Frequent patterns of childhood overweight from longitudinal data on parental and early-life of infants health

Beatriz López[1][0000−0001−9210−0073], David Galera[1],
Abel López-Bermejo[2][0000−0002−5828−8911], and
Judit Bassols[2][0000−0001−8821−6655]

[1] eXiT research group, University of Girona, Spain
beatriz.lopez@udg.edu, raizen74@gmail.com
[2] Girona Biomedical Research Institute Dr. Josep Trueta (IDIBGI), Spain
{jbassols,alopezbermejo}@idibgi.org

Abstract. Childhood obesity is considered one of the main public health concerns. Research in the field of obesity detection and prevention is moving towards promising solutions thanks to the use of Artificial Intelligence applied to data from cohorts of children. Previous studies have analyzed the data without taking into account the relationship of data regarding when they are collected. In this work, frequent pattern mining is used to find the risk factors of childhood obesity, taking into account the relationship among the data gathered in different visits. The experiments carried out on the data collected from 386 children from Girona and Figueres (Spain) demonstrate the relevance of discriminant frequent patterns for childhood overweight prediction.

Keywords: Childhood obesity · Frequent pattern mining · Discriminant patterns.

1 Introduction

Childhood obesity is one of the most important public health problems of the 21st century worldwide due to its prevalence and its impact on both short and long-term health. In addition to its association with non-communicable diseases such as cardiovascular diseases, diabetes or cancer, children who suffer from it usually experience psychological and social problems, such as low self-esteem or bullying [11].

Artificial Intelligence, and in particular Machine Learning approaches, have been used to understand the factors of childhood obesity [2]. Most of the approaches conduct cross-sectional studies in which socioeconomic and healthcare data of the parents and children are analyzed to understand the key factors of obesity. Few approaches dealt with longitudinal data. One of the exceptions is [8] that uses prenatal, perinatal, postnatal and 2 to 90 months age infant data, to identify different groups of risks (stable obesity, latent-onset obesity, moderate decline obesity, and non-obese), and their corresponding profiles. However, no evidence was found regarding that pre and perinatal, and postnatal data predicts obesity. This is an interesting finding since previous cross-sectional studies identify the Body Mass Index (BMI) paternal values as key factors. Nevertheless, [8] does not consider the relation of the variables according to the different visits in which they have been gathered. Our research is involved in understanding frequent patterns for childhood obesity based on longitudinal data similar to the one studied in [8]. To that end, we employ frequent pattern mining approaches, so that sets of variables are identified as key features instead of variables in isolation. Moreover, managing longitudinal data poses a challenge in dealing with missing values, because the follow-up of the visits is sometimes discontinued. On the other hand, feature selection methods are required to achieve a good discriminating representation of patterns regarding the target variable (overweight/obesity or normal weight). So the contribution of this work involves coping with all the challenges to apply frequent pattern mining to the longitudinal data for characterizing childhood overweight.

The paper is organized as follows. First, the work is contextualized in the previous literature. In Section 3, details of the methodology and dataset are provided. Next, in Section 4 the results are described and discussed. We end with some conclusions in Section 5.

2 Related work

The management of longitudinal data for determining risk factors of childhood obesity has been a matter of concern in recent years. In [8], prenatal, perinatal, postnatal and 2 to 90 months age infant data, similar to the one proposed in this work, has been considered. However, [8] does not take into account characterizing the baby according to frequent patterns as we do. A more recent approach is [15], which accounts for a similar imbalanced dataset as ours (14% of obese children). Our main concern with this approach is that they are using the AUC value as a performance metric. Conversely, we are analysing recall together with precision and accuracy, taking into account the minority class. Moreover, they are using logistic regression and multivariable fractional polynomials, while we are trying to capture relations among the data with the use of frequent patterns. [4] utilized real-world electronic health record (EHR) data from the first two years of life to predict the obesity status at age five. This work also ignores the relation of the data. An interesting work is [6], conducted in Korea, where the results pointed out the importance of cultural factors; for example, maternal self-esteem has been identified as highly important, in detriment of physical activity among children that has been found as less relevant. We do not have this information gathered in our dataset, and we annotate that factor for future research.

