

# RecSys Challenge 2023: From data preparation to prediction, a simple, efficient, robust and scalable solution

MAXIME MANDERLIER, University of Mons (UMONS), Belgium

FABIAN LECRON, University of Mons (UMONS), Belgium

The RecSys Challenge 2023<sup>1</sup>, presented by ShareChat, consists to predict if an user will install an application on his smartphone after having seen advertising impressions in ShareChat & Moj apps. This paper presents the solution of 'Team UMONS' to this challenge, giving accurate results (our best score is 6.622686) with a relatively small model that can be easily implemented in different production configurations. Our solution scales well when increasing the dataset size and can be used with datasets containing missing values.

CCS Concepts: • **Information systems** → **Decision support systems**.

Additional Key Words and Phrases: online advertising, neural networks, missing values, embeddings, binary classification

## 1 INTRODUCTION

In the ShareChat & Moj apps, ads are shown to users and invite them to install applications on their smartphone. By collecting data about the users and the ads we might be able to predict if one ad will conduct the user to install the advertised application on his smartphone. In this challenge, we were given a dataset of impressions (ads showed to users) and the labels 'is\_installed' and 'is\_clicked'. The challenge is evaluated based on prediction of the 'is\_installed' (Log-Loss is the evaluated metric) label but prediction of the second label is also important for further analysis. We will show how to predict these two important output variables. [Section 2](#) describes how data is prepared, [section 3](#) describes our predicting model, [section 4](#) discusses about the achieved results, metrics, and observations, and [section 5](#) concludes this work. The code for predicting the results of the challenge as well as the code that was used to compute metrics described in [section 4](#) are available in our github repository<sup>2</sup>.

## 2 DATA PREPARATION

As any machine learning practitioner should know, the best model will give bad results if it is fed bad data. Data preparation takes most of the time in many cases and includes some steps as handling missing values, normalization, feature selection, etc. We show in this section what steps we apply to our data.

As described by ShareChat RecSys Team, the dataset consists of records that capture user and ad features (categorical, binary, numerical) and whether a click and/or install was generated by the user. Description of the features is not given and therefore cannot be used to understand the data. With features descriptions, other modeling choices could be made and probably improve the quality of the predictions. In particular, we could model users and ads separately if we had information on which features correspond to the users and which correspond to the ads. We split the features by categories that we handle separately. Each one of them receive adequate data preparation steps as showed in [subsection 2.1](#), [subsection 2.2](#) and [subsection 2.3](#).

### 2.1 Categorical data

There is 30 different categorical features ('f\_2' to 'f\_32'). First of all, we remove 'f\_7' from the data as that feature has one unique value (in the proposed dataset). It will therefore add no information

<sup>1</sup><http://www.recsyschallenge.com/2023/>

<sup>2</sup><https://github.com/MaximeUM/RecSysChallenge2023>

and can be ignored. Some values are missing (NaN) and values exist in the test set but are never seen in the training set. These categorical values will be handled by embedding layers in our model (see [section 3](#)) that will allow us to use the value 0 for 'no information' and avoid training on that information. Thanks to that, estimating missing values for the categorical features is not necessary. In the columns, numbers are associated to the features but are not continuous. We will re-code the information to have each feature (column) containing numbers from 0 to n where n is the number of distinct values for that feature. That step is necessary as the embedding layers need continuous integer values. NaN are replaced by 0 and values in the test set that don't exist in the train set are also replaced by 0 (treated as missing as we never saw them).

## 2.2 Binary data

9 different binary features ('f\_33' to 'f\_41') are given in the dataset. There are no missing values for these features (train set and test set). If we were confronted to a scenario with missing values, we would have replaced the missing values with coherent values as shown in [subsection 2.3](#). Being binary values, all values are between 0 and 1 for each feature and normalization is not necessary. No (further) treatment is needed and these data will be directly given as inputs to the model.

## 2.3 Numerical data

Finally, the dataset contains 38 numerical features ('f\_42' to 'f\_79'). This category of data will be the one for which we really need to apply data preparation steps.

*2.3.1 Missing values estimation.* This part of the dataset contains missing values. One way to handle it would be to remove each row containing a NaN value (and remove the corresponding row from the binary data and categorical data). Doing that means that we would lose all the features for only one missing feature. Another idea would be to use masking in the neural network to avoid training for unknown values. However, proceeding like that doesn't allow the model to get a full representation of the inputs. Moreover, the problem of missing values is pushed to the next step (fitting a model) and will therefore not allow us to easily change from one model to another. Well done data preparation should allow us to use other models in [section 3](#) without much change. As discussed in [[1](#), [2](#), [5](#), [8](#)], better ways to handle missing values are available to us and will be used in this work.

