

# Strategies to Harness the Transformers' Potential: UNSL at eRisk 2023

Horacio Thompson<sup>1,2</sup>, Leticia Cagnina<sup>1,2</sup> and Marcelo Errecalde<sup>1</sup>

<sup>1</sup>Universidad Nacional de San Luis (UNSL), Ejército de Los Andes 950, San Luis, C.P. 5700, Argentina

<sup>2</sup>Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET), San Luis, Argentina

## Abstract

The CLEF eRisk Laboratory explores solutions to different tasks related to risk detection on the Internet. In the 2023 edition, Task 1 consisted of searching for symptoms of depression, the objective of which was to extract user writings according to their relevance to the BDI Questionnaire symptoms. Task 2 was related to the problem of early detection of pathological gambling risks, where the participants had to detect users at risk as quickly as possible. Finally, Task 3 consisted of estimating the severity levels of signs of eating disorders.

Our research group participated in the first two tasks, proposing solutions based on *Transformers*. For Task 1, we applied different approaches that can be interesting in information retrieval tasks. Two proposals were based on the similarity of *contextualized embedding vectors*, and the other one was based on *prompting*, an attractive current technique of machine learning. For Task 2, we proposed three fine-tuned models followed by decision policy according to criteria defined by an early detection framework. One model presented extended vocabulary with important words to the addressed domain. In the last task, we obtained good performances considering the decision-based metrics, ranking-based metrics, and runtime. In this work, we explore different ways to deploy the predictive potential of *Transformers* in eRisk tasks.

## Keywords

Transformers, Information Retrieval, Prompting, Early Risk Detection

## 1. Introduction

The Early Risk Prediction on the Internet (eRisk) laboratory proposes solving different challenges to explore evaluation methodologies, effectiveness metrics, and practical applications for risk detection in social networks. Through its editions [1, 2, 3, 4, 5, 6, 7, 8], several tasks have been proposed on different domains, promoting the participating teams to propose innovative solutions that solve the tasks in the best possible way. Our research group has actively participated in eRisk editions with notable contributions [9, 10, 11, 12, 13]. In the 2023 edition [8], a new task was introduced: Task 1, which involved searching for symptoms of depression in a collection of user writings. Task 2 was a continuation of the 2022 edition of the problem on early risk detection of pathological gambling. Finally, Task 3 consisted of estimating the severity level of signs of eating disorders.


---

CLEF 2023: Conference and Labs of the Evaluation Forum, September 18–21, 2023, Thessaloniki, Greece

✉ [hjthompson@unsl.edu.ar](mailto:hjthompson@unsl.edu.ar) (H. Thompson); [lcagnina@unsl.edu.ar](mailto:lcagnina@unsl.edu.ar) (L. Cagnina); [merreca@unsl.edu.ar](mailto:merreca@unsl.edu.ar) (M. Errecalde)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

The neural architectures known as *Transformers* proposed by Vaswani et al. [14] have caused a true revolution in the artificial intelligence field. Numerous studies have shown the performance of *Transformers* to solve a wide variety of natural language processing tasks, with models such as BERT [15], GPT-2 [16], and GPT-3 [17]. Motivated by the relevant role of *Transformers*, we have proposed to address this year’s tasks by applying approaches based on this architecture.

Our research group participated in Tasks 1 and 2, focusing on strategies that take advantage of the predictive power of *Transformers*. For Task 1, we presented three proposals. The first two were based on measuring the similarity of embeddings extracted from pre-trained models. We represented the writings of the collection and the symptoms of the BDI Questionnaire using verbs, adjectives, and nouns. From these terms, we obtained word embeddings using a BERT model adjusted for depression tasks, which allowed us to measure the closeness between the writings and the symptoms considering the task domain and then obtain the final rankings. The third solution consisted of applying one of the most attractive approaches currently known as *prompting*, which takes advantage of the predictive power of a pre-trained language model, adapting it to a particular task [18]. We applied the *Fixed-prompt LM Tuning* technique [19] by fitting the RoBERTa model [20] on samples containing pre-defined prompts. Since labeled data were unavailable, we created a dataset of 200 samples per symptom using *ChatGPT* [21]. We created prompts by concatenating each sample to a simple template with slots to fill. Then, we tuned the RoBERTa model to solve the missing word on these prompts, continuing its previous training. We evaluated the texts of the collection by predicting the prompts, and we associated the responses of the model with the probability of belonging to each one of the symptoms. Finally, we improved the ranking of two of the 21 symptoms with a multiclass classifier obtained by *fine-tuning* on the created dataset. For Task 2, we used an early detection framework [22], which led us to remarkable results in previous editions [12, 13]. The method defines that, to solve an early detection problem, it is necessary to consider two components: one dedicated to solving a user classification problem (classification with partial information or CPI), and the other involves a decision policy to decide when to stop evaluating a user (deciding the moment of classification or DMC). On this occasion, we presented three proposals, for which we applied the BERT model with some variants (CPI component) and defined a decision policy based on the history of predictions that a model performs during user evaluation (DMC component).