The use of frequent pattern mining is not new in Medicine, however. For example, [3] uses pattern mining algorithms to identify trajectories of patients from EHR; [1] develops a new method for dealing with cardiovascular diseases

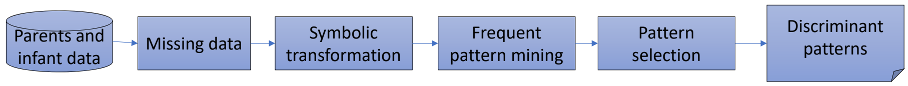

Fig. 1. Overview of the methodology.

management from EHR data; and [9] uses administrative data to predict medical in-hospital mortality regarding acute coronary syndrome.

The selection of frequent patterns has been also studied in the machine learning community, with special attention to their discriminant capability regarding some classification tasks. For example, [7] uses the FeatureMine method to predict all-cause mortality in T2D which has been particularly designed for sequential pattern mining methods. We are using in this work mRMR [14] since we are not considering sequential but frequent patterns. In that regard, some frequent pattern mining algorithms have been particularly designed for classification. For example, [5] defines a discriminative pattern mining method, and [13] introduces the concept of contrast sequence, which is later used in [12]. In our work we have kept two separate steps in the methodology: pattern discovery and pattern ranking and selection.

3 Materials and method

The methodology we propose is shown in Figure 1. We are given a dataset of paternal and early-life health data of the babies, with some overweighted/obese and non-obese children (class). In the first step of the methodology we deal with the missing data; next we transform the original data into symbols in order to apply pattern mining methods; afterwards, we obtain the frequent patterns for each class. Finally, a feature selection method is used to choose the best patterns according to their relevance for obesity prediction.

Frequent patterns can then be used for classification purposes, instead of the variables in isolation. An example of frequent pattern is the following: [$56 \leq$ Height at 2 months < 58, Smoking status of the father $= 0$, 3kg $\leq$ Birth weight mother $\leq$ 4kg].

3.1 Dataset description

The dataset contains two cohorts of babies, from Hospital Dr. Josep Trueta (Girona) and Hospital of Figueres, both in Spain, collected from 2008 to 2014. The inclusion and exclusion criteria for participation in the study was described elsewhere (study with code 2010056 of the Clinical Research Ethics Committee of the Girona University Hospital Dr. Josep Trueta).

A total of 1175 infants were provided, 212 variables each. Doctors who have been part of the project, select 36 variables as being the most related to the study: 10 categorical, 25 numerical, and the target variable (see Table 1 on

Appendix A for a description of the variables). Some of the variables are related to the mother, before and during pregnancy; some other ones about the father, and some other ones about the infant. The target variable is OBSESITY, which indicates if the baby has a normal weight (OBSESITY=1) or has overweight or obesity (OBSESITY=2), according to the BMI measured when the child is 5 years old. Column "Visit" indicates the different visits in which the variables have been gathered: 0 is related to variables at the moment the mother gets pregnant; 1 the information regarding the second trimester of pregnancy; 2 the third trimester of pregnancy; 3 the information at birth; 4 when the baby has 2 months old... and so on.

On the other hand, the samples of the infants where values were missing in more than 60% of the variables were eliminated. This high rate of missing values is due to the length of the study (more than 6 years) during which dropouts occurred, especially after the first 2 years of life. As a result of applying this filter, there were finally data from 386 infants (only 33% of the initial cases); 54 cases correspond to obese child (14%), so the dataset is quite imbalanced.

3.2 Missing data processing

The initial dataset contains 1304 missing values (9.6% of the information contained in the dataset), 947 of which belong to numerical variables and 357 to categorical variables (see Figure 5 of Appendix B for the distribution). Noteworthy to observe that we have only 101 instances (26%) with complete data (0 missing values). Therefore, we have conducted an imputation process to try to fill up as many missing values as possible.