In our case, we will replace the missing values by doing estimations based on the features of the dataset. Different methods are tested (replacement by the mean, by the median, by zero or by using the other features). The last method gave the more robust results and is the more elaborate. For each feature, we represent that feature as a function of the other features (only the numerical features here). Each time we encounter a missing value, we use the function defined for the feature to predict the value based on values of the other features. This method supposes that the variables are not independent and that similar values on some attributes should reflect a similar value on the attribute we try to estimate. The implementation we used is given by *scikit-learn*<sup>3</sup>.

*2.3.2 Data scaling.* Also, for these features the ranges of values are really diverse. For instance, 'f\_74' ranges from 0 to 0.1157 while 'f\_64' ranges from 0 to 415,706,830,424. If the model is trained with non-normalized data, it would be difficult for features as 'f\_74' to get importance (these features would merely be ignored). All features will be normalized between 0 and 1 using the min-max normalization given by the following formula:

$$x_i = \frac{x_i - \min(X)}{\max(X) - \min(X)}$$

<sup>3</sup><https://scikit-learn.org/stable/modules/generated/sklearn.impute.IterativeImputer.html>

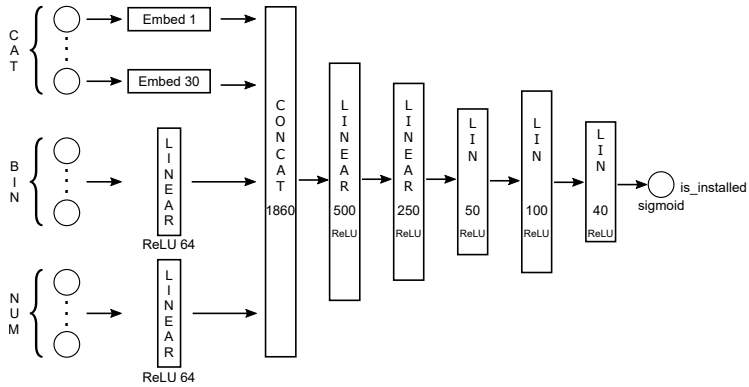


Fig. 1. Predict 'is\_installed'

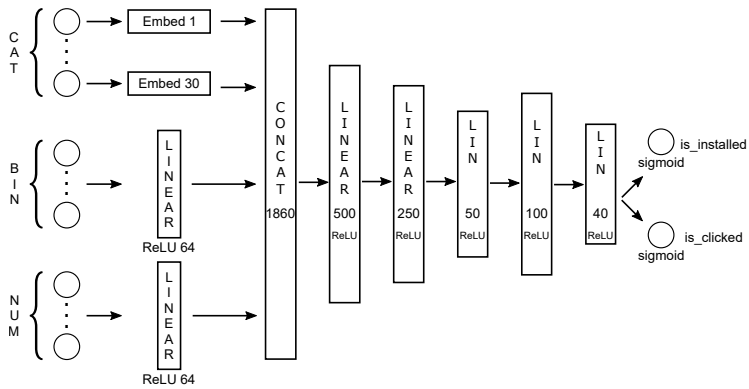


Fig. 2. Predict 'is\_clicked' and 'is\_installed'

where  $X$  represents one feature vector (e.g. 'f\_74') and  $x_i$  one value of this vector. The formula is applied to each value of the feature vector and repeated for all feature vectors.

### 3 FITTING THE MODEL

For this challenge, as evaluation was associated to the Log-Loss, a model that is good at predicting probabilities was necessary. For that reason, we decided to use a deep neural network with a final *sigmoid* function. This choice is also motivated by the use of neural networks in other works related to online advertising [4, 6, 10, 11], proof that this kind of model is suitable for the given problem. Using the neural network was also the opportunity for us to use each type of data (categorical, binary and numeric) in a different way and without doing too much transformation. In other words, the neural network is appropriate to use with multiple inputs (and also multiple outputs as we will see). We describe the design of our models (Figure 1 and Figure 2) in the following of this section.

#### 3.1 Defining the inputs

**3.1.1 Categorical data.** The number of distinct values differs from one categorical feature to another one. One easy way to proceed would be to define, for each feature,  $n$  input neurons (with  $n$  being the number of distinct values for that feature). Some features have a lot of distinct values and each time we have a value of 1 for the corresponding neuron we have  $n-1$  zeros input neurons. For

that reason, we will use embedding layers. An embedding layer encodes the values of the features given  $m$  neurons ( $m$  is to be chosen). Instead of having one neuron for each distinct value, we use a combination of a smaller number of values to encode the same information (compression). If we have  $n \leq 256$  distinct values (excluding the value 0 that is reserved for missing values), we choose  $n$  as the number of outputs of the embedding layer as it would make no sense to choose more output nodes than there is distinct values. For features with more than 256 distinct values, we encode the information using 256 neurons.