Our main contribution in this edition was the application of different strategies based on *Transformers*. We applied novel techniques that may be of interest to works related to information retrieval. For the early detection problem, combining fine-tuned models with a decision policy based on a historic rule allowed us to maximize performance and obtain good results.

## 2. Task 1: Search for symptoms of depression

From a collection of user posts, the task consisted of building rankings of relevant posts considering the 21 depression symptoms of the BDI Questionnaire. A *relevant sentence* for a symptom S was defined as the user-generated text that provides information about their condition related to S. According to official data, the collection contains 3,807,115 sentences extracted from 3,107 users. Teams could submit up to five proposals, each with 21 rankings of up to 1,000 writings in descending order by score. Our research group put forward three

solutions: two based on the similarity of contextualized embeddings and one employing the *prompting* technique.

## 2.1. Similarity-based proposals

The strategies consisted of measuring the closeness of the writings to each of the symptoms.

**Sentences filter.** Due to the large number of samples in the collection, a filter was initially applied to reduce the search space. For this, we analyzed the sentiment of the texts with the VADER processor [23], and we selected those whose negative polarity score was greater than 0. This fact allowed us to reduce the search space from almost 4 million writings to 1 million, which meant that 25% had some indication of negativity.

**Preprocessing steps.** Characters were converted to lowercase, while Unicode and HTML codes were transformed into their corresponding symbols. Web pages and numbers were replaced by the *weblink* and *number* tokens, respectively. Repeated words and spaces were also removed, and emojis were manually replaced with text representations related to the symptoms.

**Symptoms and writings representation.** One way to understand a text is to direct the analysis toward important words that describe its semantics [24]. To represent each symptom, ten verbs, ten adjectives, and ten nouns were chosen, which were manually selected based on the information provided in the BDI Questionnaire. To represent writings, we used the *Spacy* parser to extract verbs, adjectives, and nouns. *Stopwords* were discarded, except those words that were present in the symptoms. The next step was to represent the symptoms and writings through *word embeddings* considering the context of the addressed domain. We employed a transformer-based approach instead of non-contextualized methods such as *Word2Vec* [25] and *FastText* [26]. Thus, a context was defined for each word, and a language model was used to extract its embedding vector. For instance, to represent *feeling* in the context of the symptom *sadness*, it was defined as *feeling is linked to the symptom sadness*. Subsequently, embeddings were obtained using the last layer of a BERT model tuned on samples from users with depression<sup>1</sup>.

**Embeddings similarity.** The verbs, adjectives, and nouns of the texts were compared with the verbs, adjectives, and nouns of the symptoms using *cosine similarity*, obtaining a table of scores as follows:

```

∀ text T, symptom S:
  ∀ VT in verbs(T), VS in verbs(S):
    context_VT = "VT is linked to the symptom S"
    emb_VT = extract_embedding(BERT, context_VT)
    context_VS = "VS is linked to the symptom S"
    emb_VS = extract_embedding(BERT, context_VS)
    similarity = similarity(emb_VT, emb_VS)
  Repeat for adjectives and nouns
```

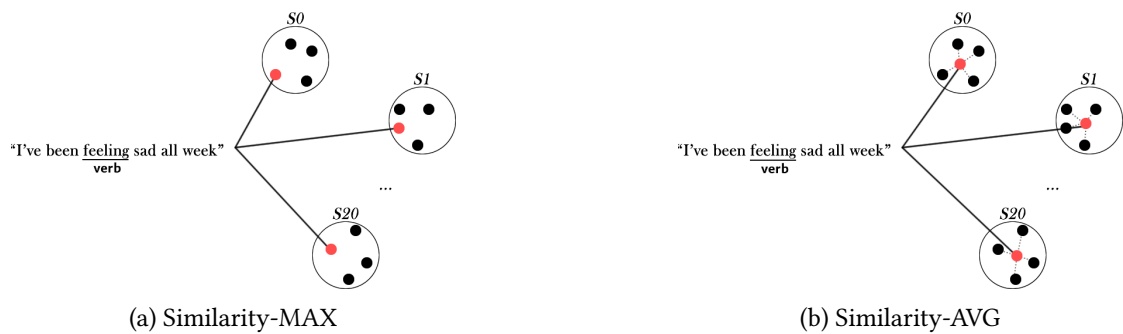
---

<sup>1</sup>Available in: <https://huggingface.co/BitanBiswas/depression-detection-bert>.


