

Early Detection of Depression and Eating Disorders in Spanish: UNSL at MentalRiskES 2023

Horacio Thompson^{1,2}, Marcelo Errecalde¹

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

²Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET), San Luis, Argentina

Abstract

MentalRiskES is a novel challenge that proposes to solve problems related to early risk detection for the Spanish language. The objective is to detect, as soon as possible, Telegram users who show signs of mental disorders considering different tasks. Task 1 involved the users' detection of eating disorders, Task 2 focused on depression detection, and Task 3 aimed at detecting an unknown disorder. These tasks were divided into subtasks, each one defining a resolution approach.

Our research group participated in subtask A for Tasks 1 and 2: a binary classification problem that evaluated whether the users were positive or negative. To solve these tasks, we proposed models based on Transformers followed by a decision policy according to criteria defined by an early detection framework. One of the models presented an extended vocabulary with important words for each task to be solved. In addition, we applied a decision policy based on the history of predictions that the model performs during user evaluation.

For Tasks 1 and 2, we obtained the second-best performance according to rankings based on classification and latency, demonstrating the effectiveness and consistency of our approaches for solving early detection problems in the Spanish language.

Keywords

Early Risk Detection, Classification Problem, Transformers, Decision Policy

1. Introduction

According to the World Health Organization, one in every eight people worldwide suffers a mental disorder. Anxiety, depression, bipolar disorder, and eating behavior disorders are the most frequent [1]. Different social networks have become mass media chosen by people to share information and express their emotions. Several studies show the relationship between the use of social networks and mental disorders [2, 3, 4]. Therefore, there is a growing interest in the early identification of users suffering from these disorders to provide them with appropriate help. Evaluation conferences such as CLEF eRisk have promoted research groups to solve early detection challenges considering different domains [5, 6, 7, 8, 9, 10, 11]. However, there is currently no challenge of these characteristics for Spanish, highlighting the urgent need to promote initiatives in this language.

MentalRiskES is a novel challenge that proposes to solve problems of early risk detection of mental disorders in Spanish [12]. In this first edition, three tasks were defined with the same

IberLEF 2023, September 2023, Jaén, Spain

✉ hjthompson@unsl.edu.ar (H. Thompson); 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)

objective: to detect Telegram users who show signs of mental disorders as early as possible. Task 1 consisted of the detection of users with eating disorders, Task 2 was related to depression detection, and Task 3 to the detection of an unknown disorder. Each task was divided into subtasks for solving the problem considering different approaches:

Binary classification (subtask A): To decide whether or not a user suffers from a mental disorder by considering the positive and negative classes.

Simple regression (subtask B): To provide an affectation probability on positive and negative classes.

Multi-class classification (subtask C): To decide whether a user suffers from a mental disorder and evaluate their attitude towards it by considering additional classes.

Multi-output regression (subtask D): To provide a confidence probability for the additional classes.

Subtasks A and B were defined for Tasks 1, 2, and 3, while subtasks C and D were also included for Task 2.

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. Our research group achieved notable results in the 2021 [13], 2022 [14], and 2023 (article currently under review) editions of the CLEF eRisk. In these last two editions, we used an early detection framework [15] which defines that 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). In particular, we applied the framework by using a BERT model [16] with an extended vocabulary (CPI component) and a decision policy based on a historic rule (DMC component). In this first edition of MentalRiskES, we participated in Tasks 1 and 2 according to subtask A. Following our Transformers-based approaches, we used the BETO model [17], a variant of BERT that was trained on large Spanish corpora, and we adjusted the historic rule according to the tasks to be solved.

The present work describes the approaches used by our research group to solve Tasks 1 and 2. Section 2 details the datasets, classification models, and decision policies applied. Section 3 shows the results obtained in both tasks and Section 4 presents the conclusions and future works.

2. Resolution Method

The challenge was divided into two stages: a *training stage*, where the participants experimented with data provided by the Organizers, and a *test stage*, where a client application interacted with a server, defining an early environment. This last process was carried out in rounds, during 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.

