

SINAI at eRisk@CLEF 2023: Approaching Early Detection of Gambling with Natural Language Processing

Notebook for the Early Risk Prediction on the Internet Lab at CLEF 2023

Alba María Mármol-Romero¹, Flor Miriam Plaza-del-Arco² and Arturo Montejo-Ráez¹

¹Computer Science Department, SINAI, CEATIC, Universidad de Jaén, 23071, Spain

²Bocconi University, Via Sarfatti 25, 20100, Milan, Italy

Abstract

This paper describes the participation of the SINAI team in the eRisk@CLEF lab. Specifically, one of the proposed tasks has been addressed: Task 2 on the early detection of signs of pathological gambling. The approach presented in Task 2 is based on pre-trained models from Transformers architecture with comprehensive preprocessing data and data balancing techniques. Moreover, we integrate Long-short Term Memory (LSTM) architecture with automodels from Transformers. In this Task, our team has been ranked in seventh position, with an F1 score of 0.126, out of 49 participant submissions and achieves the highest values in recall metrics and metrics related to early detection.

Keywords

Early risk prediction, Gambling detection, Natural Language Processing, Transformers, LSTM

1. Introduction

The large amount of content posted daily on social media has made them a significant source of data for the early detection of mental disorders and risky behaviours. The eRisk@CLEF 2023 lab [1] focuses on early risk prediction on the Internet and its goal is to promote the development of automatic systems for the detection of mental disorders such as depression, self-harm or eating disorders. In this edition, three tasks have been proposed:

- Task 1: Search for symptoms of depression. It is a new task consisting of ranking sentences from a collection of user writings according to their relevance to a depression symptom. Then, the participants will have to provide rankings for the 21 symptoms of depression from the BDI Questionnaire
- Task 2: Early Detection of Signs of Pathological Gambling. It involves sequentially processing writings and detecting as early as possible the first signs of pathological

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

✉ amarmol@ujaen.es (A. M. Mármol-Romero); flor.plaza@unibocconi.it (F. M. Plaza-del-Arco); amontejo@ujaen.es (A. Montejo-Ráez)

🆔 0000-0001-7952-4541 (A. M. Mármol-Romero); 0000-0002-3020-5512 (F. M. Plaza-del-Arco); 0000-0002-8643-2714 (A. Montejo-Ráez)

© 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)

gambling. It is a continuation of Task 1 proposed for eRisk 2021 and for eRisk 2022, but the difference is that in this edition a larger amount of training data has been provided. The training data of this edition comprises all test users of the 2021 and 2022 tasks.

- Task 3: Measuring the severity of the signs of Eating Disorders. Its aim is to estimate a user's level of disordered eating from his or her history of posts. For this purpose, for each user, a standard eating disorder questionnaire (EDE-Q) has to be filled in.

Currently, our research group SINAI¹ is working on the PRECOM project² focused on the early detection of gambling addiction risk behaviour. Therefore, our interest in developing systems as those expected to answer eRisk tasks is high, as it is a perfect playground to test our approaches.

In this sense, our main goals are not only to produce systems reporting high performance but to understand the best methods and approaches that can be applied in similar scenarios. It is not our aim to put as many features and as many systems as possible all together in an ensemble of predictors to gain the top ranking position, but rather to find out the best approaches that can be applied to our project's objectives. The design of online and monitoring tools, as is requested in Task 1, along with the ability to understand user's disorder, the main pursuit in Task 3, fully matches our research interests.

This work presents the participation of our research group, the SINAI team, in *Task 2: Early Detection of Signs of Pathological Gambling*.

The rest of the paper is organized as follows. Section 2 describes the details of our participation in Task 2. It is divided into subsections in which, first, we introduce what the task consists of, the data provided and the evaluation measures used. Secondly, the system developed and the methodology used is presented. Thirdly, the experimental setup is detailed. Subsequently, the results obtained and a discussion of them are presented. Finally, Section 3 shows the conclusions obtained after participation in the eRisk lab and the perspectives for future work.

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

2.1. Task description

This particular task focuses on identifying signs of gambling addiction at an early stage by analyzing posts from social media in the exact order they were published. The participating systems were required to read the posts, process them, and generate a response to proceed to the next set of posts. The dataset used for this task consists of 33,719 posts from 103 individuals who were categorized as having a positive association with gambling addiction, as well as approximately 1 million posts from 2,071 individuals who were not categorized as addicts [2].

