

# Bagged Boosted Trees for Classification of Ecological Momentary Assessment Data

## Citation for published version (APA):

Spanakis, G., Weiss, G., & Roefs, A. (2016). Bagged Boosted Trees for Classification of Ecological Momentary Assessment Data. In *ECAI 2016 - 22nd European Conference on Artificial Intelligence, 29 August-2 September 2016, The Hague, The Netherlands - Including Prestigious Applications of Artificial Intelligence (PAIS 2016)* (pp. 1612-1613) <https://doi.org/10.3233/978-1-61499-672-9-1612>

## Document status and date:

Published: 01/01/2016

## DOI:

[10.3233/978-1-61499-672-9-1612](https://doi.org/10.3233/978-1-61499-672-9-1612)

## Document Version:

Publisher's PDF, also known as Version of record

## Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

## General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

[www.umlib.nl/taverne-license](http://www.umlib.nl/taverne-license)

## Take down policy

If you believe that this document breaches copyright please contact us at:

[repository@maastrichtuniversity.nl](mailto:repository@maastrichtuniversity.nl)

providing details and we will investigate your claim.

# Bagged Boosted Trees for Classification of Ecological Momentary Assessment Data

Gerasimos Spanakis and Gerhard Weiss<sup>1</sup> and Anne Roefs<sup>2</sup>

**Abstract.** Ecological Momentary Assessment (EMA) data is organized in multiple levels (per-subject, per-day, etc.) and this particular structure should be taken into account in machine learning algorithms used in EMA like decision trees and its variants. We propose a new algorithm called BBT (standing for Bagged Boosted Trees) that is enhanced by a over/under sampling method and can provide better estimates for the conditional class probability function. Experimental results on a real-world dataset show that BBT can benefit EMA data classification and performance.

## 1 Background & Motivation

This work focuses on classification trees and how their ensembles can be utilized in order to set up a prediction environment using Ecological Momentary Assessment (EMA) data from a real-world study. EMA [8] refers to a collection of methods used in many different disciplines by which a research subject repeatedly reports on specific variables measured close in time to experience and in the subject's natural environment (e.g. experiencing food craving is measured again and again on the same subject). EMA aims to minimize recall bias, maximize ecological validity and allow microscopic analysis of influence behavior in real-world contexts. EMA data has a different structure than normal data and account for several dependencies between them, since e.g. many samples belong to the same subject so they are expected to be correlated. However, most decision trees that deal with EMA data do not take these specificities into account.

Bagging involves having each tree in the ensemble vote with equal weight while boosting involves incrementally building an ensemble by training each new model instance to emphasize the training instances that previous models mis-classified. Major differences between bagging and boosting are that (a) boosting changes the distribution of training data based on the performance of classifiers created up to that point (bagging acts stochastically) and (b) bagging uses equal weight voting while boosting uses a function of the performance of a classifier as a weight for voting.

There are limited studies on combining bagging and boosting ([10], [3], [6] and [11]), however, none of these approaches have been applied to longitudinal data or take into account the EMA structure. Efforts to apply decision trees to EMA data have been attempted but they are mostly focusing on regression tasks ([7], [4], [2]) and on the other hand they do not use bagging or boosting for improving performance. Work in current paper aims at bridging this gap by combining

bagging and boosting with the longitudinal data structure.

## 2 BBT: The proposed algorithm

Let the training data be  $x_1, \dots, x_n$  and  $y_1, \dots, y_n$  where each  $x_i$  is a  $d$ -dimensional vector and  $y_i \in \{-1, 1\}$  is the associated observed class label. To justify generalization, it is usually assumed that training data as well as any test data are *iid* samples from some population of  $(x, y)$  pairs. Our goal is to as accurately predict  $y_i$  given  $x_i$ .

The first step to fit a BBT is to select the loss function, which in the case of a classification problem is based on the logistic regression loss. After some initial parameter selection (number of trees to be grown in sequence, shrinkage (or learning) rate, size of individual trees and fraction of the training data sampled) we grow BBT (say using  $M$  trees) on the training data using the following process and by growing single Boosted Trees (BT):

