# Prediction of suicidal behaviors in hospitalized children and adolescents in middle-income countries: a case study of Brazil

**Isis F. Carvalho[1], Debora Miranda[2], Ana Paula Couto da Silva[1], Anisio M. Lacerda[1], Wagner Meira Jr.[1], Marco A. Romano-Silva[2], Maria Carolina Lobato[2], Gisele L. Pappa[1]**

[1]Department of Computer Science – Universidade Federal de Minas Gerais (UFMG)
Belo Horizonte – MG – Brazil

[2]Faculty of Medicine –Universidade Federal de Minas Gerais (UFMG)
Belo Horizonte – MG – Brazil

{isisferreira,ana.coutosilva,anisio,meira,glpappa}@dcc.ufmg.br

{debora.m.miranda,romanosilva,mcarollobato}@gmail.com

***Abstract.*** *Suicide is the first leading cause of death among children and adolescents worldwide. Predictors of suicide-related behaviours might help in the task of intervening to avoid or monitor future suicide risks. In this paper, a sample of individuals who were taken to a Child Psychiatry Facility in Brazil was analyzed. Machine learning algorithms were used to generate models for predicting suicidal behaviour, and the features that better explain this complex behaviour were also analyzed. Results show a sensitivity of 0.83 and a specificity of 0.97.*

## 1. Introduction

One person dies by suicide every 40 seconds in the world, and for each one, 60–135 people are impacted by his/her death [Knipe et al. 2022]. Among children and adolescents, suicide is the first leading cause of death in developed countries [Knipe et al. 2022]. Both suicide and related behaviours, such as self-harm, have an enormous and underestimated burden in low- and middle-income countries [Knipe et al. 2022, Naghavi and Collaborators 2019]. Self-harm occurs at a frequency 20 times higher than suicide and tends to precede suicide attempts. Because of that, in high-income countries, self-harm is the clearest predictor of suicide [Carroll et al. 2014].

Given the impacts of the death of children and teenagers in the familiar and societal structure, investigating the factors that lead to this condition is especially relevant. Predicting suicide is a difficult task even for competent psychiatrists and mental health specialists [Harris et al. 2019], as it goes beyond health-related issues. There is a need to understand how social and environmental factors interact with mental health. This is especially relevant in countries such as Brazil, since inequality and social failures may amplify psychiatric risks.

Obtaining data to understand the factors involved in suicide risk is a challenging task, and most data rely on self-reports. In this work, we analyze and generate simple models to predict suicidal ideation and suicide attempt based on a dataset collected from the emergency unit of the Child and Adolescent Psychiatry Center (CEPAI), associated with the hospital network FHEMIG, in Belo Horizonte, Brazil.

Dealing with this task is difficult because (i) the data has an intrinsic bias – all the people that have arrived at the emergency already had psychological problems; (ii) the suicide attempt itself is a rare event, generating very unbalanced data; (iii) the dataset is limited to a low number of patients; (iv) the risk of wrongly predicting someone with a high-risk of the suicide attempt has more serious consequences than mistaking a low-risk person, given that it is a cost-sensitive classification task.

Identifying good predictors of suicide-related behaviours can help identify ways to intervene in the behaviour of a person or monitor certain variables to reduce the risk of these behaviours. Our final goal is to be able to change variables capable of modifying risk with the help of good suicide risk predictors.

This present work expands the current literature by applying traditional machine learning techniques on very rich data derived from an emergency mental health care institution for children and adolescents. The emergency context is a big differential in our work, given that the patients CEPAI takes care of are often in a moment of crisis, and dealing with all the risks involved is an extremely complex task. That is also why the identification of predictors of suicide-related behaviours are a essential part of this work. Results show a sensitivity value over 0.83 and a specificity of 0.97 for suicide attempt prediction.

## 2. Related Work

As previously stated, predicting suicide is a difficult task even for competent psychiatrists and mental health specialists. This fact is evidenced by [Harris et al. 2019], which makes a thorough systematic review of risk assessment tools for future suicide prediction specifically in the context of adolescence. They compare ten risk assessment tools used in the US and the UK and concluded that none of them were able, on their own, to predict suicidal behaviour. This illustrates that suicidal behaviour prediction is a complex problem and one that demands a lot of attention since suicide has been growing on a global scale.

Another work that deals with very rich data, in order to predict suicide behaviour, is [Navarro et al. 2021]: it uses traditional machine learning (ML) models to analyze the impact of early life factors on suicide attempt prediction using data from the Quebec Longitudinal Study of Child Development (QLSCD) [Orri et al. 2021]. QLSCD has greatly contributed to understanding how early childhood factors impact the development of general behaviour and mental health problems.