2.1. Datasets

Three corpora were available to solve Tasks 1 and 2, as shown in Table 1. The *Train* and *Trial* corpora were available for the participants to implement their proposals. The *Trial* corpora were proposed to test the connection between the client application and the server. The *Test* corpora were used for the Organizers to evaluate the participating models. It should be noted that, for both tasks, the number of corpora samples is limited. In the training stage, 185 samples between *Train* and *Trial* were available, while in the test stage, 150 users were evaluated. Furthermore, in contrast to what typically occurs in these classification problems, the classes exhibit a relatively balanced distribution, as evidenced by the number of positive and negative users. On the other hand, for Tasks 1 and 2, the median number of posts per user is approximately 21 and 31, respectively. This fact is relevant because a model with an acceptable performance should finish the evaluation of the users in a smaller number of posts. The maximum number of posts per user in the *Test* corpus indicates the number of total rounds that the test stage had: 50 for Task 1 and 100 for Task 2. Finally, it is observed that the posts were relatively short (between 8 and 9 words per post for each task).

Table 1

Details of the corpora of Tasks 1 and 2. The number of users (total, positives, and negatives) and the number of posts in 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
Task 1	Train	175	74	101	5931	35.0	11	50	9.0	2	899
	Trial	10	5	5	389	48.5	18	50	9.0	3	753
	Test	150	64	86	4179	21.5	11	50	9.0	2	894
Task 2	Train	175	94	81	6248	26.0	11	100	9.0	1	783
	Trial	10	6	4	624	68.0	11	100	9.0	3	201
	Test	149	68	81	5164	31.0	11	100	8.0	1	368

2.2. CPI components: Models

Training set. Due to the limited data available, we augmented the number of samples. For each user, we divided the list of posts into three equal parts according to the list length. Each portion was labeled using the user’s label and added to the training set. In this way, we obtained approximately 500 new samples to train the models. Besides, this allowed the models to be trained considering different contexts of the users’ history and trying to overcome the limitation of BERT architectures that only admit 512 input tokens. Then, each model was trained and validated using an 85/15 split of the *Train* corpus with added samples.

Preprocessing. Some preprocessing actions were performed before the *fine-tuning* process. 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.

Classifier type. We used the BETO model (version: *dccuchile/bert-base-spanish-wwm-uncased*), applying the *fine-tuning* process to adjust it to each task. Different hyperparameters were considered, and a scheduler was used to automatically adjust the learning rate during *fine-tuning*, improving the model convergence. For Tasks 1 and 2, we presented two proposals:

- **Classic BETO model.** We imported the pre-trained model and applied the *fine-tuning* process. It was a baseline model.
- **BETO model with an extended vocabulary.** Important words were added according to the task to be solved. They were extracted from an external model known as SS3 [18]. We trained SS3 to classify users on the available corpora, and we selected the best words according to the confidence values on the positive class. For Task 1, *ayuno* (fasting), *cals* (calories acronym), *atracones* (binge eating), and for Task 2, *decepcionada* (disappointed), *suicidarme* (to commit suicide), and *daño* (damage) are some examples of important words. We evaluated the number of words added to each model in a range of 5 to 50 by considering the validation performances.

Finally, the best CPI model for each proposal was chosen according to the F1 metric over the positive class (F1+). Table 2 shows a summary of the hyperparameters selected for each task.

Table 2

Hyperparameters of the models presented for Tasks 1 and 2.

	Team#Run	Model type	#Batch size	Learning rate	#Epochs	#Added words	Optimizer	Scheduler
Task 1	UNSL#0	Classic BETO	8	3E-5	5	-	AdamW	Linear Scheduler Warmup
	UNSL#1	BETO with extended vocabulary	8	2E-5	3	25		
Task 2	UNSL#0	Classic BETO	8	2E-5	5	-		
	UNSL#1	BETO with extended vocabulary	8	2E-5	5	25		

2.3. DMC component: Decision Policy

The next step was to find the best decision policy for each task using a mock server (available in: https://github.com/jmloyola/erisk_mock_server). This tool simulates the eRisk challenge through the rounds of posts and answers submissions, and it allows the calculation of the final results according to metrics based on decision and ranking. It was helpful 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 BETO architecture admits, the application uses the last $N = 10$ posts (posts window), linking the current with previous posts. With the mock server, the client application, and the predictive models, different decision policies were evaluated using the F1+, ERDE-5, ERDE-50, and latency-weighted F1 metrics. It should be noted that the client application was also used in the test stage of MentalRiskES.

Decision policy: Historic rule

“If the current prediction and last M predictions exceed T times a *Threshold*, the client application must issue a risky user alarm; otherwise, it is necessary to continue the user evaluation”.

The parameter M is the number of past predictions that the rule considers, T is the tolerance, i.e., how many predictions can exceed the *Threshold* before issuing an alarm, and *Threshold* is the limit probability to predict a user as positive. In addition, the rule has the *min_delay* parameter, which defines the moment when it will start to apply. Table 3 shows the best parameters for each task, which were found by evaluating the models with the mock server on the *Trial* corpus.

