

Tackling a Challenging Corpus for Early Detection of Gambling Disorder: UNSL at MentalRiskES 2025

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

Gambling disorder is a complex behavioral addiction that is challenging to understand and address, with severe physical, psychological, and social consequences. Early Risk Detection (ERD) on the Web has become a key task in the scientific community for identifying early signs of mental health behaviors based on social media activity. This work presents our participation in the MentalRiskES 2025 challenge, specifically in Task 1, aimed at classifying users at high or low risk of developing a gambling-related disorder. We proposed three methods based on a *CPI+DMC approach*, addressing predictive effectiveness and decision-making speed as independent objectives. The components were implemented using the SS3, BERT with extended vocabulary, and SBERT models, followed by decision policies based on historical user analysis. Although it was a challenging corpus, two of our proposals achieved the top two positions in the official results, performing notably in decision metrics. Further analysis revealed some difficulty in distinguishing between users at high and low risk, reinforcing the need to explore strategies to improve data interpretation and quality, and to promote more transparent and reliable ERD systems for mental disorders.

Keywords

Early Risk Detection, SS3, BERT, Sentence-BERT, Decision Policy, Mental Health

1. Introduction

According to the World Health Organization, an estimated 1.2% of the adult population suffers from a gambling disorder, with the risk growing even for young people and children. In recent years, this disorder has been recognized as a behavioral addiction, prompting the Diagnostic and Statistical Manual of Mental Disorders (DSM-5) and the International Classification of Diseases (ICD-11) to reclassify it alongside substance-related disorders [1], replacing the term *pathological gambling* with *gambling disorder*. This condition encompasses a wide spectrum of physical, psychological, and social consequences [2], including high substance use [3], symptoms of anxiety, depression, stress, and impulsivity [4, 5, 6], as well as work-related and financial conflicts, relationship deterioration, and criminal behavior [7]. Technological advancements and widespread access to digital platforms have contributed to the increasing prevalence of this disorder [8], alongside other behaviors such as compulsive shopping and problematic social media use [9]. Furthermore, numerous studies have highlighted the challenges in establishing precise criteria and consistent methods for estimating the prevalence of gambling disorder due to the diversity of assessment tools, risk factors, and controversies surrounding the validity of diagnostic criteria [10, 11, 12].

In this context, Early Risk Detection (ERD) on the Web has become a significant research area in recent years, aiming to identify users who exhibit signs of developing a mental health condition as early as possible. Initiatives such as MentalRiskES have fostered research on ERD in Spanish [13, 14], while CLEF eRisk has promoted similar efforts primarily in English [15]. Our research group has actively participated in these challenges, addressing the detection of depression and eating disorders in the Spanish language [16], as well as the detection of depression [17], pathological gambling [18, 19], and anorexia [20] in English. In the MentalRiskES 2025 edition [21, 22], a challenge focused on early detection

IberLEF 2025, September 2025, Zaragoza, Spain

*Corresponding author.

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

ORCID: 0009-0009-2103-2275 (H. Thompson)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

of gambling disorder was proposed, consisting of two tasks: binary classification (Task 1), aimed at identifying users at high (positive) or low (negative) risk; and multiclass classification (Task 2), designed to distinguish the specific type of addiction associated with the disorder, such as *Betting*, *Online Gaming*, *Trading*, and *Lootboxes*. The challenge was conducted in two phases: a training stage, using a labeled corpus provided by the Organizers, and an online evaluation stage, where teams analyzed user posts progressively while interacting with a server in an early environment. Our team participated in Task 1, presenting three proposals based on a *CPI+DMC approach* [23]. This approach conceptualizes ERD as a multi-objective problem, where the goal is to optimize classification effectiveness and decision-making speed independently. It consists of two components: a *Classification with Partial Information* (CPI) model that processes user content progressively and a policy for *Deciding the Moment of Classification* (DMC) that determines when to make a final decision based on the accumulated evidence. While alternatives exist that simultaneously address both objectives [24], we opted for a modular design due to the complexity of the problem. For the CPI component, we implemented three different models: SS3, BERT with extended vocabulary, and SBERT. For the DMC component, we designed decision policies that evaluate users based on historical analysis.

According to the official ranking released by the Organizers, two of our proposals achieved first and second place, with remarkable results in the Macro F_1 score and other decision-making metrics. A detailed analysis of the results highlighted the inherent complexity of the task: the differentiation between users at high and low risk is subtle and difficult to define, posing a challenge for both the proposed models and potential human evaluators. The structure of the paper is as follows: Section 2 presents details of the corpus and a preliminary analysis of the data; Section 3 describes the methodology adopted and the models used; Section 4 discusses the results obtained; and Section 5 offers conclusions and future work.

2. Corpus

To address Task 1, the Organizers developed a corpus [25] divided into three parts, deployed in the different phases of the challenge (Table 1). The Train and Trial sets were provided for model training and server connection testing, while the Test set was reserved for the final evaluation of participating teams. Each set presents a balanced class distribution, with a similar mean number of posts per user (around 60) and a relatively short post length, averaging fewer than 10 words per post, although some posts are considerably longer. Additionally, users come from the Telegram and Twitch platforms, with a balanced distribution across classes.