Given that the problem we want to tackle is a challenging one, the use of machine learning models is a very interesting way to start. [Ji et al. 2021] review many of the existing approaches to tackle suicide prediction, focusing on methods that use deep learning and natural language processing (NLP). This work also highlights the general complications related to the problem of suicide prediction and other healthcare-related tasks. Some of these difficulties are also present in this work, with data imbalance and data quality (and reliability) being the most important ones.

The authors in [van Mens et al. 2020], in contrast, focused on suicidal behaviour prediction using traditional ML techniques. Their model was built using data from a population-based longitudinal study with Scottish young adults (ages 18 to 34) who have undergone an interview that collected psychological and social measures. Their results show that algorithms based on decision trees were better at predicting suicidal behaviour than regular logistic regression, achieving a sensitivity of 0.47 and a specificity of 0.91 for suicide attempt prediction.

Another work focused on suicide risk prediction that relates to our study, but explores risk assessment in a temporal manner, is [Su et al. 2020]. It targets short and long term prediction of suicide risk using longitudinal data collected from structured electronic health records (EHRs) from the Connecticut Children's Medical Center, and compares the performance of their proposed model to L1 logistic regression. Their proposed model shows a good predictive power, achieving a sensitivity of 0.62 for a specificity of 0.90. It shows a better performance for shorter prediction windows.

This present work differs from [van Mens et al. 2020] and [Su et al. 2020] in many ways. Our work is the first to explore data from CEPAI (the Brazilian emergency mental health centre) to understand factors related to suicide risk prediction in a emergency context. For this reason, we start with simple models, such as logistic regression, random forest and gradient boosting (as was done in [van Mens et al. 2020]), to thoroughly understand what can be learned from the data. We extend previous literature by providing models for predicting suicide risk with a focus on children's and teenagers' health records using a richer set of features, e.g., the reasons why the patient decided to look for help (that are not often used by standardized risk assessment tools), which were considered the most important predictors in the models evaluated and were not found on related works. Due to the characteristics of the data used in this present work (i.e., they do not come from a longitudinal study), the metrics achieved by our models can not be directly compared to other works previously cited, but they are very impressive on their own and show the excellent predictive power of our approach.

## 3. Dataset Characterization

The dataset was created from a set of 2,365 paper health records of patients admitted to the Psychic Center for Children and Adolescents (CEPAI), in Belo Horizonte, Brazil. From these, there were 1,689 patients, i.e., some patients were admitted more than once. The individuals were aged between 0 and 20 years and data were collected from June 2017 to May 2018. Population data consist of 27 personal and socio-demographic features (race, gender, place of birth, residential location, school situation, living situation, the patients' legal responsible) and 123 clinical features (including the reasons psychiatric assistance was sought, family history of mental disorders, information on substance abuse, psychiatric diagnosis received after the treatment in the facility, neuro-psychomotor development delay and previous traumatic events, among others).

After creating the dataset, in a data preparation phase, we removed 74 admissions where patients: (i) either had no information at all (5); (ii) had only personal information (37); (iii) had no information about the motivations for looking for help or about the diagnosis (31); (iv) or had not completed the screening stages at the centre (1). Regarding the features, 53 of them contained textual information (for example, there's an attribute containing the written reason why the patient left school), so those were discarded. Out of the remaining ones, we selected, with the assistance of psychiatry professionals, 57 features that were considered more informative for the tasks at hand. Also, 21 of the features discarded concerned suicidal behaviour, but consisted mainly of missing values.

We understood data imputation was risky in this context, so we just discarded those. The 33 categorical features (among the 57) were binarized to be further used in our models: for example, the five original features regarding the different diagnoses that a patient could receive originated 21 “new” binary features, as there were 21 possible diagnoses in the original dataset. By the end of this process, the dataset consisted of 2,291 admissions regarding 1,689 unique patients, 1,072 male and 617 female, described by 154 features.

To provide a brief description of our population, Figure 1 shows the age distribution in the dataset. Most patients ($\approx 62\%$) are aged between 13 and 18 years, which corroborates Brazilian and International reports revealing that suicidal thoughts and suicide rates are highest among this age group [1]. Moreover, Figure 2 presents the number of times (admissions) a patient was admitted to CEPAI. Although most of the patients were admitted only once, $21\%$ of the individuals were admitted at least twice in the studied period, revealing that suicidal thoughts and attempts may persist and should be seriously taken into consideration to prevent the worst outcome.