The task is approached from two different perspectives: as a binary decision problem and as a ranking (regression) decision problem. In the binary decision problem, the posts need to be classified as either positive (label 1, indicating the presence of addiction) or negative (label 0, indicating no addiction). The earlier the system detects the presence of addiction, the better it performs, as reflected by the evaluation metrics proposed by the organizers, namely

¹<https://sinai.ujaen.es>

²<https://precom.ujaen.es/>

ERDE and $F_{latency}$, along with well-known measures such as precision, recall, and F1 scores. In the ranking decision problem, instead of assigning binary labels, the system computes a score representing the estimated risk of developing a gambling addiction. Various metrics commonly used in information retrieval, such as P@10 or NDCG, are employed to evaluate this alternative perspective of the task.

2.2. System and methods

To tackle this task, we adopted a supervised learning approach. Our models were trained using the provided training dataset from the eRisk organizers, which combined data from the years 2021 and 2022. This dataset consists of a time series of posts authored by various users. Our main contributions in this task have been the data processing work done as well as the handling of the unbalanced data and the development of a new system that integrates a transformer with an LSTM network.

We performed two different methods of pre-processing on the training data: one comprehensive method and another lighter method tailored to the specific system structure where it is employed. The lighter method simply replaces URLs with tags, while the comprehensive method involves a series of steps. Here is a detailed explanation of the extensive processing we conducted:

- We replaced all numeric HTML entities in the post string with their corresponding Unicode character equivalents.
- Users, URLs, and emails mentioned in the messages were identified and replaced with #USER, #URL, and #EMAIL, respectively.
- We removed any empty characters and replaced special Unicode white-space characters that may have been present in the text.
- Emojis were substituted with strings using the emoji library³.
- If more than seven identical punctuation marks appeared consecutively, they were reduced to only three instances.
- Markdown formatting was removed from the text.
- We standardized the use of double and single quotes in the text string, replacing curved quotes and curved single quotes with straight double and single quotes, respectively.

For our participation in eRisk Task 2, we trained a total of five models. Two systems utilized the XLM-Roberta-Large [3] model, while the other three employed the RoBERTa-Large [4] linguistic model. Four of these systems followed the architecture depicted in Figure 1.

In these four systems, although the models used and the datasets differed, they followed the same structure. The prediction process for Task 2 involved post extraction. All posts by each user were concatenated, and if the concatenated posts exceeded 500 words, the subject was split into sub-subjects. This ensured that the maximum word limit was not exceeded. The concatenated posts served as the input to either the RoBERTa-Large automodel or the XLM-RoBERTa Large automodel. Subsequently, a Feed Forward Neural Network (FFNN) with a

³<https://pypi.org/project/emoji/>

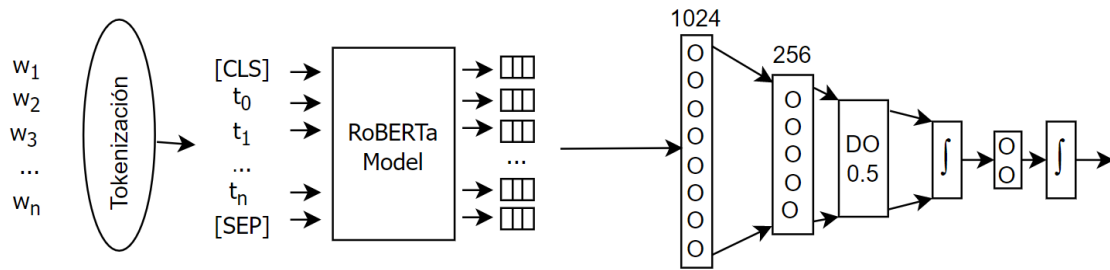


Figure 1: Model architecture for the binary decision task involving the RoBERTa-Large model or the XLM-RoBERTa Large model. The input representation is based on RoBERTa-Large tokenization or XLM-RoBERTa Large tokenization. Each embedding of the model output is passed through a feed-forward neural network with a hidden layer, generating a binary decision and a score as the logits.

single hidden layer was applied to generate the final predictions. The initial layer of the FFNN had an input dimension of 1024 from the embeddings vector generated by RoBERTa. The output size of this layer was set to 256. During the learning process, these outputs were passed through a dropout layer with a probability of 0.5. Finally, a ReLU activation function was applied before feeding the last FFNN with 256 inputs and 2 outputs.