- Divide the data into  $B$  (typically 5 – 10) subsets and construct  $B$  training data sets each of which omits one of the  $B$  subsets (the ‘out-of-bag’ data). Each one of the  $B$  subsets is created by bootstrap sampling data points from the set of subjects ( $p = 1, \dots, P$ ). To create the learning set we introduce the strategy  $\mathcal{S}$  according to which one observation is drawn per subject. This strategy is based on a simple rationale: When only one observation per subject is selected, the probability that different observations are used for the training of different trees is increased, although the same subjects might be selected which further reduces similarity between trees. By this way, we manage to incorporate advantages of subject-based bootstrapping and observation-based bootstrapping into the final BBT ensemble. Also, this approach can be applied to unbalanced data points per subject.
- Grow  $B$  BT; one for each of the  $B$  training sets, based on the AdaBoost algorithm [1]: First let  $F_0(x_i) = 0$  for all  $x_i$  and initialize weights  $w_i = 1/d$  for  $i = 1, \dots, d$ . Then repeat the following for  $m = 1, \dots, M$  for each one of the  $B$  BT:
  - ★ Fit the decision tree  $g_m$  to the training data sample using weights  $w_i$  where  $g_m$  maps each  $x_i$  to -1 or 1.
  - ★ Compute:
    - the weighted error rate  $\epsilon_m = \sum_{i=1}^n w_i I\{y_i \neq g_m(x_i)\}$
    - half its log-odds and derive  $\alpha_m = \frac{1}{2} \log \frac{1-\epsilon_m}{\epsilon_m}$
  - ★ Let  $F_m = F_{m-1} + \alpha_m g_m$
  - ★ Replace the weights  $w_i$  with  $w_i = w_i e^{-\alpha_m g_m(x_i) y_i}$  and then renormalize by replacing each  $w_i$  by  $w_i / (\sum w_i)$ .
- Calculate the PE for each BT for tree sizes 1 to  $M$  from the corresponding out-of-bag data and pool across the  $B$  boosted trees. Predictions for new data are computed by first predicting each of

<sup>1</sup> Department of Data Science and Knowledge Engineering, Maastricht University, email: {jerry.spanakis,gerhard.weiss}@maastrichtuniversity.nl

<sup>2</sup> Faculty of Psychology and Neuroscience, Maastricht University, email: a.roefs@maastrichtuniversity.nl

the component trees and then aggregate the predictions (e.g., by averaging), like in bagging.

- The minimum PE estimates the optimum number of trees  $m^*$  for the BT. The estimated PE of the single BT obtained by cross-validation can thus also be used to estimate PE for the BBT. BBT thus require minimal additional computation beyond estimation of  $m^*$ .
- Reduce the number of trees for each BT to  $m^*$ .

For a classification problem, we use an estimate  $p_m(x)$  of the Conditional Class Probability Function (CCPF)  $p(x)$  that can be obtained from  $F_m$  through a logistic link function:

$$p_m(x) = p_m(y = 1|x) = \frac{1}{1 + \exp(-2F_m(x))} \quad (1)$$

Classifying at the 1/2 quantile of the CC PF works well for binary classification problems but in the case of EMA data, sometimes classification with unequal costs or, equivalently, classification at quantiles other than 1/2 is needed. Strategies about correctly computing the CC PF are considered [5] by over/under-sampling which convert a median classifier into a q-classifier.

### 3 Experiments

In order to illustrate the effect of BBT, we now apply this method to an EMA dataset obtained by a study designed by the authors [9]. The EMA study followed 100 participants over the course of 14 days using experience & event sampling questionnaires ending up with over 5000 data points containing information about participants' eating events, emotions, circumstances, locations, etc. (in total there are 9 variables) for several time moments during each day that they participated in the study. Each data point is used to predict whether the next data point (provided that they both occur on the same day) will be a healthy or an unhealthy eating moment. Figure 1 shows an example of how data points (belonging to user "pp5") are converted and combined in order to enable early prediction using a classification algorithm (class can be either "healthy" or "unhealthy"). Then the BBT algorithm can be applied.

user	Date/time	crv	negE	posE	sp_cr	time	week	clrc	loc	sp_est
pp5	26/01/15 23:57	LOW	NO	LOW	N	evening	NO	LowLevel	Home	N
pp5	27/01/15 09:32	LOW	NO	HIGH	N	morning	NO	LowLevel	Home	N
pp5	27/01/15 12:17	MID	NO	HIGH	H	noon-after	NO	ComputerRelated	Work	U
pp5	27/01/15 14:43	LOW	YES	MID	N	noon-after	NO	Work	Work	H

↓                      ↓                      ↓

user	crv	negE	posE	sp_cr	time	week	clrc	loc	sp_est	NextEating (y class)
pp5	LOW	NO	HIGH	N	morning	NO	LowLevel	Home	N	U
pp5	MID	NO	HIGH	H	noon-after	NO	ComputerRelated	Work	U	H