Figures 3 and 4 show the distribution of the motivations for looking for help and diagnoses given by the psychiatric professional after the assessment, respectively. The most common reasons patients (or their guardians) look for help are related to agitation, aggressiveness, irritability and learning difficulties, which are not traditionally associated with suicidal behaviour. The factors commonly related to suicidal behaviour are depression, self-harm, anxiety and self-aggressiveness, which are present in 19,07%, 14,7%, 14,36% and 5,28% of the admissions registered. For the diagnoses, depression and general mood disorders appear in 21,69% of all admissions, being the most common diagnosis, even though there are many other symptoms the patients present that are related to impulsive and aggressive behaviours.

Recall that we are interested in predicting suicide attempts and suicide ideation. These are among the motivations given by patients or their guardians to seek psychiatric care. Given that the data used in this paper is not derived from a longitudinal study, the “suicide attempt” and “suicidal ideation” motives were modelled as the target variable in the classification task.

Following this approach, the dataset had 309 out of 2,291 cases of suicide ideation (13,48%) and 323 out of 2,291 cases of suicide attempts (14,09%). There are more occurrences of suicide attempts than suicide ideation because sometimes the attempt is an impulsive act, and not one that was planned out.

## 4. Methodology

In order to build our suicide attempts and suicide ideation prediction models, we follow the classic methodology for dealing with data. Our first step was data preparation, as reported in the previous section, followed by a feature selection phase. Next, we tested a few models to choose the most appropriate. These tests also evaluated the performance of the models with data oversampling and context-sensitive approaches (such as parameter tuning). Finally, we analyzed, for each task, the most important features according to the model with the best performance and used the Shapley Additive Explanations (SHAP for short) [Lundberg et al. 2020] to better understand factors associated with both outcomes.

[1] Brazil registers more than 6,000 suicides in teenagers in 5 years., Preventing suicide: A global imperative

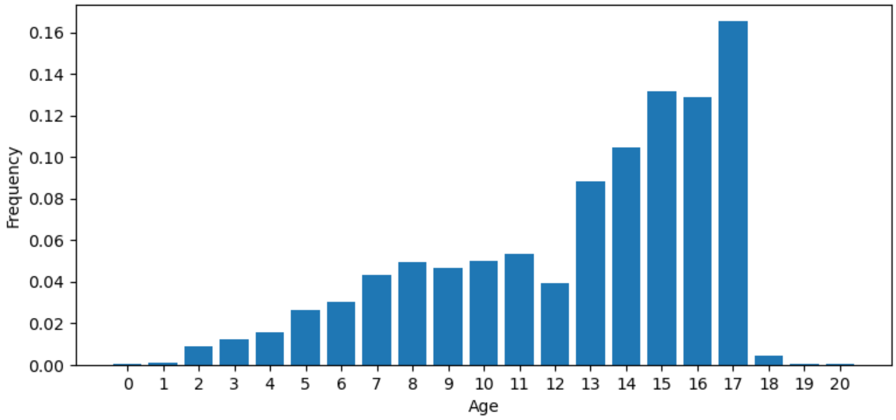


**Figure 1. Patients age distribution.**

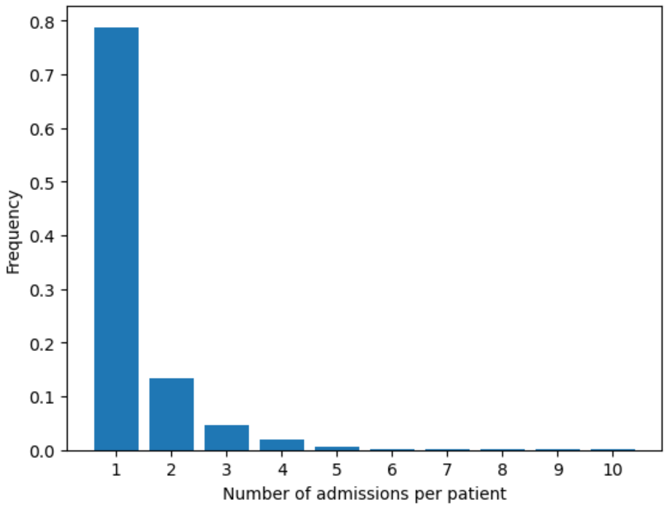


**Figure 2. Number of admissions per patient.**

The models were trained on two different sets of features. The first considers all features available, and is referred to in the results table as "All Features". The second (referred to as "Feature Selection") contains only the reasons patients looked for help (motivations), the diagnoses received in the medical assessment, if the patient was hospitalized, if it was the first time the patient was admitted to CEPAI, and personal data, including gender, age and how many people live with the patient. This latter subset was generated using expert knowledge and was inspired by the work of [van Mens et al. 2020], where models with pre-selected features performed better than the models trained with all features available. When training the models, the target variables were obviously excluded from the features used as predictors in the training process. Also, intuitively, suicide ideation comes before a suicide attempt. Hence, ideation was used as a predictor for a suicide attempt, but the contrary was not done, even though there are instances in the data where the attempt happens without any ideation being reported.