However, the most innovative system among the others is the one that combines the RoBERTa-Large model with the LSTM architecture, as illustrated in Figure 2.

Regarding the system that integrates RoBERTa-Large and the LSTM architecture, the process begins by concatenating the last 50 posts from each subject. This is because we have studied this number is the most appropriate according to this type of dataset as we saw in our last year participation [5]. This concatenation is subsequently split into rounds, representing the input to the RoBERTa-Large automodel. This approach allows us to simulate the sequencing that will occur during the testing phase. The model's output is then passed to an LSTM layer. The output of the LSTM layer is further processed by a Feed Forward Neural Network (FFNN) with a single hidden layer, generating the final predictions. The initial layer of the FFNN has an input dimension of 1024, and the output of this layer is of size 256. These outputs are subjected to a dropout layer (with a probability of 0.5) during the learning process. Finally, a ReLU activation function is applied before feeding the last FFNN with 256 inputs and 2 outputs. Moreover, the cell state and hidden state are passed to the next round in each batch.

The systems were implemented using Python packages such as scikit-learn [6], Transformers [7], and PyTorch [8].

The training was performed on a 2xGPU NVIDIA V100 server.

2.3. Experimental setup

First, all the parameters established in these systems and presented in this document are the result of a previous hyperparameter optimisation search carried out with Optuna⁴, including the dataset selected for data balancing. This means that given a dataset with a number of subjects, these subjects were divided into further subjects due to the sequence length limit given by the

⁴<https://optuna.org/>

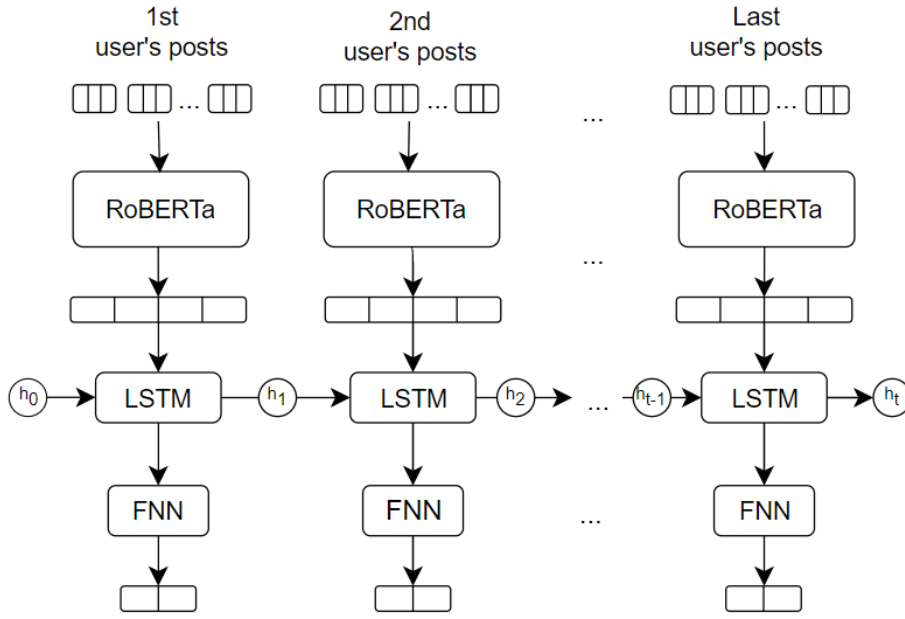


Figure 2: Model architecture for the binary decision task involving the RoBERTa model and LSTM architecture. The input representation consists of RoBERTa-Large tokenization of each subject’s message in each round. Each embedding of the model output is passed through an LSTM layer. The output of the LSTM is then fed into a feedforward neural network with one hidden layer, producing a binary decision.

Transformers models. For this new set, only a part of the subjects belonging to the control group was chosen due to the initial imbalance of the data. Table 1, shows the information concerning the search space. However, for the latest system that includes the LSTM (run 4), the parameters were chosen according to what our hardware architecture and training time limit allowed.

Table 1

Search space in the hyperparameter optimisation process. In this case, with the variable "Data" we refer to the number of sub-subjects belonging to the control group we are working with (in the order from the first generated to the last).