User information is organized at the post level and includes structured metadata such as message ID, round number, user pseudonym, text content, date, and origin platform. For example, a post from user subject1 is represented as: {id_message: 123, round: 1, nick: "subject1", message: "...", date: "2021-01-06 04:02:48+01:00", platform: "Telegram"}. Each user was labeled by the Organizers with a binary class (high or low risk) based on their complete posting history, considering that all users show some level of gambling behavior, though their risk levels vary.

Table 1

Overview of the three subsets of the corpus used in Task 1. The number of users (Total, Positive, and Negative) is reported, along with statistics on posts per user and words per post (Mean, Minimum, and Maximum). The distribution of users across platforms (Telegram and Twitch) is also indicated.

Corpus	#Users			#Posts per user			#Words per post			Platform			
	Total	Pos	Neg	Mean	Min	Max	Mean	Min	Max	Telegram		Twitch	
										Pos	Neg	Pos	Neg
Train	350	172	178	64	8	146	7	1	467	105	115	67	63
Trial	7	3	4	63	1	135	5	1	148	2	2	1	2
Test	160	83	77	59	8	140	7	1	202	50	40	33	37

We conducted a preliminary exploration of the available training sets (357 users across the Train

and Trial sets), aiming to analyze the user textual content. First, we calculated the cosine similarity on the TF-IDF representations generated from the complete vocabulary of each class, obtaining a value of 0.854, indicating a high lexical similarity between positive and negative users. Next, we used the Jaccard index to assess the lexical overlap between the classes, considering the 1,000 most frequent words in each. The resulting value was 0.581, corresponding to 735 shared words out of 1,265 unique words, further reinforcing the significant lexical overlap between the two classes. Inspection of these shared words revealed topics such as *cryptocurrencies*, *financial markets*, *games*, *betting*, *digital platforms*, and various *emotional states*. The remaining words unique to each class were strongly linked to these topics, reflecting subtle differences. For instance, positive users tended to use more technical and advanced language, referencing specific platforms (e.g., *BingX*, *Winamax*, *Ledger*, *Discord*, *OKE*), financial market concepts (e.g., *Elliot*, *scalping*, *velas*—candles, *store*, *liquidez*—liquidity), and more active participation in forums and communities. In contrast, negative users displayed less technical language, with expressions suggesting a more cautious attitude (e.g., *aprender*—learn, *consejo*—advice, *ojalá*—hopefully, *imposible*—impossible, *paciencia*—patience), possibly reflecting less experience with these topics. As part of the study, we also found that 80% of users made most of their posts during nighttime hours (between 6 p.m. and 6 a.m.), with this tendency being slightly stronger among positive users (82%) compared to negative users (78%). It should be noted that the timestamps are in UTC+1 (+01:00), although it is unclear whether this timezone corresponds to the actual location of each user. After manually inspecting posts, personal contexts, and user dynamics, no clear pattern emerged that could differentiate high-risk from low-risk users. Therefore, this exploration suggests that the task is indeed challenging, which motivated the models proposed by our team.

3. Methodology

Given the domain complexity and the high lexical similarity between classes, we explored three distinct methods, each aimed at capturing different levels of representation and implementing alternative early classification strategies. We designed three proposals that combine different classifiers and decision-making strategies in line with the *CPI+DMC approach*. Below, we describe the models employed and the experiments conducted for their training.

3.1. Models

3.1.1. SS3 Using a Global Decision Policy

The SS3 model [26] is a supervised classifier created for ERD problems. It is a robust method that enables incremental user analysis and facilitates the interpretability of the decisions. During the training, SS3 builds a vocabulary with term frequencies for each class. It employs a *global value* (gv) function that assigns a score to each term relative to the target classes, considering three parameters: σ (smoothness), ρ (sanction), and λ (significance). The model aims to emulate human behavior by focusing on key terms when classifying a text, thereby contributing to its interpretability. Internally, it performs a hierarchical analysis at multiple levels (words, sentences, and paragraphs) and applies summary operators to obtain a global value of each. Then, the final classification depends on the sum of the gv scores of all the terms in a user's text. This model constitutes the CPI component of our first proposal, where the representation of samples is created from a frequentist and interpretable model based on term relevance without relying on deep learning.

To implement the DMC component, we adopted a *global decision policy* previously proposed by our laboratory [17]. We defined a value $score_u$ that estimates a user's overall risk level based on their post history and the target classes. During user evaluation, we maintain two *confidence values*, $cv_{positive}$ and $cv_{negative}$, which accumulate the gv of each observed term across the user's writing history. At post round (*delay*), the user's current risk level is estimated by normalizing these cumulative values through

a *softmax*-based transformation:

$$score_u = softmax \left(\left[\frac{cv_{positive}}{delay}, \frac{cv_{negative}}{delay} \right] \right)_{positive} \quad (1)$$

This normalization ensures that the resulting scores are within the range $[0,1]$ and allows a fairer comparison between users by mitigating the impact of very short or very long posting histories. In this way, $score_u$ represents the relative likelihood that the user belongs to the positive class and serves as the basis for the decision-making process that determines whether the user should be classified as positive or negative:

$$decision_u = \begin{cases} 1, & \text{if } score_u > \text{median}(scores) + \gamma \cdot \text{MAD}(scores) \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

This policy uses a dynamic threshold defined by the median of all users' scores ($scores = \{score_u | u \in \text{Users}\}$) and the median absolute deviation (MAD), which together define an uncertainty interval: $\text{median}(scores) \pm \gamma \cdot \text{MAD}(scores)$. The hyperparameter γ controls how much a user's $score_u$ must deviate from the median to be classified as positive. In other words, a user is considered at risk if their $score_u$ is significantly higher than most users' scores.