Both datasets were given to three machine learning (ML) models: Logistic Regression (LR), Random Forest (RF), and eXtreme Gradient Boosting (XGBoost). LR was chosen for being easily interpretable, while RF and XGboost are tree-based and capture

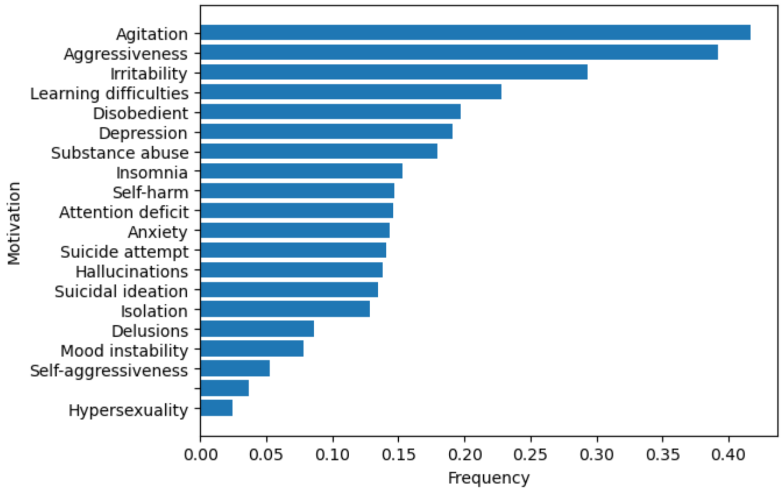


**Figure 3. Motivation distribution in the data.**

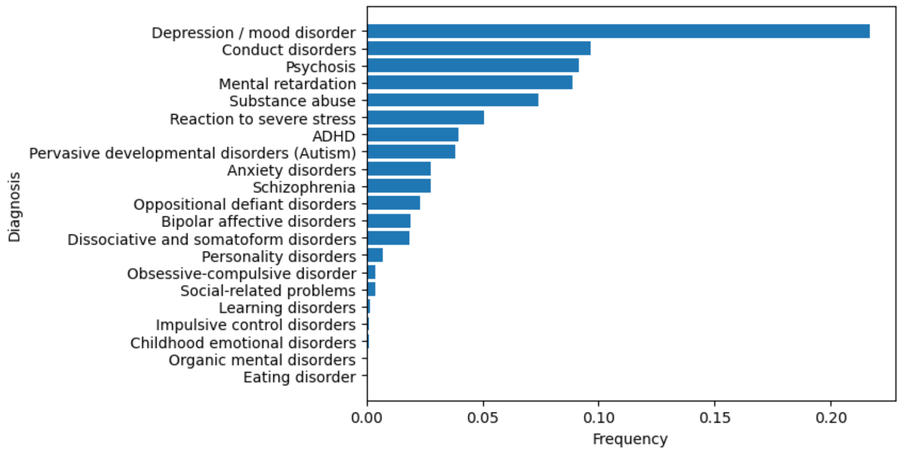


**Figure 4. Diagnosis distribution in the data.**

non-linear relationships in the data and their model can be interpreted. In our experiments, we used scikit-learn[2] and XGBoost[3] libraries. The user is referred to [Bishop 2006] for a description of the aforementioned algorithms.

Next, due to the imbalanced nature of the data – although the number is higher than the cases of suicide in the whole population, which can be attributed to data bias (data was collected in an emergency room scenario), we oversample instances with suicidal behaviour to reach 30% of the training set. Note that data in the test set were kept with their original distribution.

Finally, because every model used in this study has some kind of feature importance attributed to the features it receives as predictors, we were able to compare the best features for each of the classification tasks. In addition to that, SHAP values were generated. SHAP is an approach based on game theory used to explain the output of any machine learning model. They show how much each feature contributes, positively or

[2]Scikit-learn user guide
[3]XGBoost user guide

negatively, to the target feature. The interpretation of SHAP is similar to the one of feature importance, but SHAP is able to show whether each feature has a positive or negative relationship with the predicted value.