**3.1.2 Binary data.** For the binary data, we use one input neuron per feature. After that input, we connect a 64 neurons linear layer with *ReLU* activation function.

**3.1.3 Numerical data.** For numerical data, the same logic as for binary data is followed, one input neuron for each numerical feature. We also connect these neurons to a layer of 64 neurons with *ReLU* activation function.

**3.1.4 From the inputs to the output(s).** Each categorical feature embedding, the binary linear layer and the numerical linear layer are concatenated together. Some linear hidden layers are piped (with *ReLU* activation functions). Finally, one (Figure 1) or two (Figure 2) output layers are proposed at the end of our model. A first version with only one output layer is proposed to predict 'is\_installed'. The idea behind that is that it is the only output evaluated in the challenge and it accelerates the training for different configurations during the experiments. We then propose a version of the algorithm that contains two outputs ('is\_clicked' and 'is\_installed'). In section 4, we will investigate how this second output affects the predictions.

## 4 EXPERIMENTS

In this section, we describe the experiments and the achieved results of our model. First, we discuss about the training and validation procedure and secondly about how we obtained the results for the challenge. We then compare the models defined in section 3. All our experiments were made using *TensorFlow* for the model (neural network) and *scikit-learn* for data preparation and evaluation.

### 4.1 Validation procedure

Neural networks are prone to overfitting [9]. If we train the model for too much epochs, the loss will surely decrease. However, the model will perform very poorly on the test dataset as it will only be able to predict what it has already seen.

We randomly split our training data into a training set and a validation set (ratio 3:1). The loss on the validation set will be used to monitor training and to keep the weights for which the validation loss is the best (lowest). We monitor the loss (binary crossentropy) that is the equivalent of the Log-Loss metric used for evaluation in this challenge. If this loss stops decreasing, training will be stopped (after some epochs to be sure that it was not a momentary increase) and the best model (lowest validation Log-Loss) kept for inference.

### 4.2 Metrics

In this challenge, solutions are evaluated based on the Log-Loss, meaning that the probability (confidence) of our predictions is important. The probabilities are for sure important for business because different actions could be pursued based on the confidence we have in our model. For example, if the probability of installing the application is 0.51, one more ad (or more) could be showed to the user to increase the chances that he installs the app. On the contrary, if the probability is high it might not be useful to show another ad. However, the model can be not so good to predict probabilities but still very good at predicting binary answers. We will monitor the 'no information

Table 1. Metrics for the training set and validation set (predict 'is\_installed')

	Training set (75%)	Validation set (25%)	Training set (100%)
Log-Loss	0.3106	0.3177	0.3088
NIR	0.8258	0.8265	0.8260
Accuracy	0.8237	0.8207	0.8266
TPR (Recall)	0.5425	0.5324	0.5318
TNR (Specificity)	0.8830	0.8812	0.8888
Precision	0.4943	0.4849	0.5018
F <sub>1</sub> Score	0.5173	0.5075	0.5164

rate' (NIR), the accuracy for a model always predicting the majority class; the accuracy; the true positive rate (TPR) or recall; the true negative rate (TNR) or specificity; the precision and the F<sub>1</sub> Score that are standards metrics used for evaluation of binary classification problems [3, 7].

The results obtained on the dataset are given in Table 1. We show results for the scenario where we split the dataset in a training and validation set, only way for us to compute metrics on data that were not seen by the model. The test set cannot be used as we don't have the associated labels and will only be used to submit predictions for this challenge. We also give the values of the metrics for the case in which we use all the data for training as it is the one we use for predicting labels of the test set. For training a model on the whole dataset, we use the same number of epochs that were found by the train/val experiment (3 epochs).

### 4.3 Results interpretation

The NIR value in our case corresponds to always predict 0. We can observe that in the scenario where we split the data, the NIR is better than the accuracy (training set and validation set); the accuracy is slightly better than the NIR when we use the full dataset to train the model. It might then seem useless to use a model for prediction facing this reality. However, it might also seem really important to predict well the positive class, meaning to know when a combination of user and ad characteristics will lead to installing an application on the smartphone. If we always predict zeros, the true positive rate (TPR) is equal to 0. By using our model, we can observe that the TPR is higher than 0.5, meaning that we are able to detect more than half of the positive instances. The precision is also close to 0.5 meaning that half the time we say that a combination of features will lead to an install, it will. The TNR is much higher, meaning that it is easier to detect the instances of the negative class (majority class). Our model is then able to distinguish between positive and negative instances, which is good news. The Log-Loss is given in Table 1 but is difficult to interpret as it is. We can compare it to other models developed for this challenge.