**Summary of scores.** To summarize the values of a text according to symptoms, two strategies were considered:

- **Similarity-MAX:** The score of each text is summarized considering the verb, adjective, and noun with the maximum similarity, that is, those closest to the symptom, and then the three values are averaged.
- **Similarity-AVG:** The score of each text is summarized considering the verb, adjective, and noun with the best average, and then the values obtained are averaged. Unlike the previous one, the word that, on average, was closest to the symptom is taken to reach the final score.

**Final rankings.** For both strategies, the final ranking was created by ordering each symptom by score and extracting the first 1,000 writings. Figures 1 and 2 show the graphic scheme and the implementation in a tabular format of similarity-based proposals.



**Figure 1:** Graphic scheme of similarity-based proposals. The distance of a text is observed considering its verb and the verbs of each one of the symptoms (S0, S1,..., S20) for (a) *Similarity-MAX* and (b) *Similarity-AVG*.



**Figure 2:** Tabular scheme of the similarity-based proposals implementation. The similarities between text T and the symptoms are saved in a scores table and summarized according to *Similarity-MAX* and *Similarity-AVG*.

## 2.2. Prompting-based proposals

*Prompting* aims to reformulate the original task as a masked language problem and take advantage of the predictive models' ability to complete the missing words of an input text. The

method is divided into two steps: 1) *create prompts*, where a text fragment with slots to fill (*template*) is added to each original sentence that the model must complete; 2) *derive the final answer*, where the prediction of the model is mapped to an adequate answer for the original task. Finding suitable prompts can be challenging as it influences the solution's success. It can be solved by manually searching for the best prompt [27] or by automatic techniques such as *soft-prompting* [28, 29]. On the other hand, *Fixed-prompt LM Tuning* is based on fixing a prompt and continuing the training of the language model to improve its predictions. In a *few-shot* scenario, large language models are generally used due to the lack of labeled samples but can only be accessed via an API interface. Because of this, and according to Task 1, we decided to use a small and adjustable model based on *Transformers*, applying the *prompting* paradigm through the *Fixed-prompt LM Tuning* technique.

**Sentences filter and preprocessing steps.** As in the similarity-based proposals, writings that presented negative scores were selected to reduce the search space with the VADER processor, and the same preprocessing was also applied.

**Dataset creation.** We created a dataset of 200 samples per symptom to continue training the language model. An application to automatically interact with ChatGPT was implemented as follows:

*Query: Generate a list of user texts linked to S. For reference, use words from this list: [w<sub>1</sub>, w<sub>2</sub>, ...]*  
*Answer: Text<sub>1</sub>, Text<sub>2</sub>, ... , Text<sub>200</sub>.*

Consequently, for each symptom S, we obtained labeled samples <Text<sub>1</sub>, S>, <Text<sub>2</sub>, S>, ... , <Text<sub>200</sub>, S>. The w<sub>i</sub> words were extracted from the terms used to represent the symptoms in the similarity-based proposals.

**Language Model Tuning.** The RoBERTa pre-trained model was imported and tuned based on the masked language problem as follows:

```

∀ sample <T, S>
    label = T + "This is linked to S"
    prompt = T + "This is linked to MASK"
    pred = RoBERTa(prompt)
    loss = CrossEntropy(pred, label)
    ...

```

In this way, the model learned to complete missing words, considering the predefined prompt and the samples associated with the symptoms, directing the model's answers toward the task to be solved. For instance, for the sample <"I've been feeling sad all week", sadness>, the prompt "I've been feeling sad all week. This is linked to MASK" is formed; then, the model should predict that the best word for MASK is *sadness*.

**Evaluation of the writings.** The texts of the collection were evaluated using the *prompting* scheme. We defined the prompt using the same template with which the model was tuned. We also created a words dictionary (*verbalizer*) for mapping model predictions to the final output, including the words used to create the dataset. In this way, the tuned model was used to predict the probability that the writings were linked to each symptom.