First, for numerical values, a model-based imputation has been chosen (see Figure 2). We have discarded the use of the median or mean to avoid a bias to the high among of missing values we have. First, given the dataset D, for each numerical variable X_i that contains missing values, a subset of variables S, $X_i \in S$, with a high correlation with X_i has been identified. Second, a subset $D_i \subset D$ is obtained with all the instances that have the variable X_i also without information (missing). Third, for each instance I_j of D_i, a subset A is selected with the variables of S, $A \subseteq S$, that for the sample I_j do not have any missing value, $A \cap \{NA\} = \emptyset$. Fourth, a subset $D_j \subset D$ is gathered from the original dataset D, in which all the instances do not have missing values either in X_i or A, $D_j \cap \{NA\} = \emptyset$. Finally, a regression model is built from the dataset D_j to predict X_i. This process is repeated for each Instance I_j and missing value X_i.

With this procedure, the missing values of all the numerical variables are imputed, except for those of the months breastfeeding variable, which does not correlate with any other variable; for this variable the average has been directly imputed.

Regarding categorical data, missing values represent less than 3% of the dataset. To carry out the imputation, we follow a similar process as for numerical variables, but using the highest correlation variable instead of a regression model. In general, when a categorical variable X_i is missing for an instance, the variables correlated with X_i tend to also be missing; for this reason only 60 missing values

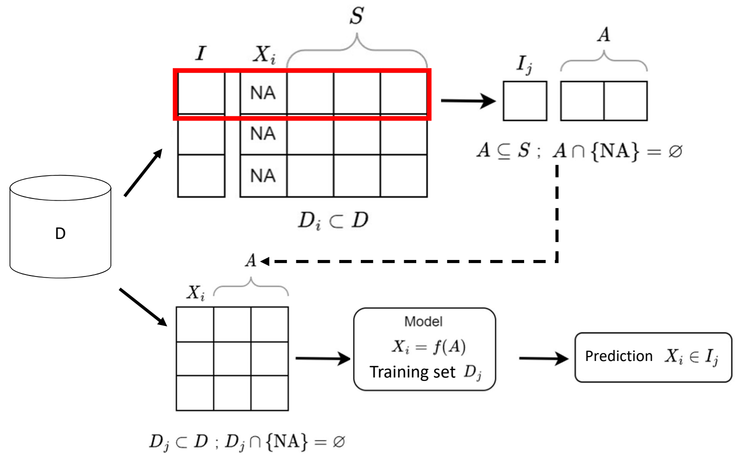

Fig. 2. Imputation of numerical values.

have been imputed with this method; for the rest, the mode has been directly imputed.

3.3 Symbolic transformation

Variables and their values should be transformed into symbols in order to apply pattern mining methods. Regarding numerical variables (a total of 25), a discretization method is applied. The variables BMI and systolic and diastolic blood pressure were discretized according to the thresholds established by the World Health Organization (WHO) and the American Heart Association (AHA) correspondingly. The remaining variables have been equally discretized in 4 bins (as BMI, and systolic pressure have four bins too).

Regarding categorical variables (10), 8 of them are boolean and have been transformed into one symbol per value. The remainder 2, have four possible categories, so 4 symbols have been also used for each. At the end we get 124 symbols, and the dataset is transformed according to this new codification. Note, that the data obtained after the symbol transformation process correspond to different visits (see column "Visit" in Table 1 of Appendix A). In particular, the data have information about 9 different visit periods, so each instance has 9-length episodes of data (itemsets).

3.4 Frequent pattern mining

Frequent pattern algorithms are applied to each class separately. The DefMe algorithm [10] has been used. It is an algorithm that mines generative patterns, a type of patterns that avoid generating the entire set of frequent patterns in classification tasks to increase the accuracy of predictions. The algorithm, as any frequent mining one, requires a support threshold (defined in [0,1]) that is related to the amount of patterns that could be mining (how many instances in the original dataset support the pattern; the higher value, the more instances are required). We have selected a support threshold of 0.3, since lower values may not generalize in new instances and the patterns with higher supports are difficult to find.

From the two sets of frequent patterns, one set per class, the symmetrical difference of the 2 sets is calculated. In this way there are non-common patterns for both classes. Finally, the original dataset is transformed using the patterns discovered. For each pattern, a binary variable has been defined. For each instance in the dataset, the value of each pattern is either 1 or 0, depending on whether this pattern is present or not. is met or not.