3.1.2. Extended BERT Using a History-Based Decision Policy

We used this transformer-based version as a baseline. It involves fine-tuning a pre-trained BERT model by expanding its original vocabulary, enabling it to represent domain-specific terms previously unknown to the model and whose semantics can contribute to the classification task. Specifically, we used BETO [27], a BERT model pre-trained on Spanish corpora. To identify the new vocabulary terms, we relied on the SS3 model (from our previous proposal) to rank the words according to their relevance to the positive class, from which we selected those to incorporate into the model. In this way, the CPI component of our second proposal aims to obtain distributed and contextualized text representations enriched with domain-relevant terms.

For the DMC component, we employed a decision policy based on the model prediction history during early user detection, referred to as the *history-based rule* [18, 19, 16, 20, 24]. Since transformer models have limitations on the number of tokens they can process, we used a sliding window that concatenates the current post with the previous N ones. At each step, the model predicts the current window and applies the *history-based rule*:

$$decision_u = \begin{cases} 1, & \text{if } \sum_{i=1}^t \mathbb{I}(p_i \geq \tau) \geq T \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

Where p_i is the predicted probability at round i , τ is the decision threshold, T is the number of required positive predictions, and $\mathbb{I}(\cdot)$ is the indicator function. Then, τ and T are hyperparameters that determine the sensitivity and tolerance of the policy, respectively, and were tuned based on the model behavior during the early evaluation of the users. To achieve this, we used the *mock-server* tool¹, which simulates an early detection environment through rounds of posts and response submissions, enabling the evaluation of the model's performance through various metrics.

3.1.3. SBERT Using a History-Based Decision Policy

The third variant relied on *Sentence-BERT* (SBERT) [28], a BERT-based model adapted with a Siamese architecture to generate dense, sentence-level representations, capturing semantic relationships to solve tasks, such as classification and semantic search. For the CPI component, we used SetFit (Sentence Transformer Fine-tuning) [29], an efficient framework designed for few-shot scenarios based on

¹Available at: https://github.com/jmloyola/erisk_mock_server

contrastive learning. The SBERT encoder is fine-tuned by automatically generating pairs of examples (positive and negative) from the original dataset and training the model to produce embeddings that are closer for examples from the same class and distant from those of different classes. This process results in a discriminative semantic space that supports class separation according to sentence-level representations obtained by the fine-tuned encoder. Then, an independent classifier is trained on the resulting embeddings without further modifying the encoder, leading to an efficient and effective method for classification tasks in complex domains. For DMC, we applied the same *history-based rule* (Equation 3), considering the model prediction history to decide when to trigger a risk alert.

3.2. Experiments

Following the *CPI+DMC approach*, the experimentation was organized in two stages. In the first stage, we explored different configurations and hyperparameters to find optimal models for user classification. In the second stage, we evaluated different decision-making policies by testing the previously selected models in an early detection environment using the *mock-server* tool. The data usage and specific model configurations adopted by our team are detailed below.

3.2.1. Data and preprocessing

We used the textual content of user posts and discarded metadata such as date and platform. Although we considered this information, preliminary results showed no performance improvements, and no substantial patterns were found to justify its use. We merged the Train and Trial sets for the experiments, resulting in 357 samples: 257 for model training and validation and 100 for evaluation in an early detection environment while maintaining a balanced distribution between classes. Preprocessing included converting texts to lowercase, transforming Unicode and HTML sequences into corresponding symbols, normalizing URLs using the 'weblink' token, removing repeated words, and applying other basic text-cleaning operations.

3.2.2. Model setup

UNSL#0. SS3 model trained on character trigrams using the hyperparameters $\sigma=0.44$, $\rho=0.5$, and $\lambda=0.86$, selected via grid-search optimized by the F_1 metric. A global decision policy was applied, configured with $\gamma=0.5$.

UNSL#1. We used the BETO model (checkpoint: *dccuchile/bert-base-spanish-wwm-uncased*) and included 25 domain-relevant words extracted from the SS3 model, considering confidence values assigned to the positive class. This extension allowed us to include terms originally outside the BETO vocabulary, such as *rebote* (rebound), *combi* (combo bets or parlays), *divergencia* (divergence), *BingX*, *scalping*, and *velita* (candlestick), among others. The remaining hyperparameters were: optimizer = AdamW, learning_rate = 5E-5, scheduler = LinearSchedulerWarmup, batch_size = 32, and n_epochs = 10. The checkpoint with the highest F_1 score on the validation set was selected. We used a *history-based rule* configured with $T=10$ and $\tau=0.6$.

UNSL#2. SBERT model, based on BETO and pre-trained on semantic similarity tasks in Spanish (checkpoint: *hiiamsid/sentence_similarity_spanish_es*). We fine-tuned the encoder using the *CosineSimilarityLoss* function, followed by a logistic regression classifier trained on the resulting embeddings. The configuration included batch_size = 16, num_epochs = 1, num_iteration = 20, and learning_rate = 2E-5. Finally, we defined a *history-based rule* configured with $T=10$ and $\tau=0.7$.