### 4.4 Increasing Recall and Precision

If increasing the recall is the goal we want to pursue, maybe that monitoring the val\_loss is not what we should do. We show in Table 2 that if we train the model for more than three epochs, we get worse results for accuracy and log-loss (which is not a good news from the challenge perspective) but increase other important metrics. As reflected in Table 2, increasing TPR tends to decrease Precision. F<sub>1</sub> Score regroups these metrics and doesn't significantly decrease (but still is). From a business perspective, the most important metrics should be defined and used to monitor training. These important metrics are likely to evolve with time as new considerations will be taken into account. Our model is easily trainable for the metrics chosen in the future.

Table 2. Metrics for some training scenarios ('is\_installed')

	5 epochs		10 epochs		50 epochs	
	Train. set (75%)	Val. set (25%)	Train. set (75%)	Val. set (25%)	Train. set (75%)	Val. set (25%)
Log-Loss	0.2931	0.3242	0.2364	0.3930	0.0663	1.2485
NIR	0.8258	0.8265	0.8258	0.8265	0.8258	0.8265
Accuracy	0.8294	0.8208	0.8313	0.8016	0.8597	0.7829
TPR (Recall)	0.5718	0.5426	0.6626	0.5679	0.8190	0.5777
TNR (Specificity)	0.8838	0.8792	0.8668	0.8507	0.8682	0.8259
Precision	0.5092	0.4854	0.5121	0.4441	0.5673	0.4107
F <sub>1</sub> Score	0.5387	0.5124	0.5777	0.4984	0.6703	0.4801

Table 3. Metrics for the training set and validation set (two outputs model)

	Output 'is_clicked'		Output 'is_installed'	
	Train. set (75%)	Val. set (25%)	Train. set (75%)	Val. set (25%)
Log-Loss	0.3117	0.3192	0.3117	0.3191
NIR	0.7800	0.7807	0.8258	0.8265
Accuracy	0.6751	0.8217	0.8242	0.8213
TPR (Recall)	0.1909	0.5293	0.5440	0.5347
TNR (Specificity)	0.8117	0.8831	0.8833	0.8815
Precision	0.2224	0.4874	0.4957	0.4865
F <sub>1</sub> Score	0.2054	0.5075	0.5187	0.5095

#### 4.5 Results for challenge submission

For the test set, we are not able to compute the metrics defined in [subsection 4.2](#) as we do not have the labels. We can only indicate the best Log-Loss value obtained in the challenge: 6.622686 (Log-Loss value with transformation factor introduced by ShareChat RecSys Team to avoid cheating).

#### 4.6 Two outputs model: does it make a difference?

We analyze if asking the model to predict one more output affects its ability to correctly predict the most interesting output ('is\_installed'). For that, we will compute the metrics described in [subsection 4.2](#). The same metrics are given for both outputs in [Table 3](#). This experiment shows that we can achieve good results with the two outputs model and the values are similar to the ones in [Table 1](#). For the 'is\_clicked' output, the metrics get values that are significantly worse than the one obtained for 'is\_installed'. This observation should be explained by the model not being sufficiently trained to correctly predict this output (we monitor the output 'is\_installed' to stop training). We will investigate that hypothesis in [subsection 4.7](#).

#### 4.7 Predict 'is\_clicked'

Following the poor results we got for the 'is\_clicked' output using our two outputs model, we try to predict only 'is\_clicked' with the one output model defined in [Figure 1](#) in which we simply replace the output 'is\_installed' by the output 'is\_clicked'. [Table 4](#) demonstrates that similar performance can be obtained for 'is\_clicked' (compared to the one output model predicting 'is\_installed'). The results in [Table 3](#) were poor as the training was not completely done for the output 'is\_clicked' in the two outputs models. Indeed, training the model with only 'is\_clicked' as output required 4 epochs while all the other models required 3 epochs.