3.5 Pattern selection

To select the best patterns, first we obtain a ranking of them using the variant of the mRMR algorithm described in [14], and next we follow an incremental process in which the different patterns are successively considered, according to their ranking, to build a predictive model. Regarding mRMR, a decision tree classifier model is used to evaluate the discriminatory power of the first N patterns of the ranking. This algorithm has been chosen because it is the one that has obtained better scores compared to XGBoost, Random Forest, Support Vector Machine and Logistic regression (see Appendix C).

The selection of the best patterns has been done incrementally. In the first iteration, the pattern in the highest ranking position is used to build a classification model and obtain the performance metrics. In each iteration, the pattern that occupies the next position in the ranking is added, and its impact on the performance of the model is measured. When all the patterns have been introduced, or when the performance shows no further improvement, the process is stopped. The performance of the model is visualized based on the patterns and with the collaboration of the medical team, the best N patterns are finally decided.

3.6 Experimental set up

To test our methodology, a total of three experiments have been carried out:

- Baseline: we compare the classification results obtained with the plain dataset (dataset resulted from the process of missing values imputation) and the results obtained using the top 20 ranking frequent patterns. The number of patterns has been determined experimentally (see ablation analysis on Section 4.2). A 5-fold cross-validation (CV) is used 386 times (equivalent to the number of instances), emulating a bootstrapping process to estimate the mean and standard deviation. This is the main experiment that would support our contribution. Our hypothesis is that frequent patterns improve information for classification.

– Sensitive analysis: we analyse the contribution of the addition of frequent patterns in the dataset. To carry out the experiment, patterns were added incrementally according to their order in the ranking until they reached the first 100. A 5-fold CV is used.
– Hybrid data versus frequent patterns alone: We consider combining plain data with patterns (in their binary representation) in a hybrid dataset. We follow the same experimental procedure than in the previous experiments, but with both datasets.

4 Results

2514 patterns were obtained for the obese class and 1777 for the non-obese class. The symmetrical difference of both sets results in a total of 2569 frequent patterns.

4.1 Baseline experiment

Figure 3 shows the histograms obtained. Using the top 20 ranking patterns to train the model, we can expect to obtain an average recall of 29% and an accuracy of 82%, while using the plain dataset the expected values are 22% and 76,5% respectively. The results are conclusive, the model makes better predictions if it is trained with the 20 most discriminant patterns.

4.2 Sensitive analysis

Figure 4 (purple line) shows the results obtained in this experiment. Analyzing the recall of this second experiment, it can be deduced that using the first 20 patterns, the capacity of the model to detect instances of the obese (minority) class increases progressively, which confirms the discriminatory power of the first 20 patterns. Recall increases when considering 43 patterns, but at the expenses of diminishing significantly accuracy. Therefore, 20 patterns seems to be the best alternative.

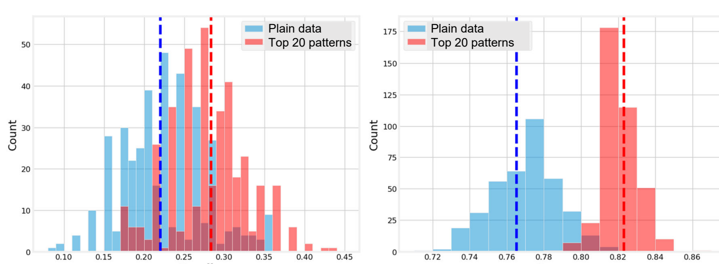

Fig. 3. Results: Value added by patterns.