Parameter	Search space
Learning rate	[1e-4, 1e-5, 1e-6]
Batch size	[8, 16]
Epoch	[1, 2, 3, 4]
Weight decay	[0.1, 0.01, 0.001]
Max length	[256, 512]
Data	[5, 10, 15, 20, All]

After the search for hyperparameters, in the learning phase, a maximum length of 512 tokens per document was set. For runs 0 to 3, the document is the concatenation of an undetermined

number of posts whose sum of their words is no bigger than 500 and the optimizer used was AdamW. For run four, the document is the concatenation of the last 50 posts of each subject and the optimizer used was also AdamW. In table 2 there is a summary of hyperparameters and models used in each run.

Table 2
Summary of hyperparameters and models used in each run

Run	Batch size	Learning rate	Weight decay	Epochs	Model
Run 0	8	1e-5	0.01	3	RoBERTa-Large
Run 1	8	1e-5	0.01	2	RoBERTa-Large
Run 2	8	1e-5	0.01	3	XLM-RoBERTa-Large
Run 3	8	1e-5	0.01	2	XLM-RoBERTa-Large
Run 4	8	1e-6	0	7	RoBERTa-Large & LSTM

- **Run 0.** This run utilized a binary classification model based on RoBERTa-Large. The document used for training was the concatenation of an undetermined number of posts, where the sum of their words did not exceed 500. However, due to the unbalanced nature of the provided data, only the first 15 sub-subjects generated by our preprocessing, with no risk of gambling pathology, and all the sub-subjects generated with the risk of gambling pathology were considered. The training dataset consisted of 14.78% positive subjects and 85.22% negative subjects. The model’s output score determined whether a user was considered a potential gambler or not. Users with a score lower than 0.5 were not considered potential gamblers, while those with a score equal to or higher than 0.5 were classified as potential gamblers.
- **Run 1.** Similar to Run 0, this run employed a binary classification model based on RoBERTa-Large. The training dataset used the same structure as the other systems, considering the first 10 sub-subjects generated by our preprocessing with no risk of suffering gambling pathology, and all the sub-subjects generated with the risk of gambling pathology. The training dataset for this run contained 18.84% positive subjects and 81.16% negative subjects.
- **Run 2.** In this run, the model used was XLM-RoBERTa-Large, while maintaining the other settings from Run 0.
- **Run 3.** Similar to Run 1, this run utilized the XLM-RoBERTa-Large model instead of RoBERTa-Large, while keeping the remaining configurations the same.
- **Run 4.** This run involved a binary classification model based on RoBERTa-Large and the LSTM architecture. The document used for training differed from the other systems due to the longer training time required. Only the last 50 concatenated posts were considered, with the only preprocessing step being the replacement of URLs with #URL. Additionally, 500 negative subjects were randomly selected from the unbalanced training dataset. The training dataset for this run comprised 245 positive subjects (the sum of positive subjects from the previous year’s data) and 500 negative subjects. The model’s output score determined whether a user was considered a potential gambler or not. Users with a score lower than 0.5 were not considered potential gamblers, while those with a score equal to or higher than 0.5 were classified as potential gamblers.

2.4. Results and discussion

Salient results have been achieved with the approaches explored by our team. From the reported results provided by the organizers, we have extracted our scores, which are shown in Tables 3 and 4.

Table 3

Results of SINAI team for Task 2 in decision-based evaluation

Run	P	R	$F1$	$ERDE_5$	$ERDE_{50}$	$latency_{tp}$	$speed$	$latency_w F1$
0	0.115	1.000	0.206	0.029	0.021	1.0	1.000	0.206
1	0.124	1.000	0.221	0.028	0.020	2.0	0.996	0.220
2	0.108	1.000	0.195	0.030	0.022	1.0	1.000	0.195
3	0.126	1.000	0.224	0.029	0.020	2.0	0.996	0.223
4	0.092	0.981	0.168	0.044	0.027	3.0	0.992	0.166

In decision-based evaluation, the value of the Recall score obtained by our Runs 0, 1, 2 and 3 (1.000) is the highest among all submissions by participants (49 submissions were reported in total). Our ERDE values are the second and fifteenth highest (0.020 and 0.028). In terms of $speed$ and $latency_{TP}$, we also reach top values for run 0 and run 2. Thus, it seems that a larger unbalance of the data and a larger number of epochs contribute slightly to the improvement in prediction. In general, the systems developed demonstrate a high capacity to correctly identify the majority of positive subjects at a fast speed.