4. Results

A total of 38 proposals were submitted by thirteen teams for Task 1. The Organizers evaluated the models using classification and latency metrics, and published an official ranking based on the Macro

F_1 . Table 2 summarizes the results obtained by our models and compares them with some of the most relevant proposals (complete official results reported in [21]). We highlight the following observations:

- UNSL#2 reached first place in the overall ranking, achieving a Macro F_1 of 0.567. It also obtained the second-best results in Accuracy, Macro Recall, Micro Precision, Micro Recall, and Micro F_1 . Additionally, it showed acceptable $F_{latency}$, considering the top three models.
- UNSL#0 was ranked second, with a Macro F_1 of 0.563, and delivered comparable performance to UNSL#2. It excelled in several metrics, obtaining the best scores in Accuracy, Macro Recall, Micro Precision, Micro Recall, and Micro F_1 , as well as the best $F_{latency}$ among the top-ranked models. It also achieved an $ERDE_{30}$ of 0.284, better than the overall average (0.325) and comparable with the best, PLN-PPM-ISB#0.
- UNSL#1 model achieved 16th with a Macro F_1 of 0.444, outperforming both the mean (0.426) and the median (0.429) of all submissions, and providing an acceptable baseline for comparison.
- Among other proposals, I2C-UHU-Rigel#1 achieved the third-best Macro F_1 , VerbaNexAI-Lab#0 (27th) the best $ERDE_5$, PLN-PPM-ISB#0 (17th) the best $ERDE_{30}$ and the second-best $F_{latency}$, while Robertuito (20th) obtained the best Macro Precision and $F_{latency}$.

Table 2

Decision-based evaluation results for Task 1 in MentalRiskES 2025. The best three models, ranked by the Organizers based on the Macro F_1 metric, are shown. Additionally, the highest performances in $ERDE_5$, $ERDE_{30}$, and $F_{latency}$ are included, along with the Mean and Median across all proposals. Values in bold and underlined indicate the first and second-best scores, respectively.

Rank	Model#Run	Acc	Macro P	Macro R	Macro F_1	Micro P	Micro R	Micro F_1	$ERDE_5 \downarrow$	$ERDE_{30} \downarrow$	$F_{latency}$
1	UNSL#2	<u>0.569</u>	0.568	<u>0.567</u>	0.567	<u>0.569</u>	<u>0.569</u>	<u>0.569</u>	0.639	0.389	0.506
2	UNSL#0	0.581	0.586	0.574	<u>0.563</u>	0.581	0.581	0.581	0.515	0.284	0.628
3	I2C-UHU-Rigel#1	0.556	0.555	0.553	<u>0.551</u>	0.556	0.556	0.556	0.600	0.432	0.496
16	UNSL#1	0.475	0.459	0.467	0.444	0.475	0.475	0.475	0.707	0.476	0.467
17	PLN-PPM-ISB#0	0.550	0.632	0.534	0.436	0.550	0.550	0.550	0.316	0.242	0.683
20	Robertuito	0.550	0.657	0.533	0.428	0.550	0.550	0.550	0.329	0.252	0.685
27	VerbaNexAI-Lab#0	0.519	0.259	0.500	0.342	0.519	0.519	0.519	0.274	0.250	0.677
Mean		0.524	0.457	0.512	0.426	0.524	0.524	0.524	0.473	0.325	0.587
Median		0.519	0.490	0.500	0.429	0.519	0.519	0.519	0.445	0.283	0.657

In addition, the Organizers requested that the teams provide estimates of energy consumption and resource usage to assess the computational and environmental impact using the CodeCarbon library². As shown in Table 3, our models were executed on the same hardware configuration, with each inference taking an average of 2.3 seconds, consuming approximately 7E-05 kWh (kilowatt-hour) of energy, and producing 1.66E-07 kgCO₂eq (kilograms of carbon dioxide equivalent), all values significantly lower than the recorded mean.

Table 3

Efficiency and resource usage metrics for our team’s submissions in Task 1. Mean values across all proposals are reported when applicable.

Team#Run	Duration (sec)	Emissions	CPU energy	GPU energy	RAM energy	Consumed energy	CPU count	GPU count	Total RAM size	CPU model	GPU model
UNSL#0	2.27E+00	1.66E-07	2.68E-05	4.29E-05	2.88E-07	6.99E-05	32	1	63.730	Intel(R) Core(TM) i9-14900HX	1 x NVIDIA GeForce RTX 4090 Laptop GPU
UNSL#1	2.27E+00	1.66E-07	2.68E-05	4.29E-05	2.88E-07	6.99E-05	32	1	63.730		
UNSL#2	2.27E+00	1.66E-07	2.68E-05	4.29E-05	2.88E-07	6.99E-05	32	1	63.730		
Mean	3.03E+01	6.59E-04	6.01E-04	5.98E-04	9.45E-05	1.30E-03	16	1	40.27	-	-

4.1. Error analysis

Figure 1 presents the confusion matrix for each proposal, based on the first time each model predicted a user as positive during its post-by-post interaction with the server. Among the three models, UNSL#0

²Available at: <https://github.com/mlco2/codecarbon>

detected the most true positives (TPs), while UNSL#2 identified the most true negatives (TNs). UNSL#2 also had the lowest number of false positives (FPs) and a similar number of false negatives (FNs), achieving a balance between precision and recall, suggesting a more conservative approach to detecting positive cases. In contrast, UNSL#0 reduced FN but increased FP, reflecting a more sensitive yet less precise strategy that prioritizes early detection, even at the expense of generating more incorrect alerts. Meanwhile, UNSL#1 exhibited intermediate performance in FNs, the highest number of FPs, and the lowest TNs, indicating difficulties distinguishing between the two classes.