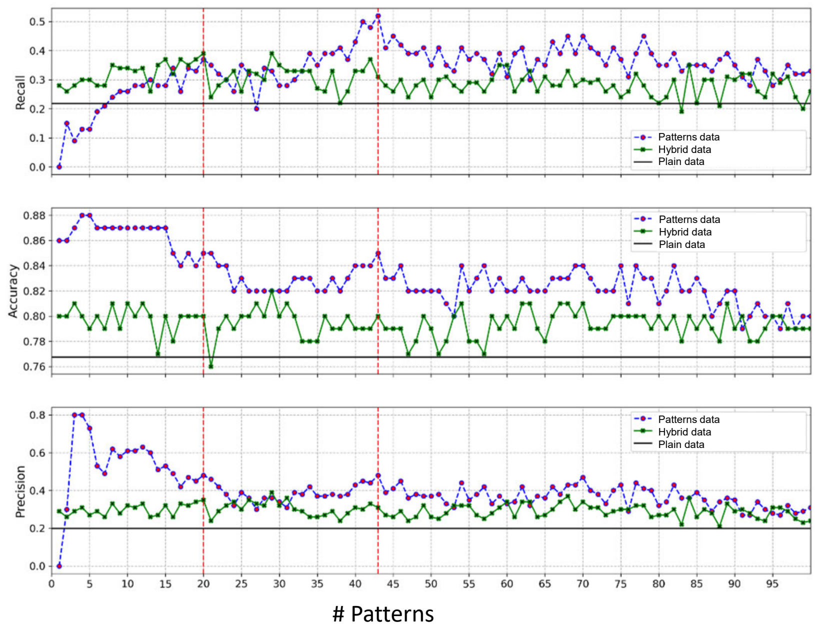

Fig. 4. Results: Sequential pattern selection.

4.3 Hybrid data versus frequent patterns alone

Figure 4 shows the comparison with the different datasets. We see how when adding patterns to the hybrid dataset, the recall increases considerably to a maximum value of 0.39 obtained by adding the first 20 patterns in the ranking. Accuracy and precision move in stable ranges between 0.78-0.82 and 0.2-0.4 respectively without showing any clear upward trend. Even so, all the values obtained in the experiment are higher than expected using only the plain dataset.

4.4 Discussion

Several conclusions can be drawn from the results obtained. The first is that adding frequent patterns considerably increases the quality of predictions compared to using only the original dataset. Regarding the results obtained for the accuracy measure (0.82), we can see that are close to the results of previous works: [15] 0.83 AUC; [4] 81.7 AUC for girls, and 76.1 for boys. On the other hand, we consider the recall measure due to the high imbalance dataset we have (14% of obese infants). Using only frequent patterns gives consistently better results in the 3 metrics accuracy, recall, and precision, except in the recall metric when using a quantity of patterns lower than the top 20 in the ranking, where the results are better if are used in conjunction with the original data (hybrid).

Analyzing the evolution of the values of the metrics obtained in the experiments carried out when we successively were adding new patterns, we can say that adding frequent discriminatory patterns causes the model to increase the number of predictions of the obese class (minority) causing an increase in the overall percentage of instances of this class classified as such (recall), which confirms the discriminatory power of the patterns.

From the 20 frequent patterns selected, it is noteworthy that only one pattern in the ranking (the second most relevant) belongs to the healthy class, which suggests that the values recorded in the healthy class have a higher redundancy load. Based on the ranking of frequent patterns, the medical team could also extract obesity prevention and risk metrics.

4.5 Limitations

The results obtained demonstrate the initial hypothesis, that is to say, that the frequent patterns provide discriminatory power. Even so they represent certain limitations in clinical practice, due to the low values obtained in the performance metrics. One possible direction under study is to separate girls from boys, as in [4], but previous results do not augur a high improvement. Another interesting research path is to deal with gradient values of variables between visits, instead of actual ones, and use sequence pattern mining algorithms.

5 Conclusions

Childhood obesity is a disease that continues to grow around the world in an alarming way. That's why many researchers are looking for key risk factors to prevent this health condition. This work proposes the use of frequent patterns learned from a dataset that contains longitudinal data of children and parents, in order to determine which are the most discriminated patterns to predict the obesity of a baby. The results of the experiments demonstrate the great discriminatory power of frequent patterns that capture the relationship between the data, versus other techniques.

However, the results obtained have limited values of accuracy, recall and precision, which implies continuing to reconsider other methodologies in the future. One of the lines of future work could be to explore alternative methods as sequence pattern mining algorithms.

Acknowledgements. We would like to thank Marina Rodriguez for her support in the initial dataset exploration. This work received joint funding from the European Regional Development Fund (ERDF), the Spanish Ministry of the Economy, Industry and Competitiveness (MINECO) and the Carlos III Research Institute, under grants no. PI23/00545 and PI22/00366. The work was carried out with support from the Generalitat de Catalunya 2021 SGR 01125.