Table 3

Best parameters to define the decision policy based on the historic rule. The rule was applied for the UNSL#0 and UNSL#1 models in both tasks. M : number of past predictions; T : tolerance; *Threshold*: the limit probability to predict a positive user; *min_delay*: the moment when the rule will apply.

	M	T	Threshold	min_delay
HistoricRule_T1	all predictions	5 predictions	0.7	after 5 predictions
HistoricRule_T2	all predictions	10 predictions	0.7	after 10 predictions

In summary, the final models to solve Tasks 1 and 2 were:

Task 1: Eating disorders
UNSL#0: Classic BETO (CPI) + HistoricRule_T1 (DMC)
UNSL#1: BETO with an extended vocabulary (CPI) + HistoricRule_T1 (DMC)
Task 2: Depression
UNSL#0: Classic BETO (CPI) + HistoricRule_T2 (DMC)
UNSL#1: BETO with an extended vocabulary (CPI) + HistoricRule_T2 (DMC)

3. Results

The Organizers evaluated the teams considering metrics based on classification and latency for subtask A of Tasks 1 and 2. The first metrics evaluate the models according to classification performance, while the second ones penalize performance considering the number of posts required to detect positive users. The Organizers published a results report with team rankings ordered according to the Macro-F1 (classification-based evaluation) and ERDE-30 (latency-based evaluation) metrics.

3.1. Task 1 - Subtask A

Table 4 shows the results obtained considering the classification metrics. The models with the best Macro-F1 were CIMAT-NLP-GTO#0 with 0.966, followed by UMUTeam#0 (0.918) and UNSL#1 (0.913). Considering the mean and median values among all the proposals (in total, 25), the three models showed excellent classification performance. For its part, UNSL#0 obtained

0.751, a similar performance to the teams’ average. According to the latency metrics (Table 5), the best ERDE-30 was obtained by CIMAT-NLP-GTO#0 with 0.018, followed by UNSL#1 (0.045) and CIMAT-NLP-GTO#1 (0.065). The best latency-weighted F1 results were achieved by CIMAT-NLP-GTO#0 (0.863), BaseLine-RobertaLarge#1 (0.792), and UNSL#1 (0.776), while the best ERDE-5 was obtained by BaseLine-RobertaLarge#1 with 0.163. The UNSL#0 model achieved a better ERDE-30 than the mean and median of the proposals. In summary, the most outstanding models for Task 1 were CIMAT-NLP-GTO#0 and UNSL#1, achieving notable performance in classification and latency.

Table 4

Classification-based evaluation for Task 1 (subtask A). The best team according to the Accuracy, Macro-P, Macro-R, and Macro-F1 metrics is shown (values in bold), as well as the *mean* and *median* values of the results report for MentalRiskES 2023. The second and third-best teams are also included.

Ranking	Model#Run	Accuracy	Macro P	Macro R	Macro F1
1	CIMAT-NLP-GTO#0	0.967	0.964	0.969	0.966
2	UMUTeam#0	0.920	0.922	0.914	0.918
3	UNSL#1	0.913	0.912	0.920	0.913
15	UNSL#0	0.753	0.817	0.785	0.751
	<i>Mean</i>	0.765	0.820	0.786	0.750
	<i>Median</i>	0.810	0.830	0.829	0.810

Table 5

Latency-based evaluation for Task 1 (subtask A). The best teams according to the ERDE-5, ERDE-30, and latency-weighted F1 metrics are shown (values in bold), as well as the *mean* and *median* values of the results report for MentalRiskES 2023. The second and third-best teams are also included.

Ranking	Model#Run	ERDE-5	ERDE-30	latencyTP	speed	latency-weighted F1
1	CIMAT-NLP-GTO#0	0.334	0.018	6	0.898	0.863
2	UNSL#1	0.433	0.045	8	0.857	0.776
3	CIMAT-NLP-GTO#1	0.379	0.065	6	0.898	0.761
10	BaseLine-RobertaLarge#1	0.163	0.099	2	0.979	0.792
11	UNSL#0	0.502	0.105	8	0.867	0.673
	<i>Mean</i>	0.322	0.122	6	0.909	0.707
	<i>Median</i>	0.306	0.112	4	0.938	0.704

