological Association (APA). 1996
- Kroenke K, Spitzer RL, Williams JB. The PHQ-9: validity of a brief depression severity measure. *J General Internal Med*. 2001;16(9):606–613
- Pérez A, Parapar J. Barreiro: Automatic depression score estimation with word embedding models. *Artif Intell Med*. 2022;132:102380.
- Nguyen T, Yates A, Zirikly A, Desmet B, Cohan A. Improving the generalizability of depression detection by leveraging clinical questionnaires. In: Proceedings of the 60th annual meeting of the association for computational linguistics, vol. 1 (Long Papers). Dublin: Association for Computational Linguistics; 2022. p. 8446–59
- Zhang Z, Chen S, Wu M, Zhu K. Symptom identification for interpretable detection of multiple mental disorders on social media. In: Proceedings of the 2022 conference on empirical methods in natural language processing. Abu Dhabi: Association for Computational Linguistics; 2022. p. 9970–85
- Zhang Z, Chen S, Wu M, Zhu KQ. Psychiatric scale guided risky post screening for early detection of depression. In: Raedt LD, editor. Proceedings of the thirty-first international joint conference on artificial intelligence, IJCAI-22. AI for Good; 2022. p. 5220–26
- Pérez A, Warikoo N, Wang K, Parapar J, Gurevych I. Semantic similarity models for depression severity estimation. In: Bouamor H, Pino J, Bali K, editors. Proceedings of the 2023 conference on empirical methods in natural language processing. Singapore: Association for Computational Linguistics; 2023. p. 16104–18.
- Smith KM, Renshaw PF, Bilello J. The diagnosis of depression: current and emerging methods. *Compr Psychiatry*. 2013;54(1):1–6.
- Doshi-Velez F, Kim B. Towards a rigorous science of interpretable machine learning. 2017.
- Markus AF, Kors JA, Rijnbeek PR. The role of explainability in creating trustworthy artificial intelligence for health care: a comprehensive survey of the terminology, design choices, and evaluation strategies. *J Biomed Inform*. 2021;113:103655.
- Miller T. Explanation in artificial intelligence: insights from the social sciences. *Artif Intell*. 2019;267:1–38.
- OpenAI: introducing ChatGPT. Accessed Aug 1 2023. 2022
- Chiang W-L, Li Z, Lin Z, Sheng Y, Wu Z, Zhang H, Zheng L, Zhuang S, Zhuang Y, Gonzalez J.E, Stoica I, Xing EP. Vicuna: an open-source chatbot impressing GPT-4 with 90%\* ChatGPT quality. 2023.
- Jiang A.Q, Sablayrolles A, Mensch A, Bamford C, Chaplot DS, Casas D, Bressand F, Lengyel G, Lample G, Saulnier L, Lavaud LR, Lachaux M-A, Stock P, Scao TL, Lavril T, Wang T, Lacroix T, Sayed WE. Mistral 7B. 2023
- Yang K, Zhang T, Kuang Z, Xie Q, Ananiadou S. Mentalllama: interpretable mental health analysis on social media with large language models. To be

- published at: WWW '24: the ACM web conference 2023, Singapore May 13–17, 2024. 2023
36. Perez A, Parapar J, Barreiro A, Lopez-Larrosa S. Bdi-sen: a sentence dataset for clinical symptoms of depression. In: Proceedings of the 46th international ACM SIGIR conference on research and development in information retrieval. 2023.
  37. Islam MR, Kabir MA, Ahmed A, Kamal ARM, Wang H, Ulhaq A. Depression detection from social network data using machine learning techniques. *Health Inf Sci Syst.* 2018;6(1):8
  38. Li Z, An Z, Cheng W, Zhou J, Zheng F, Hu B. MHA: a multimodal hierarchical attention model for depression detection in social media. *Health Inf. Sci. Syst.* 2023;11(1):6.