References

1. Dong, W., Lee, E.W., et al.: Gasp: Graph-based approximate sequential pattern mining for electronic health records. Communications in Computer and Information Science **1450**, 50–60 (2021). https://doi.org/10.1007/978-3-030-85082-1_5/FIGURES/2
2. Gou, H., Song, H., Tian, Z., Liu, Y.: Prediction models for children/adolescents with obesity/overweight: A systematic review and meta-analysis. Preventive Medicine **179**, 107823 (2 2024). https://doi.org/10.1016/J.YPMED.2023.107823
3. Guillamet, G.H., Seguí, F.L., Vidal-Alaball, J., López, B.: Cauruler: Causal irredundant association rule miner for complex patient trajectory modelling. Computers in Biology and Medicine **155**, 106636 (3 2023). https://doi.org/10.1016/J.COMPBIOMED.2023.106636
4. Hammond, R., Athanasiadou, R., et. al: Predicting childhood obesity using electronic health records and publicly available data. PLOS ONE **14**, e0215571 (4 2019). https://doi.org/10.1371/JOURNAL.PONE.0215571
5. He, Z., Zhang, S., Wu, J.: Significance-based discriminative sequential pattern mining. Expert Systems with Applications **122**, 54–64 (5 2019). https://doi.org/10.1016/J.ESWA.2018.12.046
6. Lim, H., Lee, H., Kim, J.: A prediction model for childhood obesity risk using the machine learning method: a panel study on korean children. Scientific Reports 2023 13:1 **13**, 1–8 (6 2023). https://doi.org/10.1038/s41598-023-37171-4
7. Novitski, P., Cohen, C.M., Karasik, A., Hodik, G., Moskovitch, R.: Temporal patterns selection for all-cause mortality prediction in t2d with anns. Journal of Biomedical Informatics **134**, 104198 (10 2022). https://doi.org/10.1016/J.JBI.2022.104198
8. O'Connor, T.G., Williams, J., Blair, C., Gatzke-Kopp, L.M., Francis, L., Willoughby, M.T.: Predictors of developmental patterns of obesity in young children. Frontiers in Pediatrics **8**, 506691 (3 2020). https://doi.org/10.3389/FPED.2020.00109/BIBTEX
9. Pinaire, J., Chabert, E., Azé, J., Bringay, S., Landais, P.: Sequential pattern mining to predict medical in-hospital mortality from administrative data: Application to acute coronary syndrome. Journal of Healthcare Engineering **2021** (2021). https://doi.org/10.1155/2021/5531807
10. Soulet, A., Rioult, F.: Efficiently depth-first minimal pattern mining. LNAI **8443**, 28–39 (2014). https://doi.org/10.1007/978-3-319-06608-0_3/COVER
11. WHO: Obesity and overweight (2024), https://www.who.int/news-room/fact-sheets/detail/obesity-and-overweight
12. Wu, Y., Wang, Y., Li, Y., Zhu, X., Wu, X.: Top-k self-adaptive contrast sequential pattern mining. IEEE Transactions on Cybernetics **52**, 11819–11833 (11 2022). https://doi.org/10.1109/TCYB.2021.3082114
13. Wu, Y., Wang, Y., Liu, J., Yu, M., Liu, J., Li, Y.: Mining distinguishing subsequence patterns with nonoverlapping condition. Cluster Computing **22**, 5905–5917 (5 2019). https://doi.org/10.1007/S10586-017-1671-0/FIGURES/5
14. Zhao, Z., Anand, R., Wang, M.: Maximum relevance and minimum redundancy feature selection methods for a marketing machine learning platform. Proc.IEEE Int. Conf. on Data Science and Advanced Analytics (DSAA) pp. 442–452 (10 2019). https://doi.org/10.1109/DSAA.2019.00059
15. Ziauddeen, N., Roderick, P.J., Santorelli, G., Alwan, N.A.: Prediction of childhood overweight and obesity at age 10–11: findings from the studying lifecourse obesity

predictors and the born in bradford cohorts. International Journal of Obesity 2023 47:11 **47**, 1065–1073 (8 2023). https://doi.org/10.1038/s41366-023-01356-8

Appendix A. Variable description

Table 1 describes the variables used in this work.