**Preliminary ranking.** A table of results was obtained with the first 1,000 writings ordered by probability in decreasing order for each symptom.

When inspecting the writings obtained in the rankings, we observed that samples not related to depression were found in two of the 21 symptoms: *worthlessness* and *loss of energy*. For the former, finance and economic crisis were the main topics, while for the latter, they were renewable energy, and oil, among others. On the other hand, the ranking of *indecisiveness* and *fatigue* had potential writings of *worthlessness* and *loss of energy*, respectively. Because the probabilities are correlated, the possibility of belonging to one symptom affects the rest. Therefore, we improved the quality of these two symptoms with an extension of the proposal, re-evaluating the writings using a classification model.

**Fine-tuning RoBERTa for *worthlessness* y *loss-energy*.** A multiclass classification problem was formulated considering the labels: *worthlessness*, *loss-energy*, and *others*. The training set was defined as follows: for *worthlessness* and *loss-energy*, we used samples extracted from the previously created dataset; for *others*, the remaining symptoms were used, excluding sentences with the *indecisiveness* and *fatigue* labels. Besides, we added texts related to monetary value loss. The *others* class attracted those samples that should not be part of *worthlessness* and *loss-energy*. Then, *fine-tuning* was applied to the RoBERTa model on this training set.

**Classification of the writings.** The texts were evaluated by the classification model, recording the probability of each prediction in a new table. The ranking for *worthlessness* and *loss-energy* was created, and we selected the first 1,000 sentences.

**Final ranking (*Prompting-Classifier*).** The final proposal consisted of 19 symptom rankings using *prompting* and 2 rankings using the *classifier* explained above.

### 2.3. Results

For evaluating the teams, a writing pool for each symptom was built by selecting the first 50 sentences from all proposals (in total, 37). The organizers carried out a labeling process with three assessors that manually chose which writings were relevant [8], resulting in two evaluation schemes: *majority voting* (if 2/3 agreed) and *unanimity* (if 3/3 agreed). Table 1 shows an extract of the result of the labeling process.

**Table 1**

Extract from the count of relevant sentences for each symptom of the BDI Questionnaire. *Original*: a writing pool with the first 50 sentences extracted from proposals of all teams. *Majority voting*: 2 of 3 assessors agreed. *Unanimity*: 3 of 3 agreed.

Symptom	Original	Majority voting	Unanimity
Sadness	1110	318	179
Pessimism	1150	325	104
Past Failure	973	300	160
Loss of Pleasure	1013	204	97
Guilty Feelings	829	143	83
Punishment Feelings	1079	50	21
...	...	...	...

Tables 2 and 3 show the results obtained, considering the *majority voting* and *unanimity* schemes. The best results were achieved by Formula-ML. The mean, median, and performance distribution of all proposals (Figure 3) show that most teams maintained a considerably lower performance than Formula-ML. *Prompting-Classifier* stands out among our proposals with a performance close to that of most teams, considering the AP, R-PREC, and NDCG@1000 metrics.

**Table 2**

Ranking-based evaluation for Task 1. Results are reported according to the metrics Average Precision (AP), R-Precision (R-PREC), Precision at 10 (P@10), and NDCG at 1000 (NDCG@1000), for the majority voting scheme.

Team	Run	Majority voting			
		AP	R-PREC	P@10	NDCG@1000
UNSL	Prompting-Classifier	0.036	0.090	0.229	0.180
UNSL	Similarity-AVG	0.001	0.008	0.010	0.016
UNSL	Similarity-MAX	0.001	0.011	0.019	0.019
Formula-ML	SentenceTransformers_0.25	<b>0.319</b>	<b>0.375</b>	<b>0.861</b>	<b>0.596</b>
Formula-ML	SentenceTransformers_0.1	0.308	0.359	<b>0.861</b>	0.584
	<i>Mean</i>	0.084	0.131	0.428	0.219
	<i>Median</i>	0.065	0.114	0.471	0.180

**Table 3**


Ranking-based evaluation for Task 1 according to the unanimity scheme.