## 5. Experimental setup

This section describes the way experiments were performed to evaluate the models for predicting suicide attempts or suicidal ideation. The three models aforementioned (LR, RF and XGBoost) were tested with the complete set of attributes (referred to as "All Features") and with the partial set of features (referred to as "Feature Selection"). Both experiments were also performed with and without oversampling.

All models were evaluated using a 10 fold cross-validation strategy [Bishop 2006]. This strategy runs the model 10 times, each of them using 9 out of 10 folds to train the model and the remaining one to test the model. In each run, a different one of the 10 folds is used for testing the model. In order to avoid data leakage, we made sure that all admissions of the same patient were added to a single fold, while also maintaining the same proportion of positive cases of the target variable between all folds. We are aware that the data is not independent and identically distributed (i.i.d), like it is commonly assumed by most machine learning algorithms, but the strategies of avoiding data leakage and evaluating the results based on cross-validation give out robust results.

The parameters of RF and LR were kept in their default values. For XGBoost, a grid search considering the percentage of features used to train a tree, the tree depth, and the ratio of training instances in a subsample was executed in order to find the parameters that return the best model (as we'll see shortly, the model with the biggest sensitivity).

Four metrics were used to evaluate the performance of the models: Area Under the Receiver Operating Characteristic Curve (ROC AUC) – which shows the performance of a classification model at all classification thresholds of true positive rate and false positive rate; Positive Predictive Value (PPV) – the proportion of correctly predicted positive instances out of all predicted positives; sensitivity – the proportion of actual positive instances correctly identified by the classifier; and specificity – the proportion of actual negative instances correctly identified by the classifier.

However, in the context of this project, the sensitivity score is considered more important than the other metrics, because the impact of not identifying a positive case of suicide attempt (or ideation) is considered more serious than misclassifying a non-suicide attempt (or ideation). For this reason, for the results of logistic regression, we experimented with different thresholds for classifying an example as a positive case of suicidal behaviour or not.The default value is 0.5, but we performed experiments with values in the interval from 0.3 to 0.5 in 0.05 steps.

In traditional ML models, there is a phenomenon called sensitivity-specificity trade-off. Increasing sensitivity typically involves using a lower classification threshold, which means more instances are classified as positive, including both true positives and false positives. This leads to a higher sensitivity as more actual positive instances are correctly identified.

However, when the classification threshold is lowered, it also increases the chances of including false positives, which reduces the PPV (precision to identify pos-

**Table 1. Suicide attempt prediction results.**

| All Features | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Without oversampling | | | | With oversampling | | | |
| **Models** | **AUC** | **PPV** | **Sensit** | **Specif** | **AUC** | **PPV** | **Sensit** | **Specif** |
| LR | 0.6509 | 0.4129 | 0.3949 | 0.9069 | 0.6865 | 0.3525 | 0.5377 | 0.8353 |
| RF | 0.5339 | 0.4772 | 0.0766 | 0.9913 | 0.5606 | 0.5447 | 0.1411 | 0.9801 |
| XGB | 0.6228 | 0.4259 | 0.3177 | 0.9278 | 0.6163 | 0.4061 | 0.3084 | 0.9242 |
| **Feature subset** | | | | | | | | |
| | **Without oversampling** | | | | **With oversampling** | | | |
| **Models** | **AUC** | **PPV** | **Sensit** | **Specif** | **AUC** | **PPV** | **Sensit** | **Specif** |
| LR | 0.6403 | 0.4412 | **0.3509** | 0.9298 | 0.7382 | 0.4237 | 0.6087 | 0.8677 |
| RF | 0.6062 | 0.5121 | 0.2491 | 0.9633 | 0.9067 | 0.8808 | **0.8319** | 0.9814 |
| XGB | 0.6163 | 0.4043 | 0.3036 | 0.9291 | 0.9017 | 0.8559 | 0.8261 | 0.9773 |

itive instances in the data), since it is the proportion of true positives out of all predicted positives, so a higher number of false positives can decrease the precision of the classifier. Because of this trade-off, we chose 0.35 as our threshold in all logistic regression models, since it seems to balance well between a higher sensitivity and a PPV that's not so bad.

## 6. Experimental Results

This section shows the results obtained for predicting both suicide attempts and suicidal ideation. Tables 1 and 2 show the results considering models to predict suicide attempts and ideation, respectively. In both tables, LR results consider a threshold of 0.35 for positive classification. The best sensitivity found for each task is highlighted in bold. Note there is not a noticeable trend in the results.

The best overall performance, in both tasks, was achieved by the RF model, trained on oversampled data using only the attributes present in the group that used features manually selected by the specialists. For suicide attempt prediction, the PPV for the best-performing model was 0.8808, the sensitivity 0.8319, and the specificity 0.9814. For suicide ideation, the best model achieved a PPV of 0.8360, a sensitivity of 0.8422, and a specificity of 0.9750.