Visit	Variable	Description	Who	Units
0	AgeM	Age of the mother	Mother	years
0	Height_EM	Height of the mother	Mother	cm
0	Smoking_preM	Smoking status of the mother before pregnancy	Mother	{0,1}
0	Alcohol_preM	Drinking status of the mother before pregnancy	Mother	{0,1}
0	BWM	Birth weight mother	Mother	Kg
0	Parity_EM	Is first child of the mother	Mother	{1,2}
0	BMIpre_Me	BMI of mother	Mother	kg/m2
0	AgeP	Age of the father	Father	years
0	HeightP	Height of the father	Father	cm
0	BMI_P	BMI of father	Father	kg/m2
0	SmokingP	Smoking status of the father	Father	{0,1}
0	AlcoholP	Drinking status of the father	Father	{0,1}
0	BWP	Birth weight father	Father	Kg
1	BMI_2_EM	BMI of mother second trimester	Mother	kg/m2
1	DBP_2_EM	Systolic bp of mother second trimester	Mother	mmhg
1	DBP_2_EM	Diastolic bp of mother second trimester	Mother	mmhg
2	BMI_3_EM	BMI of mother third trimester	Mother	kg/m2
2	DBP_3_EM	Systolic bp of mother third trimester	Mother	mmHg
2	DBP_3_EM	Diastolic bp of mother third trimester	Mother	mmHg
3	Obes_gest	Gestational obesity	Mother	{0,1}
3	Smoking_gestM	Smoking status of mother during pregnancy	Mother	{0,1}
3	Sex_nen	Gender of the baby	Baby	{1,2}
3	GA	Gestational age	Baby	weeks
3	Weight_B	Baby's weight at birth	Baby	gr
3	Height_B	Baby's length at birth	Baby	cm
3	Weight placenta	Weight of the placenta at birth	Baby	gr
9	Breastfeeding	Months of breastfeeding	Baby	month
4	Weight 2m_nen	Weight at 2 months	Baby	gr
4	Height 2m_nen	Height at 2 months	Baby	cm
5	Weight 4m_nen	Weight at 4 months	Baby	gr
5	Height 4m_nen	Height at 4 months	Baby	cm
6	Weight 6m_nen	Weight at 6 months	Baby	gr
6	Height 6m_nen	Height at 6 months	Baby	cm
7	Weight 12m_nen	Weight at 12 months	Baby	gr
7	Height 12m_nen	Height at 12 months	Baby	cm
9	OBSESITY	Obesity status of the child	Baby	{1,2}

Table 1. Dataset variables. Boolean values 0,1 meaning: 0=no, 1=yes. Values for Parity_EM: 1= primiparous; 2= multiparous(more than one child). Values for Sexe_nen: 1=girl; 2= boy. Values for OBSESITY (target): 1= normoweight ($BMISDS < 1$); 2= overweight or obesity ($BMISDS > 1$).

Appendix B. Missing values distribution

Figure 5 shows the number of instances by amount of missing values.

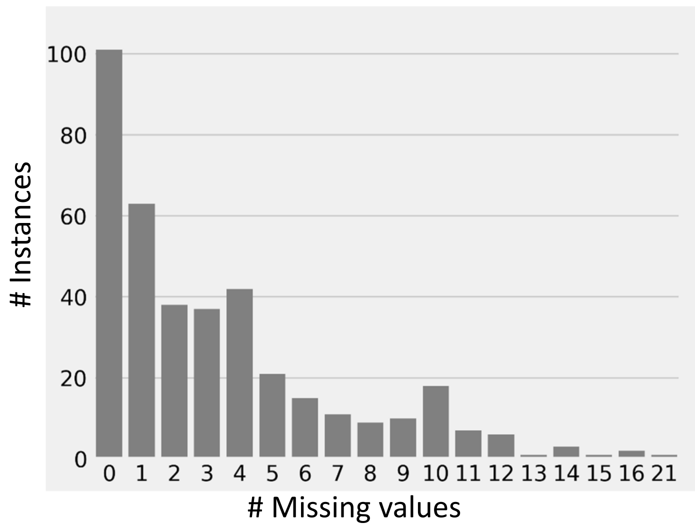

Fig. 5. Number of instances per amount of missing values.