Figure 1: Data conversion example for early prediction

In the comparison between methods, BBT gave a PE of 23.3%, whereas the single classification tree (37.3%), bagged trees (28.9%), boosted trees (adaboost) (25.9%), random forests (26.8%) and B&B combine method [11] (26.2%) have higher PE than the BBT. Table 1 summarizes these results and also presents a series of experiments made to demonstrate the effectiveness of BBT when the number of different subjects ( $P$ ) involved in the dataset increases. For relatively small numbers of subjects (10 or 20) performance of BBT and Adaboost is comparable (although variance increases and the number of data samples is not large enough) but as  $P$  increases the performance of BBT is clearly better. Larger  $P$  means that there are more subjects in the dataset, thus the complexity of longitudinal structure increases

and it is imperative to take this into account when classifying longitudinal data. This is the reason that BBT performs better than all other algorithms as  $P$  increases. However, for small  $P$  the effect of different subjects is smaller and this is the reason that Adaboost performs slightly better than all other algorithms.

Table 1: Prediction Error (Variance) % for different algorithms and different numbers of subjects ( $P$ )

	P=10	P=20	P=50	P=100
SCT	25.1 (0.10)	27.4 (0.08)	30.9 (0.10)	37.3 (0.06)
Bagging	24.4 (0.22)	23.8 (0.08)	30.1 (0.08)	28.9 (0.06)
Boosting	<b>22.0 (0.12)</b>	22.7 (0.10)	27.0 (0.06)	25.9 (0.04)
Random Forest	23.7 (0.16)	24.5 (0.14)	27.2 (0.04)	26.8 (0.04)
B&B Combine	23.2 (0.08)	25.1 (0.06)	26.8 (0.08)	26.2 (0.02)
BBT	22.4 (0.14)	<b>21.9 (0.06)</b>	<b>24.2 (0.04)</b>	<b>23.3 (0.02)</b>

### 4 Discussion

In this paper a combination of bagging and boosting was presented: Bagged Boosted Trees (BBT). BBT have the advantage of being able to deal with multiple categorical data which raises a scalability issue when dealing with classic models (like generalized linear models) that are widely used in EMA studies. Moreover, BBT can tackle potential nonlinearities and interactions in the data, since these issues are handled through the combination of many different trees of different sizes. Experimental results of BBT on a real-world EMA dataset clearly show improvement with respect to accuracy in prediction compared to other decision tree algorithms. Further work involves the evaluation of the conditional class probability function (based on over/under sampling of data), as well as the application to other EMA datasets. Finally, adjustment of boosting in order to implement weights based on subjects (and not individual observations) is a direction with promising results.

### REFERENCES

- [1] Yoav Freund, Robert E Schapire, et al., 'Experiments with a new boosting algorithm', in *ICML*, volume 96, pp. 148–156, (1996).
- [2] Wei Fu and Jeffrey S Simonoff, 'Unbiased regression trees for longitudinal and clustered data', *Computational Statistics & Data Analysis*, **88**, 53–74, (2015).
- [3] S Kotsiantis and P Pintelas, 'Combining bagging and boosting', *International Journal of Computational Intelligence*, **1**(4), 324–333, (2004).
- [4] Wei-Yin Loh, Wei Zheng, et al., 'Regression trees for longitudinal and multiresponse data', *The Annals of Applied Statistics*, **7**(1), 495–522, (2013).
- [5] David Mease, Abraham J. Wyner, and Andreas Buja, 'Boosted classification trees and class probability/quantile estimation', *J. Mach. Learn. Res.*, **8**, 409–439, (May 2007).
- [6] J. Ross Quinlan, *C4.5: Programs for Machine Learning*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1993.
- [7] Rebecca J Sela and Jeffrey S Simonoff, 'RE-EM trees: a data mining approach for longitudinal and clustered data', *Machine learning*, **86**(2), 169–207, (2012).
- [8] Saul Shiffman, Arthur A Stone, and Michael R Hufford, 'Ecological momentary assessment', *Annu. Rev. Clin. Psychol.*, **4**, 1–32, (2008).
- [9] Gerasimos Spanakis, Gerhard Weiss, Bastiaan Boh, and Anne Roefs, 'Network analysis of ecological momentary assessment data for monitoring and understanding eating behavior', in *Smart Health - International Conference, ICSH 2015, Phoenix, AZ, USA, November 17-18, 2015*, pp. 43–54, (2015).
- [10] Yuk Lai Suen, Prem Melville, and Raymond J Mooney, 'Combining bias and variance reduction techniques for regression trees', in *Machine Learning: ECML 2005*, 741–749, Springer, (2005).
- [11] Jianjun Xie, Viktoria Rojkova, Siddharth Pal, and Stephen Coggeshall, 'A combination of boosting and bagging for kdd cup 2009-fast scoring on a large database.', in *KDD Cup*, pp. 35–43, (2009).