3.2. Task 2 - Subtask A

Table 6 shows the classification performances. Considering the Macro-F1 metric, the UMUTeam#0, UNSL#1, and UNSL#0 models obtained the best results, which were very similar to each other. However, UNSL#1 achieved the best Macro-P and Macro-R. These models notably outperformed the mean and median values among all the proposals (in total, 33). According to the latency metrics (Table 7), the SINAI-SELA#0 model achieved the best ERDE-30 (0.140), followed by UNSL#1 (0.148) and BaseLine-Deberta#0 (0.153). Regarding the latency-weighted F1 metric, SINAI-SELA#0 also obtained the best performance, while the best ERDE-5 was obtained by VICOM-nlp#2 (0.275). On the other hand, the UNSL#0 model achieved a higher performance

than the mean and median among all the teams for the ERDE-30 metric. In summary, the best models for this task were SINAI-SELA#0 and UNSL#1. Our model achieved the best classification results and remarkable latency performance with the second-best ERDE-30.

Table 6

Classification-based evaluation results for Task 2 (subtask A). The best teams according to the Accuracy, Macro-P, Macro-R, and Macro-F1 metrics are shown (values in bold), as well as the *mean* and *median* values of the results report for MentalRiskES 2023. The SINAI-SELA#0 model is also included due to its results on latency-based metrics.

Ranking	Model#Run	Accuracy	Macro-P	Macro-R	Macro-F1
1	UMUTeam#0	0.738	0.756	0.749	0.737
2	UNSL#1	0.738	0.791	0.756	0.733
3	UNSL#0	0.732	0.752	0.742	0.731
5	SINAI-SELA#0	0.725	0.775	0.742	0.720
	<i>Mean</i>	0.617	0.710	0.637	0.579
	<i>Median</i>	0.631	0.731	0.658	0.616

Table 7

Latency-based evaluation results for Task 2 (subtask A). The best teams according to the ERDE-5, ERDE-30, and latency-weighted F1 are shown (values in bold). The *mean* and *median* values of the results report for MentalRiskES 2023 are shown. The second and third-best teams are also included.

Ranking	Model#Run	ERDE-5	ERDE-30	latencyTP	speed	latency-weighted F1
1	SINAI-SELA#0	0.395	0.140	4	0.951	0.720
2	UNSL#1	0.567	0.148	14	0.791	0.609
3	BaseLine-Deberta#0	0.303	0.153	2	0.984	0.719
8	VICOM-nlp#2	0.275	0.173	2	0.984	0.706
14	UNSL#0	0.551	0.188	14	0.791	0.591
	<i>Mean</i>	0.383	0.232	8	0.902	0.599
	<i>Median</i>	0.362	0.205	3	0.967	0.627

Finally, the performance of our proposals in terms of efficiency metrics for Tasks 1 and 2 is shown in Table 8. It is observed that UNSL#1 and UNSL#0 outperformed the mean among all the proposals, demonstrating the capability to solve both tasks while minimizing resource requirements and reducing environmental impact.


Table 8

Efficiency metrics for Tasks 1 and 2. The *mean* among all proposals for each metric is shown.


Task	Model#Run	Duration (min)	Emissions	CPU energy	GPU energy	RAM energy	Consumed energy	CPU count	GPU count	RAM total size	CPU model	GPU model
1	UNSL#1	4.644	2.78E-05	6.13E-05	0	1.34E-06	6.26E-05	16	1	23.545	AMD Ryzen 7	1 x
	UNSL#0	4.631	2.77E-05	6.11E-05	0	1.34E-06	6.24E-05	16	1	23.545	1700X Eight-Core Processor	GeForce GTX 1080 Ti
	<i>Mean</i>	50.840	38.45E-05	33.66E-05	48.9E-05	4.66E-06	83.02E-05	38	4	164.198	-	-
2	UNSL#1	3.349	2.01E-05	4.42E-05	0	9.98E-07	4.52E-05	16	1	23.545	AMD Ryzen 7	1 x
	UNSL#0	3.347	2.01E-05	4.42E-05	0	9.98E-07	4.51E-05	16	1	23.545	1700X Eight-Core Processor	GeForce GTX 1080 Ti
	<i>Mean</i>	34.704	50.37E-05	38.78E-05	129.97E-05	1.17E-05	169.92E-05	29	3	123.458	-	-

3.3. Error analysis

Analyzing the proposals of our team, UNSL#1 obtained better performance than UNSL#0 in both tasks. As an illustrative example, Figure 1 shows user evaluation for Task 2, where UNSL#1 correctly resolved the misclassified users by UNSL#0. Furthermore, it observes that UNSL#1 tends to minimize the probabilities of the negative user (Figure 1a) and maximize those of the positive user (Figure 1b). It also shows that UNSL#1 detected the positive user in post 23 (decision delay=23), a reasonable instance considering the number of user posts.