Table 4. Metrics for the training set and validation set (predict 'is\_clicked')

	Training set (75%)	Validation set (25%)
Log-Loss	0.3373	0.3473
NIR	0.7800	0.7807
Accuracy	0.7873	0.7827
TPR (Recall)	0.6091	0.5963
TNR (Specificity)	0.8375	0.8350
Precision	0.5139	0.5038
F <sub>1</sub> Score	0.5575	0.5462

#### 4.8 A model to predict all the outputs

Sharing all the layers between the inputs and the outputs (see Figure 2) seemed to be a good idea to share a maximum of information and also to avoid increasing the size of the model that we want to keep as small as possible. In fact, it is not a good idea as it doesn't allow the model to learn different representations of these two output variables. The hypothesis made was that we could share knowledge as the output variables are similar but our experiments showed that we can't and that the outputs need to be treated independently. To use only one model, we need to duplicate all the layers that are used between the inputs and the outputs to allow the model to learn different weights values for each output. Moreover, we should stop training of each part of the model by monitoring the corresponding 'val\_loss'. This scenario is not so much better than using two models with one output, a solution that can also be used.

## 5 CONCLUSION

In challenges like the Recsys Challenge as in the daily life of researchers in machine learning and related fields, pursuing the best results based on a metric is often set as a goal. However, developing complex models to increase the 'prediction score' is not always necessary and the pros and cons should always be considered. With our model, we show that pretty good results can be achieved with a relatively simple model that can be easily used in many production environments. We also demonstrate that we can tune our model to achieve different results (monitor different metrics) and allow easy tuning following business needs.

## REFERENCES

- [1] Alan C. Acock. 2005. Working With Missing Values. *Journal of Marriage and Family* 67, 4 (2005), 1012–1028. <https://doi.org/10.1111/j.1741-3737.2005.00191.x> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1741-3737.2005.00191.x>
- [2] Marvin L Brown and John F Kros. 2003. Data mining and the impact of missing data. *Industrial Management & Data Systems* 103, 8 (2003), 611–621.
- [3] Gürol Canbek, Seref Sagiroglu, Tugba Taskaya Temizel, and Nazife Baykal. 2017. Binary classification performance measures/metrics: A comprehensive visualized roadmap to gain new insights. In *2017 International Conference on Computer Science and Engineering (UBMK)*. 821–826. <https://doi.org/10.1109/UBMK.2017.8093539>
- [4] Michael Fire and Jonathan Schler. 2017. Exploring Online Ad Images Using a Deep Convolutional Neural Network Approach. In *2017 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*. 1053–1060. <https://doi.org/10.1109/iThings-GreenCom-CPSCom-SmartData.2017.160>
- [5] Therese D. Pigott. 2001. A Review of Methods for Missing Data. *Educational Research and Evaluation* 7, 4 (2001), 353–383. <https://doi.org/10.1076/edre.7.4.353.8937> arXiv:<https://doi.org/10.1076/edre.7.4.353.8937>
- [6] Yanru Qu, Han Cai, Kan Ren, Weinan Zhang, Yong Yu, Ying Wen, and Jun Wang. 2016. Product-Based Neural Networks for User Response Prediction. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*. 1149–1154. <https://doi.org/10.1109/ICDM.2016.0151> ISSN: 2374-8486.

- [7] Sebastian Raschka. 2014. An Overview of General Performance Metrics of Binary Classifier Systems. arXiv:1410.5330 [cs.LG]
- [8] Joseph L. Schafer. 1999. Multiple imputation: a primer. *Statistical Methods in Medical Research* 8, 1 (1999), 3–15. <https://doi.org/10.1177/096228029900800102> arXiv:<https://doi.org/10.1177/096228029900800102> PMID: 10347857.
- [9] Xue Ying. 2019. An Overview of Overfitting and its Solutions. *Journal of Physics: Conference Series* 1168, 2 (feb 2019), 022022. <https://doi.org/10.1088/1742-6596/1168/2/022022>
- [10] Shuangfei Zhai, Keng-hao Chang, Ruofei Zhang, and Zhongfei Mark Zhang. 2016. DeepIntent: Learning Attentions for Online Advertising with Recurrent Neural Networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*. Association for Computing Machinery, New York, NY, USA, 1295–1304. <https://doi.org/10.1145/2939672.2939759>
- [11] Weinan Zhang, Tianming Du, and Jun Wang. 2016. Deep Learning over Multi-field Categorical Data. In *Advances in Information Retrieval (Lecture Notes in Computer Science)*, Nicola Ferro, Fabio Crestani, Marie-Francine Moens, Josiane Mothe, Fabrizio Silvestri, Giorgio Maria Di Nunzio, Claudia Hauff, and Gianmaria Silvello (Eds.). Springer International Publishing, Cham, 45–57. [https://doi.org/10.1007/978-3-319-30671-1\\_4](https://doi.org/10.1007/978-3-319-30671-1_4)