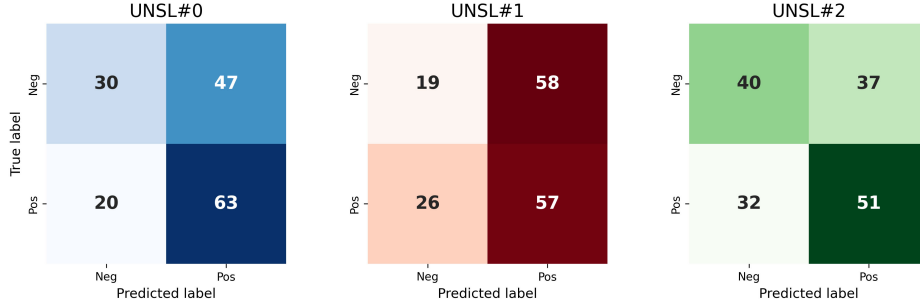


Figure 1: Confusion matrices for the three models proposed in Task 1.

Figure 2 shows the predictions of the three models when analyzing a positive user from Task 1. Despite following different strategies, all three models consistently showed signs that the user was at high risk throughout the analysis. UNSL#0 predicted the user as positive at round 7, UNSL#1 at round 39, and UNSL#2 at round 29. UNSL#0 exhibits fewer score variations, with confidence values remaining close to 0.5, and issued an alert after exceeding the uncertainty interval defined by the decision policy. UNSL#1 displays higher variability, with predictions oscillating between high and low probabilities across many rounds. In contrast, UNSL#2 initially presents isolated high scores that later stabilize toward the end of the analysis. Both UNSL#1 and UNSL#2 benefited from the history-based rule, which allowed them to tolerate fluctuations and wait for consistent signals before issuing a final decision.

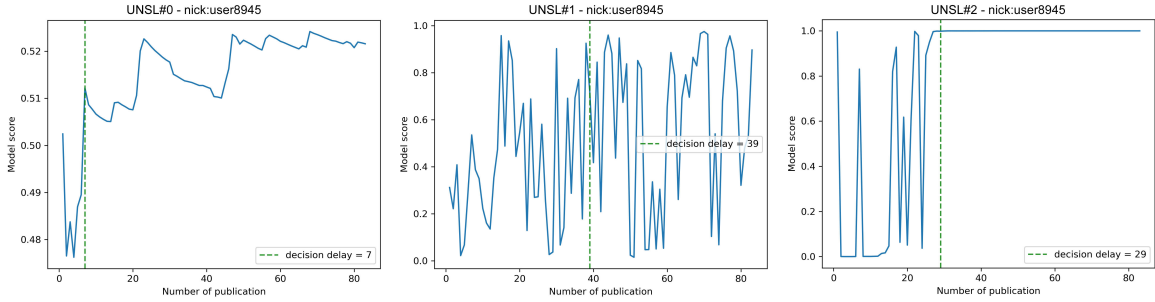


Figure 2: Comparison of UNSL#0, UNSL#1, and UNSL#2 on a positive user from Task 1. The x -axis shows the rounds of posts (number of publication), while the y -axis indicates the model's predicted score at each decision point. Dashed green lines mark the moment when each model made its final correct prediction.

The Venn diagram in Figure 3 illustrates the distribution of positive predictions across the three models, highlighting their overlaps and divergences. The three models agreed on 57 positive instances (35 correctly classified), while the remaining 22 were FPs shared by all three. Upon examining these presumably incorrect samples, we observed recurring themes such as *sports betting*, *video games with elements of chance*, and *cryptocurrency trading*. These cases often featured behaviors such as *financial speculation*, *active engagement in games*, and *intense emotional expressions related to wins and losses*. Only a minority consisted of short or ambiguous messages that required deeper contextual or linguistic interpretation. This suggests that many of these FPs may be due to the limitations of the corpus in clearly distinguishing between risk levels. The analysis highlights the intrinsic complexity of the task, where ambiguity and overlap between users at high and low risk, not only at the lexical level (as noted

in Section 2) but also semantic, can lead to misclassifications by both predictive models and human evaluators.

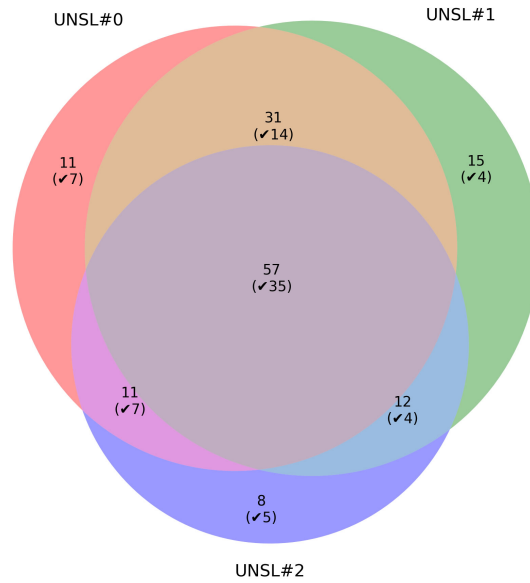


Figure 3: Venn diagram comparing the users predicted as positive by the three models proposed in Task 1. The overlaps and values between sets illustrate the agreement among models, with the number of correct predictions shown in parentheses.