Team	Run	Unanimity			
		AP	R-PREC	P@10	NDCG@1000
UNSL	Prompting-Classifier	0.020	0.063	0.090	0.157
UNSL	Similarity-AVG	0.000	0.005	0.005	0.011
UNSL	Similarity-MAX	0.001	0.006	0.010	0.012
Formula-ML	SentenceTransformers_0.25	0.268	<b>0.360</b>	<b>0.709</b>	<b>0.615</b>
Formula-ML	SentenceTransformers_0.1	<b>0.293</b>	0.350	0.685	0.611
	<i>Mean</i>	0.072	0.118	0.297	0.210
	<i>Median</i>	0.059	0.110	0.333	0.177


Several factors may have influenced the general performance of the models, such as the labeling process or the way of evaluating them. As observed in Table 1, the number of relevant sentences was considerably less than the original number of samples extracted from all the proposals, and, in turn, distant values are observed between *majority voting* and *unanimity* for each symptom. This fact could indicate that the assessors had different opinions when interpreting sentences, probably due to the concept of *relevance*. On the other hand, it would be important to consider the performance of the models at the symptom level since there were significant differences in the number of relevant sentences between symptoms.

### 3. Task 2: Early Detection of Signs of Pathological Gambling

The goal was to detect, as early as possible, users that showed signs of pathological gambling. The challenge was divided into two stages: a *training stage*, where the participants experimented



(a) Majority voting



(b) Unanimity

**Figure 3:** Performance distribution of the proposals presented by all teams for Task 1, considering (a) *Majority voting* and (b) *Unanimity*. The performances of our team are shown with X marks. The *mean* (dashed line) is included.

with data extracted from previous editions, and a *test stage*, where a client application interacted with a server, defining an *early* environment. This last process was divided into rounds in which the client requested the next post of users and, according to the number of predictive models, evaluated them and returned a response to the server.

Early risk detection can be analyzed as a multi-objective problem, where the challenge is to find an adequate balance between the precision in identifying risky users and the minimum time required for that decision to be reliable. In [13], our research group applied the early classification framework [22] by using a BERT model with extended vocabulary (CPI component) and a decision policy based on a historic rule (DMC component). In this edition, we presented three proposals, improving the CPI and DMC components to maximize the performance of the final models.



### 3.1. Datasets

Table 4 shows the detail of the corpora available to solve the task. The eRisk2021 and eRisk2022 corpora were used to train the models, as well as UNSL2021\_train and UNSL2021\_valid created in [12]. The eRisk2023 corpus was used for the organizers to evaluate the participating models. It contains 4.7% of positive users compared to 7% and 3.9% of eRisk2021 and eRisk2022, respectively. Furthermore, the number of words per post was considerably higher than in previous editions, which may be a relevant factor for the models’ performance when evaluating longer posts. On the other hand, the UNSL2021\_train and UNSL2022\_valid corpora have considerably fewer posts than the eRisk corpora.

**Table 4**

Details of the corpora used for Task 2. The corpora of the different eRisk editions are shown, as well as the corpora created by our team in previous editions. The number of users (total, positives, and negatives) and the number of posts of each corpus are reported. The median, minimum, and maximum number of posts per user and words per post in each corpus are detailed.

Corpus	#users			#posts	#posts per user			#words per post		
	Total	Pos	Neg		Med	Min	Max	Med	Min	Max
eRisk2023	2,174	103	2071	1,102,871	327.33	10	2,004	28.9	1	12,779
eRisk2022	2,079	81	1,998	1,177,590	297	3	2,001	11	0	6,728
eRisk2021	2,348	164	2,184	1,130,799	244	10	2,001	11	1	8,241
UNSL2021_train	1,746	286	1,460	158,924	51	31	1,188	20	1	7,479
UNSL2021_valid	1,746	286	1,460	161,204	53	31	1,337	20	1	3,234

### 3.2. CPI components: Models

Each model was trained and validated by combining different corpora with an 85/15 split. We used the BERT model applying the *fine-tuning* process to adjust it to the classification task. A limitation of the BERT architecture is that it only supports 512 input tokens. Thus, we improved the posts extracted from each user by selecting those posts with some indication of negativity by the VADER processor, and the first 512 tokens were taken. We also used a scheduler to automatically adjust the Learning Rate during *fine-tuning*, improving the convergence and performance of the model. Finally, the best model for each proposal was chosen considering the F1 metric over the positive class ( $F_1+$ ).

#### UNSL#0: Classic BERT model

*Training set.* Combination of the eRisk2021, UNSL2021\_train, and UNSL2021\_valid corpora. Extraction of posts with some indication of negativity.

*Preprocessing steps.* Characters were converted to lowercase, while Unicode and HTML codes were transformed into their corresponding symbols. Web pages and numbers were replaced by the *weblink* and *number* tokens, respectively. Repeated words and spaces were also removed.