In sum, the metrics achieved with the Random Forest model were satisfactory, with a PPV over 0.8 for both models, meaning that among 10 people predicted as suicidal by the model, more than 8 would actually present suicidal behaviour. Having both a high sensitivity (superior to 0.83) and a high specificity (superior to 0.97) means that we are predicting correctly 83 among 100 positive cases and 97 among 100 negatives. These results show that this model can be used to assist health professionals, acting as an additional tool to be consider in the identification of patients at risk of suicide behaviour. It could be useful as a pre-screening tool, which would give people with high suicide risk probability a higher priority to see a doctor, for example.

So far we have looked at the prediction models, but one of the main objectives of this study is to identify the most relevant factors able to predict suicide attempts and suicidal ideation so we can intervene. Hence, we now turn to the feature importance given by the methods to different attributes.

**Table 2. Suicide ideation prediction results.**

| All Features | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Without oversampling | | | | With oversampling | | | |
| **Models** | **AUC** | **PPV** | **Sensit** | **Specif** | **AUC** | **PPV** | **Sensit** | **Specif** |
| LR | 0.6828 | 0.5066 | 0.4367 | 0.9289 | 0.7182 | 0.4179 | 0.5685 | 0.8678 |
| RF | 0.5746 | 0.6683 | 0.1619 | 0.9874 | 0.6025 | 0.6011 | 0.2328 | 0.9722 |
| XGB | 0.6440 | 0.4611 | 0.3495 | 0.9384 | 0.6480 | 0.4667 | 0.3632 | 0.9329 |
| **Feature subset** | | | | | | | | |
| | Without oversampling | | | | With oversampling | | | |
| **Models** | **AUC** | **PPV** | **Sensit** | **Specif** | **AUC** | **PPV** | **Sensit** | **Specif** |
| LR | 0.6789 | 0.4807 | **0.4274** | 0.9304 | 0.7658 | 0.4468 | 0.6555 | 0.8761 |
| RF | 0.6056 | 0.4514 | 0.2587 | 0.9525 | 0.9086 | 0.8360 | **0.8422** | 0.9750 |
| XGB | 0.6343 | 0.4605 | 0.3283 | 0.9403 | 0.6639 | 0.4316 | 0.4123 | 0.9155 |

**Table 3. Rank of the most important features according to the best models for suicide attempt and ideation prediction.**

| | **Suicide attempt prediction** | **Imp** | **Suicide ideation prediction** | **Imp** |
|---|---|---|---|---|
| 1 | Number of people the patient lives with | 0.090 | Depression (motivation) | 0.152 |
| 2 | Depression (motivation) | 0.066 | Self-harm | 0.130 |
| 3 | Self-harm | 0.064 | Male | 0.049 |
| 4 | Male | 0.058 | First time in CEPAI | 0.047 |
| 5 | Suicidal ideation | 0.057 | Nervousness/irritability | 0.044 |
| 6 | Agitation | 0.056 | Agitation | 0.041 |
| 7 | First time in CEPAI | 0.054 | Isolation | 0.038 |
| 8 | Was hospitalized | 0.043 | Substance abuse | 0.035 |
| 9 | Insomnia | 0.040 | Anxiety | 0.035 |
| 10 | Substance abuse | 0.037 | Was hospitalized | 0.034 |

Table 3 shows the 10 most important features for the best model of each task, together with their attributed importance (based on the mean impurity decrease, and that appears in the column "Imp"). Observe that 7 out of the 10 most important features for predicting both suicide attempt and ideation are the same for both models, which evidences the close relationship between both events that we want to predict. Also note that, in our dataset, depression can be one of the reasons the patient looked for help (codified as "Depression (motivation)") or one of the diagnoses given by the mental health specialist (codified as "Depression/mood disorder").

We also analyzed the effect each feature has on the model output using the SHAP values. For both tasks, these values are shown in Figures 5 and 6, respectively. When reading these plots, for each feature, the horizontal plot shows whether the effect of that value is associated with a higher or lower prediction. The colours show if that attribute is high (in red) or low (in blue).