To assess the semantic consistency of the predictions made by the UNSL#2 and UNSL#0 models, we show some illustrative examples using the following sentences:

S1: *“hoy jugué durante horas en BingX buscando ese rebote que me haría recuperar lo que perdí ayer... pero me comí terrible divergencia!”* (I played for hours today on BingX looking for that rebound that would make up for what I lost yesterday... but I got caught with a terrible divergence!).

S2: *“tarde o temprano llegará el goal, pero espero que no sea en el primer tiempo porque entre con bastante”* (Sooner or later the goal will come, but I hope it’s not in the first half because I’m going in with a lot).

S3: *“he hecho algunas pequeñas inversiones, pero la verdad que no me gustan mucho estas cosas”* (I’ve made some small investments, but I’m not really into this kind of thing).

The UNSL#2 model classified the first two sentences as positive and the third as negative. To assess the semantic consistency of these predictions, we obtained the embeddings of each sentence using the UNSL#2 encoder and calculated the cosine similarity between them. We observed that S1 and S2 exhibited high similarity (0.7282), while S3 showed low similarity with both S1 (0.1202) and S2 (0.0065). This suggests that the model may be constructing a representation space where sentences related to users with gambling tendencies (e.g., impulsive investments or sports betting) tend to be closer together, while those without these characteristics are farther apart. Since distinguishing between users at high and low risk can be ambiguous and challenging, even semantically consistent representations can result in errors if the corpus does not adequately reflect these differences. However, the overall performance of UNSL#2 in detecting the positive class depends on both the learned representation space and the classifier that utilizes these representations for prediction.

The UNSL#0 model, based on SS3, facilitates interpretation by providing a tool to visualize the information it considers relevant for each prediction. For instance, Figure 4 shows that the model classified sentence S1 as positive, assigning distinct relevance scores to terms such as *durante* (during),

BingX, *rebote* (bounce), and *divergencia* (divergence) for the positive class, while the term *jugué* (I played) was slightly associated with the negative class. Additionally, the cumulative *cv* increases as the sentence progresses, especially after encountering the word *BingX*, which strongly contributed to the model’s decision.

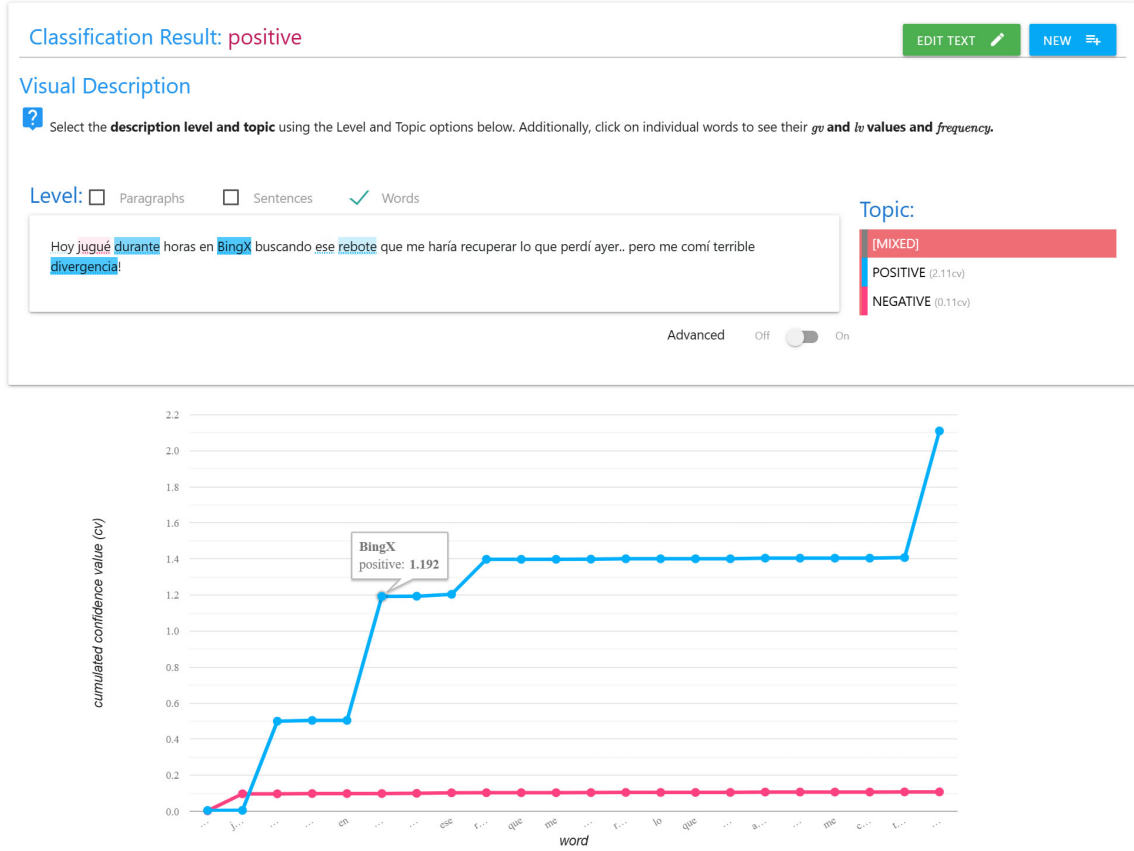


Figure 4: Visual explanation of the evaluation process used by UNSL#0. The model has a visualization tool to improve interpretability by using the cumulated confidence values (*cv*), associated with each word in the sentence. Words are color-coded according to their contribution to the positive class. As the analysis progresses, the *cv* value increases, with certain words such *BingX*, *rebote* (bounce), and *divergencia* (divergence) progressively reinforcing the model’s tendency to classify the sentence as positive.