  39. Guntuku SC, Yaden DB, Kern ML, Ungar LH, Eichstaedt JC. Detecting depression and mental illness on social media: an integrative review. *Big data in the behavioural sciences. Curr Opin Behav Sci.* 2017;18:43–49
  40. Skaik R, Inkpen D. Using social media for mental health surveillance: a review. *ACM Comput Surv* 2020;53(6)
  41. Crestani F, Losada DE, Parapar J. Early detection of mental health disorders by social media monitoring: the first five years of the ERisk project vol. 1018. 2022
  42. Wang T, Brede M, Ianni A, Mentzakis E. Detecting and characterizing eating-disorder communities on social media. In: Proceedings of the tenth ACM international conference on web search and data mining. WSDM '17. New York: Association for Computing Machinery; 2017. p. 91–100
  43. Shuai H-H, Shen C-Y, Yang D-N, Lan Y-F, Lee W-C, Yu PS, Chen M-S. Mining online social data for detecting social network mental disorders. In: Proceedings of the 25th international conference on world wide web. WWW '16. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE. 2016. p. 275–85
  44. Halfin A. Depression: the benefits of early and appropriate treatment. *Am J Manag Care.* 2007;13(4 Suppl):92–7.
  45. Guidotti R, Ruggieri S. On the stability of interpretable models. In: 2019 international joint conference on neural networks (IJCNN). 2019. p. 1–8
  46. Tulio Ribeiro M, Singh S, Guestrin C. Model-agnostic interpretability of machine learning. 2016. 1606–05386[stat.ML]
  47. Ribeiro M.T, Singh S, Guestrin C. Anchors: high-precision model-agnostic explanations. *Proc AAAI Conf Artif Intell* 2018;32(1)
  48. Baehrens D, Schroeter T, Harmeling S, Kawanabe M, Hansen K, Müller K.-R. How to explain individual classification decisions. *J Mach Learn Res.* 2010;11:1803–31
  49. Ribeiro MT, Singh S, Guestrin C. "Why should i trust you?" Explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. 2016. p. 1135–44
  50. Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. In: Advances in neural information processing systems, vol. 30. 2017
  51. Lopez-Larrosa S, Sánchez-Souto V, Losada D.E, Parapar J, Barreiro Ha A.P, Cummings E.M. Using machine learning techniques to predict adolescents' involvement in family conflict. *Soc Sci Comput Rev.* 2023;41(5):1581–607
  52. Nordin N, Zainol Z, Mohd Noor M.H, Chan L.F. An explainable predictive model for suicide attempt risk using an ensemble learning and shapley additive explanations (SHAP) approach. *Asian J Psychiatry* 2023;79:103316
  53. Ustun B, Rudin C. Supersparse linear integer models for optimized medical scoring systems. *Mach Learn.* 2016;102(3):349–91.
  54. Lakkaraju H, Bach SH, Leskovec J. Interpretable decision sets: A joint framework for description and prediction. In: Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. KDD '16. New York: Association for Computing Machinery; 2016. p. 1675–84
  55. Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. In: Proceedings of the 31st international conference on neural information processing systems. NIPS'17. Red Hook: Curran Associates Inc.; 2017. p. 4768–77
  56. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Lu, Polosukhin I. Attention is all you need. In: Guyon I, Luxburg U.V, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R, editors. Advances in neural information processing systems. 2017;30
  57. Mullenbach J, Wiegrefe S, Duke J, Sun J, Eisenstein J. Explainable prediction of medical codes from clinical text. In: Proceedings of the 2018 conference of the North American chapter of the association for computational linguistics: human language technologies, vol. 1 (Long Papers). New Orleans: Association for Computational Linguistics; 2018. p. 1101–111
  58. Song H, You J, Chung J-W, Park J.C. Feature attention network: Interpretable depression detection from social media. In: Proceedings of the 32nd Pacific Asia conference on language, information and computation. Hong Kong: Association for Computational Linguistics; 2018
  59. Amini H, Kosseim L. Towards explainability in using deep learning for the detection of anorexia in social media. In: Métais E, Meziane F, Horacek H, Cimiano P, editors. Natural language processing and information systems. Cham: Springer; 2020. p. 225–35.
  60. Wiegrefe S, Pinter Y. Attention is not not explanation. In: Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP). Hong Kong: Association for Computational Linguistics; 2019. p. 11–20
  61. Tutek M, Šnajder J. Toward practical usage of the attention mechanism as a tool for interpretability. *IEEE Access.* 2022;10:47011–30.
  62. Jain S, Wallace BC. Attention is not explanation. In: Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, vol. 1 (Long and Short Papers). Minneapolis: Association for Computational Linguistics; 2019. p. 3543–56.
  63. Serrano S, Smith NA. Is attention interpretable? In: Proceedings of the 57th annual meeting of the association for computational linguistics. Florence: Association for Computational Linguistics; 2019. p. 2931–51.