For example, higher values of "Agitation", "Male", "First time in CEPAI", and "Aggressiveness" tend to produce lower predicted values for a suicide attempt, meaning that male patients on their first admission, showing signs of agitation and aggressiveness, are less likely to perform a suicide attempt than the patients characterized by the opposite values of these binary variables (i.e., recurrent female CEPAI patients, with no signs of

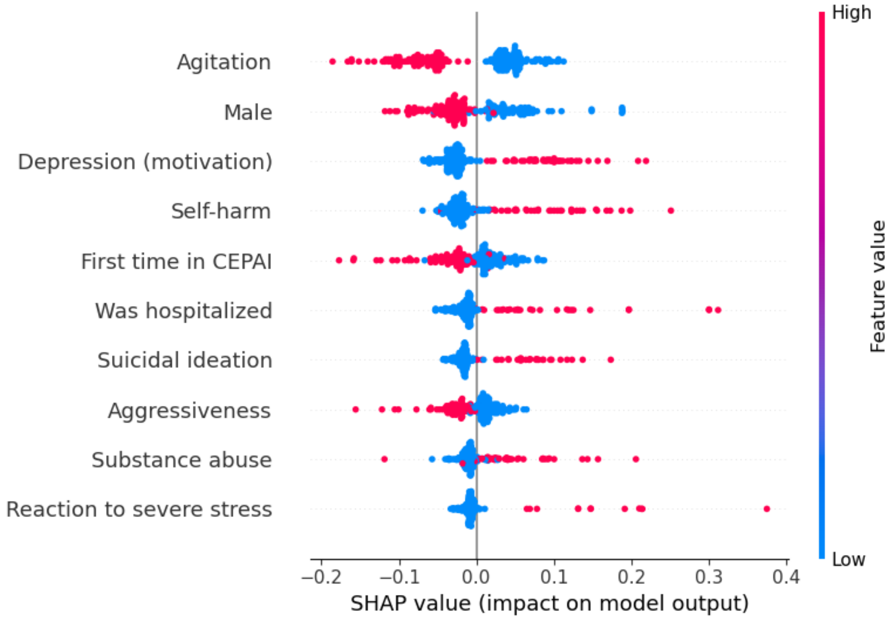


**Figure 5. SHAP values for suicide attempt prediction**

agitation or aggressiveness). On the other hand, patients with depression, self-harm, suicidal ideation and substance abuse as motives for looking for help, diagnosed with reactions to severe stress and hospitalized in CEPAI are more likely to have a suicide attempt.

Many of these tendencies are reflected in the decisions made by the suicide ideation prediction model, with three other attributes having their impact considered: positive values for isolation tend to increase the risk of suicide ideation, while learning difficulties mostly generate lower predicted values for this target. The attribute "Hallucinations" did not seem to reflect in a linear way on the prediction.

## 7. Conclusions and Future Work

Suicide is a challenging condition to predict. Both outcomes proposed in this work (suicide attempt and ideation) involve complex and multifactorial conditions that no simple clue will respond to. There are many predisposing conditions, the frequency of the tragic event is low and the error is undesirable.

Here, we evaluate mostly motives and diagnoses given by the mental health professionals at CEPAI as predictors to create a model in a high-risk patient sample. The models were reasonably predictive in comparison to other medical predictors, but we still suggest a careful approach since the outcomes are so tragic. As previously discussed, the Random Forest model classifies both positive and negative cases of suicidal behaviour accurately. The experimental results suggest that our models, after more careful study and validation in other databases, can be used at CEPAI as a pre-screening tool, prioritizing higher predicted risk patients for immediate emergency care and identifying the risk factors involved in each individual case. Models to predict suicide ideation and attempt generate different results, which might happen because of the different timing between the two events, but there are common pathways and symptoms that should be carefully evaluated and replicated in other child and adolescent populations.