(a) Evaluation of a negative user. UNSL#0 generated a false positive, while UNSL#1 classified the user correctly.



(b) Evaluation of a positive user. UNSL#0 generated a false negative, while UNSL#1 classified the user correctly.

Figure 1: Comparison between UNSL#0 and UNSL#1 to evaluate users of Task 2. The graphics show the evaluation of each model on a negative (a) and positive (b) user. The rounds of posts (*number of publication*) are observed on the *x*-axis, and the model probability at each prediction instance is on the *y*-axis. The dashed lines show the moment of the model's final decision: green (correct prediction) and red (incorrect prediction).

Considering the latency-based metrics, our proposals demonstrated satisfactory results, particularly in the ERDE-30 and latency-weighted F1 metrics. However, the results for ERDE-5 were less favorable. It would be interesting to explore potential strategies to enhance the performance of the models for ERDE-5 without compromising the other metrics. This could

involve optimizing the classification performance of the models and aligning them with the decision policy proposed in this work. Additionally, it would be worth analyzing new decision policies prioritizing speed and efficiency.

Finally, considering the mean values among all the teams, it is observed that Task 2 was more challenging than Task 1. This fact was probably due to the subjectivity level with which users expressed themselves in each domain, which may have impacted the performance of the models. For example, the post “*Me gustaría poder comer sin sentir culpa como antes*” (I wish I could eat without feeling guilty like before) could be linked to a user at risk for an eating disorder; however, the text “*Esta semana fue difícil para mí*” (This week was hard for me), it would be rushed to associate it directly with a user with depression.

4. Conclusion

In this first edition of the MentalRiskES challenge, our research group solved Tasks 1 and 2. We applied the BETO model by extending its vocabulary with important words, and we used a decision policy based on a historic rule to detect users with depression and eating disorders as early as possible. The method obtained excellent results, demonstrating its effectiveness and consistency in solving these problems in a challenging and underexplored language such as Spanish.

As future work, the classification models could be refined, analyzing the important words considered to extend the vocabulary, improving the representation of the analyzed instances during user evaluation, and testing the performance of other classification models. Furthermore, it would be interesting to evaluate other decision policies to improve the performance of the models in terms of latency.

References

- [1] F. Charlson, M. van Ommeren, A. Flaxman, J. Cornett, H. Whiteford, S. Saxena, New who prevalence estimates of mental disorders in conflict settings: a systematic review and meta-analysis, *The Lancet* 394 (2019) 240–248.
- [2] F. Aliverdi, H. Farajidana, Z. M. Tourzani, L. Salehi, M. Qorbani, F. Mohamadi, Z. Mahmoodi, Social networks and internet emotional relationships on mental health and quality of life in students: structural equation modelling, *BMC psychiatry* 22 (2022) 1–10.
- [3] P. K. Maulik, W. W. Eaton, C. P. Bradshaw, The effect of social networks and social support on common mental disorders following specific life events, *Acta Psychiatrica Scandinavica* 122 (2010) 118–128.
- [4] J. Martínez-Líbano, N. González Campusano, J. I. Pereira Castillo, et al., Las redes sociales y su influencia en la salud mental de los estudiantes universitarios: Una revisión sistemática (2022).
- [5] 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.

- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] A. M. Mármol-Romero, A. Moreno-Muñoz, F. M. Plaza-del-Arco, M. D. Molina-González, M. T. Martín-Valdivia, L. A. Ureña-López, A. Montejo-Ráez, Overview of MentalriskES at IberLEF 2023: Early Detection of Mental Disorders Risk in Spanish, *Procesamiento del Lenguaje Natural* 71 (2023).
- [13] 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.
- [14] J. M. Loyola, H. Thompson, S. Burdisso, M. Errecalde, Unsl at erisk 2022: Decision policies with history for early classification (2022).
- [15] 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.
- [16] 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).
- [17] J. Cañete, G. Chaperon, R. Fuentes, J.-H. Ho, H. Kang, J. Pérez, Spanish pre-trained bert model and evaluation data, in: PML4DC at ICLR 2020, 2020.
- [18] 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.