4.2. Final Considerations

The results highlight the inherent complexity of the task, particularly the challenge of distinguishing between positive and negative classes. In [30], we explored Large Language Models (LLMs) to address the ERD in depression by introducing a reasoning criterion grounded in expert knowledge. The goal was not only to detect positive cases, but also to generate explanations that justify the model’s decisions. While defining such criteria is difficult, this strategy can enhance data construction and interpretation, allowing for more precise identification of the specific moments when risk signals emerge. It may also contribute to more reliable evaluation metrics, such as an adaptive version of $ERDE_{\theta}$, where each user has an individual threshold θ based on their behavior, penalizing delayed decisions more fairly.

On the other hand, the ERDE metric was defined in [31] to penalize delayed TPs, while the cost associated with FPs and FNs depends on the domain. The authors also point out that negative cases correspond to non-risk situations, in which early or urgent intervention is not required. However, in our case, negative users are already at some level of risk, which changes the notion of “early detection”: it is no longer about identifying an early risk case under the assumption that risk is initially absent, but rather about recognizing the moment when a user crosses a critical threshold that justifies more

serious concern. This shift in perspective may help explain why most participants underperformed on early detection metrics, particularly $ERDE_{\theta}$. Considering that in Task 1 the average number of posts per user was around 60 (see Table 1), a higher threshold (e.g., $\theta = 50$) could have helped avoid severe penalties during the initial stages of analysis, when much of the evidence was still unavailable.

This year's edition raises relevant conceptual challenges for ERD, including how risk is defined, how it is measured, and what types of decisions we expect models to make. For instance, an FP made with very little evidence might represent a more serious false alarm than a late FP, in which the user already shows ambiguous patterns, something even possibly acceptable from a preventive perspective. The same holds for FNs: delaying a prediction may be reasonable when signals are weak, but as more information accumulates over time, the model should be able to detect evident signs of risk.

5. Conclusion

Our laboratory solved Task 1 of MentalRiskES 2025 by presenting three proposals based on the *CPI+DMC approach*. Two of these proposals achieved outstanding results among all participating teams, demonstrating that ERD can be addressed by balancing classification performance and decision-making speed through a modular and independent approach. Corpus exploration played a crucial role in the selection of the methods used. The results highlighted the complexity of distinguishing between high-risk and low-risk users, which can be challenging even from a human perspective. It is essential to continue researching strategies that enhance the quality and interpretation of data, particularly in the ERD of mental health, where transparent and reliable systems are needed to support identification and analysis in critical areas, such as gambling disorder. Furthermore, we will continue exploring new approaches that address ERD by combining predictive effectiveness and speed as a single combined objective.

Acknowledgments

This work is part of the doctoral research of Horacio Thompson, carried out at the Laboratorio de Investigación y Desarrollo en Inteligencia Computacional (LIDIC), under the project PROICO 03-0620, Argentina.

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

- [1] H. S. Kim, D. C. Hodgins, A Review of the Evidence for Considering Gambling Disorder (and Other Behavioral Addictions) as a Disorder Due to Addictive Behaviors in the ICD-11: a Focus on Case-Control Studies, *Current Addiction Reports* 6 (2019) 273–295.
- [2] A. Wöhr, M. Wuketich, Perception of Gamblers: A Systematic Review, *Journal of Gambling Studies* 37 (2021) 795–816.
- [3] M. Browne, V. Rawat, P. Newall, S. Begg, M. Rockloff, N. Hing, A Framework for Indirect Elicitation of the Public Health Impact of Gambling Problems, *BMC Public Health* 20 (2020) 1–14.
- [4] A. H. Bargerón, J. M. Hormes, Psychosocial Correlates of Internet Gaming Disorder: Psychopathology, Life Satisfaction, and Impulsivity, *Computers in Human Behavior* 68 (2017) 388–394.
- [5] C. De Pasquale, F. Sciacca, V. Martinelli, M. Chiappedi, C. Dinaro, Z. Hichy, Relationship of Internet Gaming Disorder with Psychopathology and Social Adaptation in Italian Young Adults, *International journal of environmental research and public health* 17 (2020) 8201.