Appendix C. Results on different classifiers

In this section you can see the average recall obtained by the different models XGBoost, Random Forest, SVClassifier, Logistic Regression and Decision Tree Classifier using *K-fold CV* with $K = 5$ using the first top 40 ranked patterns.

#	XGBoost	Random Forest	SVC	Logistic Regression	Decision Tree
1	0	0	0	0	0
2	**0,15**	**0,15**	0	0	**0,15**
3	0,07	0,05	0	**0,09**	0,1
4	0,15	**0,16**	0	0,07	0,13
5	**0,13**	**0,13**	0	0,09	**0,13**
6	**0,21**	**0,21**	0	0,11	0,19
7	**0,3**	0,26	0	0,06	0,21
8	**0,3**	0,24	0,05	0,15	0,24
9	**0,3**	**0,3**	0,04	0,15	0,26
10	**0,34**	0,23	0	0,19	0,26
11	**0,32**	0,21	0,05	0,11	0,28
12	0,3	0,23	0,09	0,17	**0,28**
13	**0,24**	0,17	0,09	0,19	0,3
14	0,26	0,13	0,09	0,22	**0,28**
15	0,22	0,13	0,09	0,22	**0,28**
16	0,19	0,15	0,17	0,26	**0,34**
17	0,24	0,17	0,17	0,24	**0,26**
18	0,23	0,13	0,17	0,22	**0,34**
19	0,24	0,15	0,19	0,22	**0,33**
20	0,28	0,19	0,2	0,3	**0,37**
21	0,25	0,19	0,26	0,26	**0,35**
22	0,24	0,16	**0,32**	0,26	**0,32**
23	0,24	0,21	**0,3**	0,24	**0,3**
24	0,26	0,15	**0,32**	0,26	0,26
25	0,22	0,19	0,3	0,26	**0,35**
26	0,24	0,19	0,28	0,26	**0,32**
27	0,28	0,15	**0,32**	0,26	0,20
28	0,28	0,15	0,32	0,28	**0,34**
29	0,28	0,13	0,28	0,28	**0,33**
30	0,28	0,15	**0,32**	0,28	0,28
31	0,24	0,07	**0,3**	0,28	0,28
32	0,28	0,07	**0,33**	0,28	0,3
33	0,3	0,15	0,32	0,3	**0,33**
34	0,28	0,15	0,34	0,28	**0,39**
35	0,3	0,15	0,33	0,26	**0,35**
36	0,24	0,15	0,31	0,24	**0,39**
37	0,24	0,15	0,29	0,3	**0,39**
38	0,26	0,11	0,29	0,31	**0,41**
39	0,2	0,13	0,35	0,31	**0,37**
40	0,24	0,09	0,37	0,33	**0,43**

Table 2. Mean recall for the different methods analyzed with 5-cv. In bold the maximum value of the row.

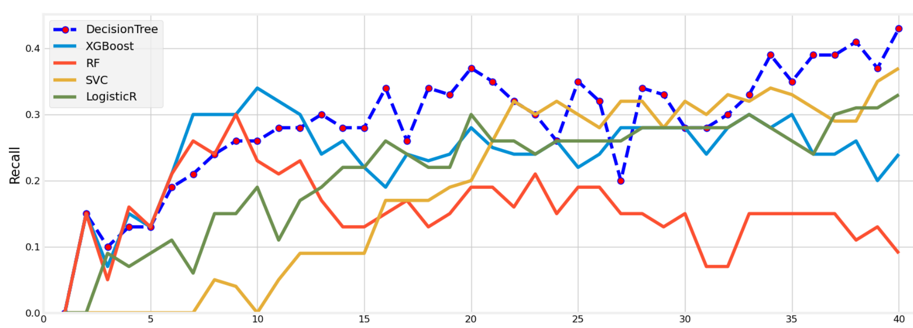

Fig. 6. Mean recall for the different methods analyzed with 5-cv.