*Hyperparameters for fine-tuning.* Architecture = ‘BERT-based-uncased’, optimizer = ‘AdamW’, LR = 3E-5, scheduler = ‘LinearSchedulerWarmup’, batch\_size = 8, and n\_epochs = 3.

#### UNSL#1: BERT model with an extended vocabulary

We extended the BERT vocabulary using important words to the addressed domain extracted

from an external model. The SS3 model [30] was trained to classify users on the available corpora, and we selected the first 40 words according to the confidence values on the positive class.

*Training set, preprocessing steps, and hyperparameters for fine-tuning.* The same as the UNSL#0 model.

#### **UNSL#2: Classic BERT model on all available data**

Considering the same hyperparameters of the UNSL#0 model, the *fine-tuning* process was applied using all available data (eRisk2021, eRisk2022, UNSL2021\_train, and UNSL2021\_valid).

### **3.3. DMC component: Decision Policy**

The best decision policy for the models described above was evaluated using a mock server<sup>2</sup>. This tool simulates the eRisk challenge through the rounds of posts and answers submissions, and then it calculates the final results according to the decision and ranking-based metrics. It was useful since the performance of CPI models can drastically change when evaluated in an *early* environment. A client application was defined to manage the interaction with the server. When it receives a round of posts, the system preprocesses the writings, invokes the predictive models (CPI), and applies a decision policy (DMC). To take advantage of the 512 input tokens that the BERT architecture admits, the application uses the last  $N=10$  posts (posts window), linking the current post with previous posts. With the mock server, the client application, and the predictive models, different decision policies were evaluated using the  $F_1+$  and  $F_{latency}$  metrics.

#### **Decision policy based on a historic rule**

The historic rule defines that if the current prediction and the last  $M$  predictions exceed a *threshold* (limit probability to predict a positive user), the client application must issue a risky user alarm; otherwise, it is necessary to continue the user evaluation. In addition, the rule has the *min\_delay* parameter, which defines the moment when it will start to apply. We obtained that the best parameters were *threshold* = 0.7,  $M$  = 10, and *min\_delay* = 10.

*Performance in an early environment.* The models were evaluated using the eRisk2022 corpus. UNSL#0:  $F_1+$  = 0.88 and  $F_{latency}$  = 0.83; UNSL#1:  $F_1+$  = 0.84 and  $F_{latency}$  = 0.81. UNSL#2 was not tested as it was trained with all the corpora, including eRisk2022.

### **3.4. Results**

Table 5 shows the results obtained by our team according to the decision-based metrics. The best results were achieved by ELiRF-UPV#0, as well as other proposals had good results. Considering the average level among all the teams, our models achieved remarkable results in the  $F_1$ , ERDE<sub>50</sub>, and  $F_{latency}$  metrics. The UNSL#1 and UNSL#2 models showed similar performance, outperforming UNSL#0 on the same metrics. Regarding ERDE<sub>5</sub>, the three models obtained the same performance as the mean among all teams.

Regarding the ranking-based metrics, the teams that achieved the best results considering 1, 100, 500, and 1000 posts were ELiRF-UPV, NLP-UNED-2, OBSER-MENH, and UNSL. As can

---

<sup>2</sup>Available in: [https://github.com/jmloyola/erisk\\_mock\\_server](https://github.com/jmloyola/erisk_mock_server).

**Table 5**

Decision-based evaluation results for Task 2. The best team taking into account the  $F_1$ ,  $ERDE_5$ ,  $ERDE_{50}$ , and  $F_{latency}$  is shown (values in bold), as well as the *mean* and *median* values of the results report for CLEF eRisk 2023. The second-best teams are also included.

Model	P	R	$F_1$	$ERDE_5$	$ERDE_{50}$	latencyTP	speed	$F_{latency}$
UNSL#0	0.752	0.767	0.760	0.048	0.017	15.0	0.945	0.718
UNSL#1	0.79	0.806	0.798	0.048	0.014	13.0	0.953	0.761
UNSL#2	0.752	0.854	0.800	0.048	0.013	14.0	0.949	0.759
ELiRF-UPV#0	1.000	0.883	<b>0.938</b>	<b>0.026</b>	<b>0.010</b>	4.0	0.988	<b>0.927</b>
NLP-UNED-2#1	0.957	0.883	0.919	0.034	0.016	13.0	0.953	0.876
NLP-UNED-2#4	0.764	0.883	0.819	0.033	0.010	13.0	0.953	0.781
<i>Mean</i>	0.390	0.796	0.367	0.048	0.035	19.12	0.932	0.362
<i>Median</i>	0.092	0.903	0.125	0.047	0.042	8.00	0.973	0.162

be seen in Table 6, our team obtained the best results for the P@10 and NDCG@10 metrics. For NDCG@100, acceptable values were obtained, mainly with 100 posts. Besides, as in the decision-based metrics, UNSL#1 and UNSL#2 achieved similar results, outperforming UNSL#0.