  64. Camburu O-M, Rocktäschel T, Lukasiewicz T, Blunsom P. e-snli: natural language inference with natural language explanations. In: Bengio S, Wallach H, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R, editors. Advances in neural information processing systems, vol. 31
  65. Yang K, Ji S, Zhang T, Xie Q, Kuang Z, Ananiadou S. Towards interpretable mental health analysis with large language models. In: Bouamor H, Pino J, Bali K, editors. Proceedings of the 2023 conference on empirical methods in natural language processing, EMNLP 2023, Singapore, December 6–10; 2023. p. 6056–77.
  66. Kojima T, Gu S.S, Reid M, Matsuo Y, Iwasawa Y. Large language models are zero-shot reasoners. In: Koyejo S, Mohamed S, Agarwal A, Belgrave D, Cho K, Oh A, editors. Advances in neural information processing systems. 2022;35:22199–213.
  67. Cambria E, Malandri L, Mercurio F, Mezzanzanica M, Nobani N. A survey on xai and natural language explanations. *Inf. Process Manag.* 2023;60(1):103111.
  68. Raffel C, Shazeer N, Roberts A, Lee K, Narang S, Matena M, Zhou Y, Li W, Liu P.J. Exploring the limits of transfer learning with a unified text-to-text transformer. *J Mach Learn Res.* 2020;21(1)
  69. Sutskever I, Vinyals O, Le Q.V. Sequence to sequence learning with neural networks. In: Proceedings of the 27th international conference on neural information processing systems-volume 2. NIPS'14. Cambridge: MIT Press; 2014. p. 3104–12
  70. Devlin J, Chang M-W, Lee K, Toutanova K. BERT: pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, vol 1 (Long and Short Papers). Minneapolis: Association for Computational Linguistics; 2019. p. 4171–86
  71. Wang A, Pruksachatkun Y, Nangia N, Singh A, Michael J, Hill F, Levy O, Bowman SR. SuperGLUE: a stickier benchmark for general-purpose language understanding systems. Red Hook: Curran Associates Inc; 2019.
  72. Lewis M, Liu Y, Goyal N, Ghazvininejad M, Mohamed A, Levy O, Stoyanov V, Zettlemoyer L. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In: Proceedings of the 58th annual meeting of the association for computational linguistics. Association for Computational Linguistics; 2020. p. 7871–80.
  73. Erikksson KF. Arpteg A, Matskin M, Payberah AH. Cross-domain transfer of generative explanations using text-to-text models. In: Natural language processing and information systems: 26th international conference on applications of natural language to information systems. NLDB 2021, Saarbrücken, Germany, June 23–25, 2021, Proceedings. Berlin, Heidelberg: Springer; 2021. p. 76–89.

74. Narang S, Raffel C, Lee K, Roberts A, Fiedel N, Malkan K. WT5?! Training Text-to-Text Models to Explain their Predictions 2020.
75. Ji S, Zhang T, Ansari L, Fu J, Tiwari P, Cambria E. MentalBERT: publicly available pretrained language models for mental healthcare. In: Proceedings of the thirteenth language resources and evaluation conference. Marseille: European Language Resources Association; 2022. p. 7184–90.
76. Brown T, Mann B, Ryder N, Subbiah M, Kaplan J.D, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, Agarwal S, Herbert-Voss A, Krueger G, Henighan T, Child R, Ramesh A, Ziegler D, Wu J, Winter C, Hesse C, Chen M, Sigler E, Litwin M, Gray S, Chess B, Clark J, Berner C, McCandlish S, Radford A, Sutskever I, Amodei D. Language models are few-shot learners. In: Larochelle H, Ranzato M, Hadsell R, Balcan MF, Lin H, editors. *Advances in neural information processing systems*, vol. 33. 2020. p. 1877–901.
77. Dettmers T, Pagnoni A, Holtzman A, Zettlemoyer L. QLoRA: efficient finetuning of quantized LLMs. 2023
78. Touvron H, Martin L, Stone K, Albert P, Almahairi A, Babaei Y, Bashlykov N, Batra S, Bhargava P, Bhosale S, Bikel D, Blecher L, Ferrer C.C, Chen M, Cucurull G, Esiohub D, Fernandes J, Fu J, Fu W, Fuller B, Gao C, Goswami V, Goyal N, Hartshorn A, Hosseini S, Hou R, Inan H, Kardas M, Kerkez V, Khabsa M, Kloumann I, Korenev A, Koura PS, Lachaux M-A, Lavril T, Lee J, Liskovitch D, Lu Y, Mao Y, Martinet X, Mihaylov T, Mishra P, Molybog I, Nie Y, Poulton A, Reizenstein J, Rungta R, Saladi K, Schelten A, Silva R, Smith EM, Subramanian R, Tan XE, Tang B, Taylor R, Williams A, Kuan J.X, Xu P, Yan Z, Zarov I, Zhang Y, Fan A, Kambadur M, Narang S, Rodriguez A, Stojnic R, Edunov S, Scialom T. Llama 2: open foundation and fine-tuned chat models. 2023