- [6] A. M. Wu, J. H. Chen, K.-K. Tong, S. Yu, J. T. Lau, Prevalence and Associated Factors of Internet Gaming Disorder Among Community Dwelling Adults in Macao, China, *Journal of behavioral addictions* 7 (2018) 62–69.
- [7] M. Browne, E. Langham, V. Rawat, N. Greer, E. Li, J. Rose, M. Rockloff, P. Donaldson, H. Thorne, B. Goodwin, et al., Assessing Gambling-Related Harm in Victoria: A Public Health Perspective, Technical Report, Victorian Responsible Gambling Foundation, 2016.
- [8] S. Núñez-Rodríguez, D. Burgos-González, L. A. Mínguez-Mínguez, F. Menéndez-Vega, J. L. Antoñanzas-Laborda, J. J. González-Bernal, J. González-Santos, Effectiveness of Therapeutic Interventions in the Treatment of Internet Gaming Disorder: A Systematic Review, *European Journal of Investigation in Health, Psychology and Education* 15 (2025) 49.
- [9] G. Mestre-Bach, U. Paiva, L. S. M. Iniguez, M. Beranuy, M. Martín-Vivar, N. Mallorquí-Bagué, E. Normand, M. C. Chicote, M. N. Potenza, G. Arrondo, The Association Between Internet-Use-Disorder Symptoms and Loneliness: A Systematic Review and Meta-Analysis with a Categorical Approach, *Psychological Medicine* 55 (2025) e77.
- [10] E. Gabellini, F. Lucchini, M. E. Gattoni, Prevalence of Problem Gambling: A Meta-analysis of Recent Empirical Research (2016-2022), *Journal of Gambling Studies* 39 (2023) 1027–1057.
- [11] Y. Allami, D. C. Hodgins, M. Young, N. Brunelle, S. Currie, M. Dufour, M.-C. Flores-Pajot, L. Nadeau, A Meta-Analysis of Problem Gambling Risk Factors in the General Adult Population, *Addiction* 116 (2021) 2968–2977.
- [12] C. J. Rash, J. Weinstock, R. Van Patten, A Review of Gambling Disorder and Substance Use Disorders, *Substance abuse and rehabilitation* (2016) 3–13.
- [13] 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) 329–350.
- [14] 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 2024: Early Detection of Mental Disorders Risk in Spanish, *Procesamiento del lenguaje natural* 73 (2024) 435–448.
- [15] J. Parapar, P. Martín-Rodilla, D. E. Losada, F. Crestani, Overview of eRisk 2024: Early Risk Prediction on the Internet, in: *Experimental IR Meets Multilinguality, Multimodality, and Interaction. 15th International Conference of the CLEF Association, CLEF 2024, Grenoble, France, Springer International, 2024.*
- [16] H. Thompson, M. Errecalde, Early Detection of Depression and Eating Disorders in Spanish: UNSL at MentalRiskES 2023, 2023. [arXiv:2310.20003](https://arxiv.org/abs/2310.20003).
- [17] 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.
- [18] J. M. Loyola, H. Thompson, S. Burdisso, M. Errecalde, UNSL at eRisk 2022: Decision Policies with History for Early Classification (2022).
- [19] H. Thompson, L. Cagnina, M. Errecalde, Strategies to Harness the Transformers’ Potential: UNSL at eRisk 2023, in: *CLEF (Working Notes), 2023*, pp. 791–804.
- [20] H. Thompson, M. Errecalde, A Time-Aware Approach to Early Detection of Anorexia: UNSL at eRisk 2024, 2024. [arXiv:2410.17963](https://arxiv.org/abs/2410.17963).
- [21] A. M. Mármol-Romero, P. Álvarez-Ojeda, A. Moreno-Muñoz, F. M. P. del Arco, M. D. Molina-González, M.-T. Martín-Valdivia, L. A. Ureña-López, A. Montejo-Ráez, Overview of MentalRiskES at IberLEF 2025: Early Detection of Mental Disorders Risk in Spanish, *Procesamiento del Lenguaje Natural* 75 (2025).
- [22] J. Á. González-Barba, L. Chiruzzo, S. M. Jiménez-Zafra, Overview of IberLEF 2025: Natural Language Processing Challenges for Spanish and other Iberian Languages, in: *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2025), co-located with the 41st Conference of the Spanish Society for Natural Language Processing (SEPLN 2025), CEUR-WS. org, 2025.*

- [23] 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.
- [24] H. Thompson, E. Villatoro-Tello, M. Montes-y Gómez, M. Errecalde, Temporal Fine-tuning for Early Risk Detection, in: Memorias de las JAIIO–Simposio Argentino de Inteligencia Artificial y Ciencias de Datos (ASAIID), volume 10, 2024, pp. 137–149.
- [25] P. Álvarez-Ojeda, M. V. Cantero-Romero, A. Semikozova, A. Montejo-Ráez, The PRECOM-SM Corpus: Gambling in Spanish Social Media, in: Proceedings of the 31st International Conference on Computational Linguistics, 2025, pp. 17–28.
- [26] 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.
- [27] 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.
- [28] N. Reimers, I. Gurevych, Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks, 2019. [arXiv:1908.10084](https://arxiv.org/abs/1908.10084).
- [29] L. Tunstall, N. Reimers, U. E. S. Jo, L. Bates, D. Korat, M. Wasserblat, O. Pereg, Efficient Few-Shot Learning Without Prompts, 2022. [arXiv:2209.11055](https://arxiv.org/abs/2209.11055).
- [30] H. Thompson, M. Sapino, E. Ferretti, M. Errecalde, Hacia la Interpretabilidad de la Detección Anticipada de Riesgos de Depresión Utilizando Grandes Modelos de Lenguaje, 2025. [arXiv:2503.20939](https://arxiv.org/abs/2503.20939).
- [31] D. E. Losada, F. Crestani, A Test Collection for Research on Depression and Language Use, in: International conference of the cross-language evaluation forum for European languages, Springer, 2016, pp. 28–39.