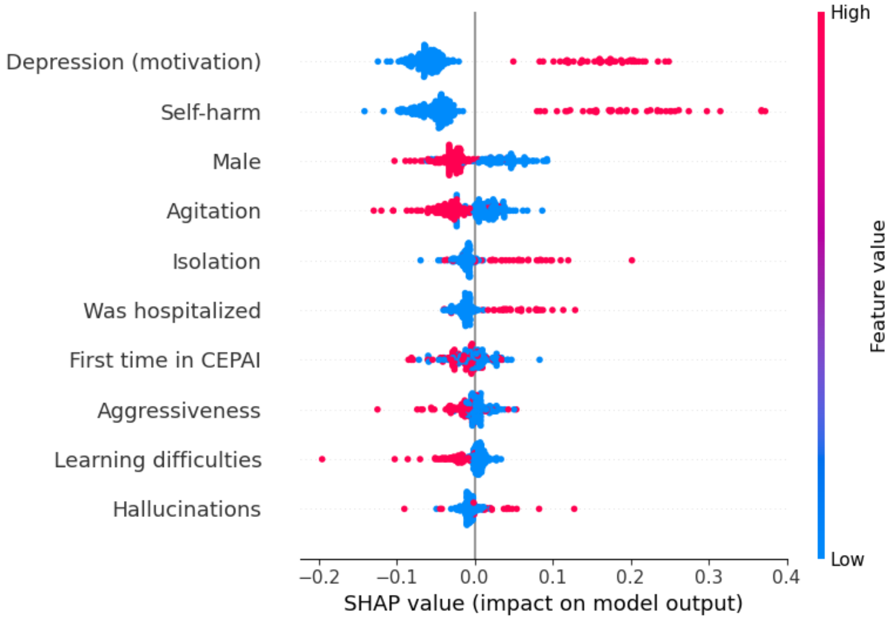


**Figure 6. SHAP values for suicide ideation prediction**

Depression is common sense and confirms many previous findings among the predictors. It was identified as a very important predictor in many of the models' tests (including the best ones), so it should be the primary target to be treated to avoid any outcome related to suicide.

When using suicidal ideation as the target feature, symptoms of delusions and learning disorders that are not so commonly associated with suicide attempts show the severity of the disorder underlying the search for help. For suicide attempts as the target feature, some interesting predictors emerged suggesting there are symptoms that should raise an alert to severe mental health cases, such as agitation and impulsivity-related disorders (such as personality disorders and bipolar disorders).

We are aware that this work presents preliminary results and still needs to be replicated and validated in other cultures and populations. In future works, we aim to explore the factors involved in multiple admissions of a single patient by, for example, adding a variable that indicates how many times a patient has been to CEPAI, and analyze how this can affect the final result. The coding of data might be a problem since it was done from medical records, in retrospective design with no contact with family or patients. However, the availability of data on suicide has been growing, indicating both that there is more data to work on and also that there is a clear and urgent need to have good predictors to anticipate and avoid this scenery.

## References

[Bishop 2006] Bishop, C. M. (2006). *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg.

[Carroll et al. 2014] Carroll, R., Metcalfe, C., and Gunnell, D. (2014). Hospital presenting self-harm and risk of fatal and non-fatal repetition: systematic review and meta-analysis. *PLoS One*, 9(2):e89944.

[Harris et al. 2019] Harris, I., Beese, S., and Moore, D. (2019). Predicting future self-harm or suicide in adolescents: a systematic review of risk assessment scales/tools. *BMJ open*, 9.

[Ji et al. 2021] Ji, S., Pan, S., Li, X., Cambria, E., Long, G., and Huang, Z. (2021). Suicidal ideation detection: A review of machine learning methods and applications. *IEEE Transactions on Computational Social Systems*, 8(1):214–226.

[Knipe et al. 2022] Knipe, D., Padmanathan, P., Newton-Howes, G., Chan, L. F., and Kapur, N. (2022). Suicide and self-harm. *Lancet*, 399(10338):1903–1916.

[Lundberg et al. 2020] Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., and Lee, S.-I. (2020). From local explanations to global understanding with explainable ai for trees. *Nature Machine Intelligence*, 2(1):2522–5839.

[Naghavi and Collaborators 2019] Naghavi, M. and Collaborators (2019). Global, regional, and national burden of suicide mortality 1990 to 2016: systematic analysis for the global burden of disease study 2016. *BMJ*, 364:l94.

[Navarro et al. 2021] Navarro, M. C., Ouellet-Morin, I., Geoffroy, M.-C., Boivin, M., Tremblay, R. E., Côté, S. M., and Orri, M. (2021). Machine Learning Assessment of Early Life Factors Predicting Suicide Attempt in Adolescence or Young Adulthood. *JAMA Network Open*, 4(3):e211450–e211450.

[Orri et al. 2021] Orri, M., Boivin, M., and Chen, C. e. a. (2021). Cohort profile: Quebec longitudinal study of child development (qlscd). *Soc Psychiatry Psychiatr Epidemiol*, 56.

[Su et al. 2020] Su, C., Aseltine, R., Doshi, R., Chen, K., Rogers, S. C., and Wang, F. (2020). Machine learning for suicide risk prediction in children and adolescents with electronic health records. *Translational Psychiatry*, 10(1):413.

[van Mens et al. 2020] van Mens, K., de Schepper, C., Wijnen, B., Koldijk, S. J., Schnack, H., de Looff, P., Lokkerbol, J., Wetherall, K., Cleare, S., C O'Connor, R., and de Beurs, D. (2020). Predicting future suicidal behaviour in young adults, with different machine learning techniques: A population-based longitudinal study. *Journal of Affective Disorders*, 271:169–177.