**Table 6**

Ranking-based evaluation results for Task 2. Results are reported according to the three classification metrics obtained after processing 1, 100, 500, and 1000 posts, respectively.

Ranking	Metric	UNSL#0	UNSL#1	UNSL#2
1 post	P@10	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
	NDCG@10	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
	NDCG@100	0.46	0.57	0.55
100 posts	P@10	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
	NDCG@10	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
	NDCG@100	0.70	0.78	0.75
500 posts	P@10	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
	NDCG@10	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
	NDCG@100	0.64	0.67	0.69
1000 posts	P@10	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
	NDCG@10	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
	NDCG@100	0.64	0.70	0.69

Finally, Table 7 shows the total time spent by each team to solve the task. Our team was the fastest, presenting three models with a final delay of *1 day and 2 hours*, followed by the ELiRF-UPV's model (almost *11 hours* of difference). The rest of the teams used five models to solve the task, with delays ranging from *4 to 54 days*.

**Table 7**

Total time spent by each team for Task 2. The team name, number of models, and number of user posts processed are shown. The teams are displayed according to the total time.

<b>Team</b>	<b>#models</b>	<b>#posts processed</b>	<b>Total time</b>
<b>UNSL</b>	<b>3</b>	<b>2004</b>	<b>1 day + 2h:17m</b>
ELiRF-UPV	1	2004	1 day + 13h:3m
Xabi_EHU	5	2004	4 days + 23h:52m
OBSER-MENH	5	2004	6 days + 3h:56m
RELAI	5	764	6 days + 9h:12m
NLP-UNED-2	5	2004	7 days + 4h:24m
NUS-eRisk	5	2004	9 days + 14h:39m
BioNLP-IISERB	5	61	10 days + 0h:49m
SINAI	5	809	10 days + 13h:0m
UMUTeam	5	2004	14 days + 0h:29m
NLP-UNED	5	1151	54 days + 19h:27m

## 4. Conclusion

In this article, the UNSL team solved Tasks 1 and 2 of the eRisk 2023 Laboratory. For Task 1, we applied different approaches, obtaining better results with the proposal based on *prompting*. Although we did not get the best results for this task, it would be interesting to improve these approaches considering the criteria used in the labeling process. For Task 2, we obtained outstanding results in all evaluation metrics, applying a model with extended vocabulary and a decision policy based on a historic rule. Our proposals harnessed the predictive potential of *Transformers*, demonstrating that these architectures can be used in information retrieval tasks and problems of early risk detection.

## References

- [1] D. E. Losada, F. Crestani, A test collection for research on depression and language use, in: Proc. of Conference and Labs of the Evaluation Forum (CLEF 2016), Evora, Portugal, 2016, pp. 28–39.
- [2] D. E. Losada, F. Crestani, J. Parapar, erisk 2017: Clef lab on early risk prediction on the internet: experimental foundations, in: International Conference of the Cross-Language Evaluation Forum for European Languages, Springer, 2017, pp. 346–360.
- [3] D. E. Losada, F. Crestani, J. Parapar, Overview of erisk: early risk prediction on the internet, in: International Conference of the Cross-Language Evaluation Forum for European Languages, Springer, 2018, pp. 343–361.
- [4] D. E. Losada, F. Crestani, J. Parapar, Overview of erisk 2019 early risk prediction on the internet, in: Experimental IR Meets Multilinguality, Multimodality, and Interaction: 10th International Conference of the CLEF Association, CLEF 2019, Lugano, Switzerland, September 9–12, 2019, Proceedings 10, Springer, 2019, pp. 340–357.
- [5] D. E. Losada, F. Crestani, J. Parapar, erisk 2020: Self-harm and depression challenges, in:

- Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14–17, 2020, Proceedings, Part II 42, Springer, 2020, pp. 557–563.
- [6] J. Parapar, P. Martín-Rodilla, D. E. Losada, F. Crestani, Overview of erisk 2021: Early risk prediction on the internet, in: International Conference of the Cross-Language Evaluation Forum for European Languages, Springer, 2021, pp. 324–344.
- [7] J. Parapar, P. Martín-Rodilla, D. E. Losada, F. Crestani, Overview of erisk 2022: Early risk prediction on the internet, in: Experimental IR Meets Multilinguality, Multimodality, and Interaction: 13th International Conference of the CLEF Association, CLEF 2022, Bologna, Italy, September 5–8, 2022, Proceedings, Springer, 2022, pp. 233–256.
- [8] J. Parapar, P. Martín-Rodilla, D. E. Losada, F. Crestani, Overview of eRisk 2023: Early Risk Prediction on the Internet, in: Experimental IR Meets Multilinguality, Multimodality, and Interaction. 14th International Conference of the CLEF Association, CLEF 2023, Springer International Publishing, Thessaloniki, Greece, 2023.
- [9] M. L. Errecalde, M. P. Villegas, D. G. Funez, M. J. G. Ucelay, L. C. Cagnina, Temporal variation of terms as concept space for early risk prediction., in: Clef (working notes), 2017.
- [10] D. G. Funez, M. J. G. Ucelay, M. P. Villegas, S. Burdisso, L. C. Cagnina, M. Montes-y Gómez, M. Errecalde, Unsl’s participation at erisk 2018 lab., in: CLEF (Working Notes), 2018.
- [11] S. G. Burdisso, M. Errecalde, M. Montes-y Gómez, Unsl at erisk 2019: a unified approach for anorexia, self-harm and depression detection in social media., in: CLEF (Working Notes), 2019.
- [12] J. M. Loyola, S. Burdisso, H. Thompson, L. C. Cagnina, M. Errecalde, Unsl at erisk 2021: A comparison of three early alert policies for early risk detection., in: CLEF (Working Notes), 2021, pp. 992–1021.
- [13] J. M. Loyola, H. Thompson, S. Burdisso, M. Errecalde, Unsl at erisk 2022: Decision policies with history for early classification (2022).
- [14] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, *Advances in neural information processing systems* 30 (2017).
- [15] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, *arXiv preprint arXiv:1810.04805* (2018).
- [16] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, et al., Language models are unsupervised multitask learners, *OpenAI blog* 1 (2019) 9.
- [17] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, *Advances in neural information processing systems* 33 (2020) 1877–1901.
- [18] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, G. Neubig, Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing, *ACM Computing Surveys* 55 (2023) 1–35.
- [19] X. L. Li, P. Liang, Prefix-tuning: Optimizing continuous prompts for generation, *arXiv preprint arXiv:2101.00190* (2021).
- [20] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, Roberta: A robustly optimized bert pretraining approach, *arXiv preprint arXiv:1907.11692* (2019).

- [21] OpenAI, Chatgpt, <https://github.com/openai/chatgpt>, 2021. Accessed: May 29, 2023.
- [22] J. M. Loyola, M. L. Errecalde, H. J. Escalante, M. Montes y Gomez, Learning when to classify for early text classification, in: *Computer Science–CACIC 2017: 23rd Argentine Congress*, La Plata, Argentina, October 9-13, 2017, Revised Selected Papers 23, Springer, 2018, pp. 24–34.
- [23] C. Hutto, E. Gilbert, 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*, volume 8, 2014, pp. 216–225.
- [24] Z. Lin, M. Feng, C. N. d. Santos, M. Yu, B. Xiang, B. Zhou, Y. Bengio, A structured self-attentive sentence embedding, *arXiv preprint arXiv:1703.03130* (2017).
- [25] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space, *arXiv preprint arXiv:1301.3781* (2013).
- [26] P. Bojanowski, E. Grave, A. Joulin, T. Mikolov, Enriching word vectors with subword information, *Transactions of the association for computational linguistics* 5 (2017) 135–146.
- [27] F. Petroni, T. Rocktäschel, P. Lewis, A. Bakhtin, Y. Wu, A. H. Miller, S. Riedel, Language models as knowledge bases?, *arXiv preprint arXiv:1909.01066* (2019).
- [28] G. Qin, J. Eisner, Learning how to ask: Querying lms with mixtures of soft prompts, *arXiv preprint arXiv:2104.06599* (2021).
- [29] B. Lester, R. Al-Rfou, N. Constant, The power of scale for parameter-efficient prompt tuning, *arXiv preprint arXiv:2104.08691* (2021).
- [30] S. G. Burdisso, M. Errecalde, M. Montes-y Gómez, A text classification framework for simple and effective early depression detection over social media streams, *Expert Systems with Applications* 133 (2019) 182–197.