79. Jackson-Koku G. Beck depression inventory. *Occup Med*. 2016;66(2):174–5.
80. Naseem U, Dunn A.G, Kim J, Khushi M. Early identification of depression severity levels on reddit using ordinal classification. In: *Proceedings of the ACM web conference 2022*. 2022. p. 2563–72
81. Nuckols CC, Nuckols CC. *The diagnostic and statistical manual of mental disorders, (DSM-5)*. Philadelphia: American Psychiatric Association; 2013.
82. Ji Z, Lee N, Frieske R, Yu T, Su D, Xu Y, Ishii E, Bang YJ, Madotto A, Fung P. Survey of hallucination in natural language generation. *ACM Comput Surv*. 2023;55(12):1–38.
83. Lin C-Y. ROUGE: a package for automatic evaluation of summaries. In: *Text summarization branches out*. Association for Computational Linguistics, Barcelona, Spain; 2004. p. 74–81
84. Papineni K, Roukos S, Ward T, Zhu W-J. Bleu: a method for automatic evaluation of machine translation. In: *Proceedings of the 40th annual meeting of the association for computational linguistics*. Association for Computational Linguistics, Philadelphia, Pennsylvania, USA; 2002. p. 311–18
85. DeYoung J, Jain S, Rajani N.F, Lehman E, Xiong C, Socher R, Wallace BC. ERASER: a benchmark to evaluate rationalized NLP models. In: *Proceedings of the 58th annual meeting of the association for computational linguistics*. Association for Computational Linguistics; 2020. p. 4443–58.
86. MacAvaney S, Desmet B, Cohan A, Soldaini L, Yates A, Ziriky A, Goharian N. RSDDD-time: temporal annotation of self-reported mental health diagnoses. In: *Proceedings of the fifth workshop on computational linguistics and clinical psychology: from keyboard to clinic*. New Orleans: Association for Computational Linguistics; 2018. p. 168–73
87. Mowery D, Smith H, Cheney T, Stoddard G, Coppersmith G, Bryan C, Conway M, et al. Understanding depressive symptoms and psychosocial stressors on twitter: a corpus-based study. *J Med Internet Res*. 2017;19(2):6895.
88. Findings of the Association for Computational Linguistics. Harrigan K, Aguirre C, Dredze M. Do models of mental health based on social media data generalize? In: *EMNLP*; 2020. p. 3774–88.
89. Pérez A, Fernández-Pichel M, Parapar J, Losada D.E. DepreSym: a depression symptom annotated corpus and the role of LLMs as assessors of psychological markers. 2023
90. Feder A, Keith KA, Manzoor E, Pryzant R, Sridhar D, Wood-Doughty Z, Eisenstein J, Grimmer J, Reichart R, Roberts ME, Stewart BM, Veitch V, Yang D. Causal inference in natural language processing: estimation, prediction, interpretation and beyond. *Trans Assoc Comput Linguist*. 2022;10:1138–58.
91. Barredo Arrieta A, Díaz-Rodríguez N, Del Ser J, Bennetot A, Tabik S, Barbado A, Garcia S, Gil-Lopez S, Molina D, Benjamins R, Chatila R, Herrera F. Explainable artificial intelligence (xai): concepts, taxonomies, opportunities and challenges toward responsible ai. *Information Fusion*. 2020;58:82–115.
92. Zhao Z, Wallace E, Feng S, Klein D, Singh S. Calibrate before use: improving few-shot performance of language models. In: *International conference on machine learning*. PMLR; 2021. p. 12697–706
93. Xun G, Jia X, Gopalakrishnan V, Zhang A. A survey on context learning. *IEEE Trans Knowl Data Eng*. 2017;29(1):38–56.
94. First JM, Shin H, Ranjit YS, Houston JB. Covid-19 stress and depression: examining social media, traditional media, and interpersonal communication. *J Loss Trauma*. 2021;26(2):101–15.
95. Lacoste A, Luccioni A, Schmidt V, Dandres T. Quantifying the carbon emissions of machine learning. [arXiv:1910.09700](https://arxiv.org/abs/1910.09700). 2019

---

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