Table 4

Results of SINAI team for Task 2 in ranking-based evaluation

Run	1 writing			100 writing			500 writing		
	P@10	NDCG@10	NDCG@100	P@10	NDCG@10	NDCG@100	P@10	NDCG@10	NDCG@100
0	1.00	1.00	0.72	1.00	1.00	0.88	1.00	1.00	0.85
1	1.00	1.00	0.73	1.00	1.00	0.90	1.00	1.00	0.85
2	1.00	1.00	0.71	1.00	1.00	0.87	1.00	1.00	0.84
3	1.00	1.00	0.72	1.00	1.00	0.89	1.00	1.00	0.86
4	0.80	0.86	0.53	0.90	0.94	0.56	0.70	0.80	0.47

Regarding ranking-based evaluation (see Table 4), our team reaches one of the highest positions. High NDCG and P@k values indicate that the systems are capable of providing relevant and quality recommendations.

The differences in the systems developed according to the metrics appear to be minimal. This makes sense since the systems have almost the same structure. The only different system presented is used in run 4, which has slightly lower values than the rest of the systems, which can be understood by the resource limitations that have been taken into account due to the time required for execution. However, it is the one that we believe has the most promising results with the most research time invested.

3. Conclusions and future work

This paper describes our participation as SINAI team in Task 2 of the eRisk@CLEF 2023 edition. The former is the continuation of the first edition in 2021 and the second edition in 2022 and aims to detect signs of pathological gambling as soon as possible. For Task 2, we have developed only classification models using state-of-the-art pre-trained language models based on Transformers. Besides, we explored LSTM architecture's integration with this model type.

In future work, we plan to analyze in depth the use of the LSTM architecture for this kind of problem as the main objective is early detection in a sequential way. We plan to test different training strategies with the integration of a transformers model together with the LSTM architecture. In addition, we expect to continue to work on data processing and the imbalance present in the data.

Acknowledgments

This work has been partially supported by WeLee project (1380939, FEDER Andalucía 2014-2020) funded by the Andalusian Regional Government, and projects CONSENSO (PID2021-122263OB-C21), MODERATES (TED2021-130145B-I00), SocialTOX (PDC2022-133146-C21) funded by Plan Nacional I+D+i from the Spanish Government, and project PRECOM (SUBV-00016) funded by the Ministry of Consumer Affairs of the Spanish Government.

References

- [1] J. Parapar, P. Martín-Rodilla, D. E. Losada, F. Crestani, erisk 2023: Depression, pathological gambling, and eating disorder challenges, in: J. Kamps, L. Goeuriot, F. Crestani, M. Maistro, H. Joho, B. Davis, C. Gurrin, U. Kruschwitz, A. Caputo (Eds.), *Advances in Information Retrieval*, Springer Nature Switzerland, Cham, 2023, pp. 585–592.
- [2] J. Parapar, P. M. Rodilla, D. Losada, F. Crestani, Overview of eRisk 2023: Early Risk Prediction on the Internet. , in: *Experimental IR Meets Multilinguality, Multimodality, and Interaction. Proceedings of the 14th International Conference of the CLEF Association, CLEF 2023*, Springer International Publishing, Thessaloniki, Greece, 2023.
- [3] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, V. Stoyanov, Unsupervised cross-lingual representation learning at scale, *CoRR abs/1911.02116* (2019). URL: <http://arxiv.org/abs/1911.02116>. arXiv: 1911. 02116.
- [4] 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, *CoRR abs/1907.11692* (2019). URL: <http://arxiv.org/abs/1907.11692>.
- [5] A. M. Mármol-Romero, S. M. Jiménez-Zafra, F. M. Plaza-del Arco, M. D. Molina-González, M.-T. Martín-Valdivia, A. Montejo-Ráez, Sinai at erisk@ clef 2022: Approaching early detection of gambling and eating disorders with natural language processing (2022).
- [6] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al., Scikit-learn: Machine learning in python, *Journal of machine learning research* 12 (2011) 2825–2830.

- [7] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, A. M. Rush, Transformers: State-of-the-art natural language processing, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Association for Computational Linguistics, Online, 2020, pp. 38–45. URL: <https://aclanthology.org/2020.emnlp-demos.6>. doi:10.18653/v1/2020.emnlp-demos.6.
- [8] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, Pytorch: An imperative style, high-performance deep learning library, in: Advances in Neural Information Processing Systems 32, Curran Associates, Inc., 2019, pp. 8024–8035. URL: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.