

Pycobra: A Python Toolbox for Ensemble Learning and Visualisation

Benjamin Guedj

Modal project-team

Inria Lille - Nord Europe

France

BENJAMIN.GUEDJ@INRIA.FR

Bhargav Srinivasa Desikan

Modal project-team

Inria Lille - Nord Europe

France

BHARGAV.SRINIVASA-DESIKAN@INRIA.FR

Abstract

We introduce `pycobra`, a Python library devoted to ensemble learning (regression and classification) and visualisation. Its main assets are the implementation of several ensemble learning algorithms, a flexible and generic interface to compare and blend any existing machine learning algorithm available in Python libraries (as long as a `predict` method is given), and visualisation tools such as Voronoi tessellations. `pycobra` is fully `scikit-learn` compatible and is released under the MIT open-source license. `pycobra` can be downloaded from the Python Package Index (PyPi) and Machine Learning Open Source Software (MLOSS). The current version (along with Jupyter notebooks, extensive documentation, and continuous integration tests) is available at <https://github.com/bhargavvader/pycobra>.

Keywords: Ensemble methods, Machine learning, Voronoi tessellation, Python, Open Source Software.

1. Introduction

Combined statistical procedures, also known as ensemble or aggregation methods, are very popular in Machine Learning competitions – the celebrated Netflix Challenge (as are repeatedly Kaggle competitions) was won by an ensemble technique (see for example [Bell and Koren, 2007](#), for a discussion). In a nutshell, ensemble methods combine different predictors (each trained on the same dataset) to improve performance. `scikit-learn` offers implementations of some ensemble methods but it does not ship with a specific analysis toolbox for aggregation. `pycobra` attempts to fill this gap by providing to the `scikit-learn` environment a toolbox to analyse and visualise the different preliminary predictors (also called machines hereafter), in addition to providing pythonic implementations of several ensemble algorithms (the COBRA algorithm introduced by [Biau et al., 2016](#); the Exponential Weighted Aggregate (EWA) introduced by [Vovk, 1990](#); a COBRA-flavored majority vote inspired by [Mojirsheibani, 1999](#)). All these algorithms are supported by oracle bounds whose main message is that their performance outperforms the performance of the best preliminary predictor, up to a remainder term which decays to zero.

COBRA (standing for COmBined Regression Alternative) is a nonlinear ensemble method designed for regression problems. To predict the response for a new data point, COBRA operates in two steps. First, it retains data points for which the prediction made by preliminary machines is close (in the Euclidean sense) to the prediction made for the new point. The final prediction is then formed by averaging responses over the retained points' indices. COBRA outputs a predictor which outperforms any combination of the preliminary predictors (as shown by see Theorem 2.1 in [Biau et al., 2016](#)). We describe the COBRA algorithm as implemented in `pycobra` in [Algorithm 1](#). Exponential weights have been used for a long time in statistical learning theory ([Vovk, 1990](#)). The EWA algorithm implemented in `Pycobra` is inspired by the description of [Dalalyan and Tsybakov \(2007\)](#). EWA amounts to forming an exponentially weighted average of preliminary predictors, where the weights include a measure of each predictor's performance. COBRA has been inspired by the work of [Mojirsheibani \(1999\)](#) in classification, and therefore `pycobra` includes a classification version of COBRA called `ClassifierCobra`. The consensus step is similar to what is done in regression but the final prediction is delivered by a majority vote among the labels of the retained data points.

`pycobra` allows the user to gauge the performance of the preliminary predictors used in the aggregation, with built-in methods to easily plot boxplots and QQ-plots. A salient feature of `pycobra` is using Voronoi tessellations for generic visualisation. By implementing ensemble algorithms which were not represented in the Python machine learning community, and providing a variety of tools to visualise and analyse their behavior and performance, we present to the machine learning open source community a toolbox designed for ensemble learning and visualisation.

2. The `pycobra` library

Our toolbox is written in Python and uses `NumPy` ([Walt et al., 2011](#)) and `scikit-learn` ([Pedregosa et al., 2011](#)) for computation and linear algebra operations. `Matplotlib` ([Hunter, 2007](#)) is used for visualisation purposes, and `SciPy` ([Jones et al., 2001](#)) is used to help create Voronoi tessellations. Tutorials and examples are created using Jupyter IPython notebooks ([Pérez and Granger, 2007](#)). Currently `pycobra` supports any machine learning algorithm with a `predict` method, casting our procedure into a very flexible and generic framework. By default, `pycobra` relies on `scikit-learn` implementations of Lasso, Random Forest, Decision Trees and Ridge regression for the EWA and COBRA implementations. As for `ClassifierCobra`, `scikit-learn` implementations of SVM classifier, KNN, Decision Tree classifier, and SGD classifier are used. `pycobra` itself and all software it relies on is open source, with no dependence on proprietary software. [Algorithm 1](#) presents the pseudo-code of the COBRA implementation. All the `pycobra` estimators are `scikit-learn` compatible and can be used as part of the existing `scikit-learn` ecosystem, such as `GridSearchCV` and `Pipeline`.

While hyperparameter initialisation is systematically done using `scikit-learn`'s `GridSearchCV`, `pycobra`'s `diagnostics` class allows us to compare between different combinations of the constituent predictors and data-splitting, among other basic parameters.

Algorithm 1: The original COBRA algorithm from [Biau et al. \(2016\)](#).

```

Data: input vector  $\mathbf{X}$ , epsilon, alpha, basic-machines, training-set-responses
Result: prediction  $\mathbf{Y}$ 
for machine  $j$  in basic-machines do
  set = [];
  pred =  $r_j(\mathbf{X})$ ; # where  $r_j(x)$  denotes the prediction made by machine  $j$  at point  $x$ ;
  for response in training-set-responses do
    | if  $|\text{point} - \text{pred}| \leq \text{epsilon}$ , collect response in machine  $M$ 's set ;
  end
end
array = [];
for point in training-set do
  | if at least alpha machines out of  $M$  have collected point in their set, store point in array ;
end
result = average(array) ;

```

The `visualisation` class allows the user to compare all the machines used to create aggregates, as well as visualise the results, for all `pycobra` estimators. `pycobra` ships with a [notebook on visualisation](#) to illustrate this.

QQ-plots and boxplots. Once all the basic machines are set up, the user can easily compare their performance with boxplots and QQ-plots. All the plotting details are handled by both the `diagnostics` and `visualisation` classes. An example is given in [Figure 1](#).

Visualisation through Voronoi tessellations. Voronoi tessellations can be used to visualise the selection of optimal machines for COBRA, as well as visualising outputs from clustering algorithms, as illustrated by the [tutorial notebook](#). An example is reproduced [Figure 2](#).

3. Project Focus

Community-based development. We intend `pycobra` to be an ongoing collaborative project. To that matter, it is referenced on [Python Package Index \(PyPi\)](#) and [Machine Learning Open Source Software](#). Moreover, `pycobra` is under active development and available on [GitHub](#) to enforce collaborative programming, issue tracking, code integration, and idea discussions.




Figure 1: Assessing the performance of regression machines and COBRA.

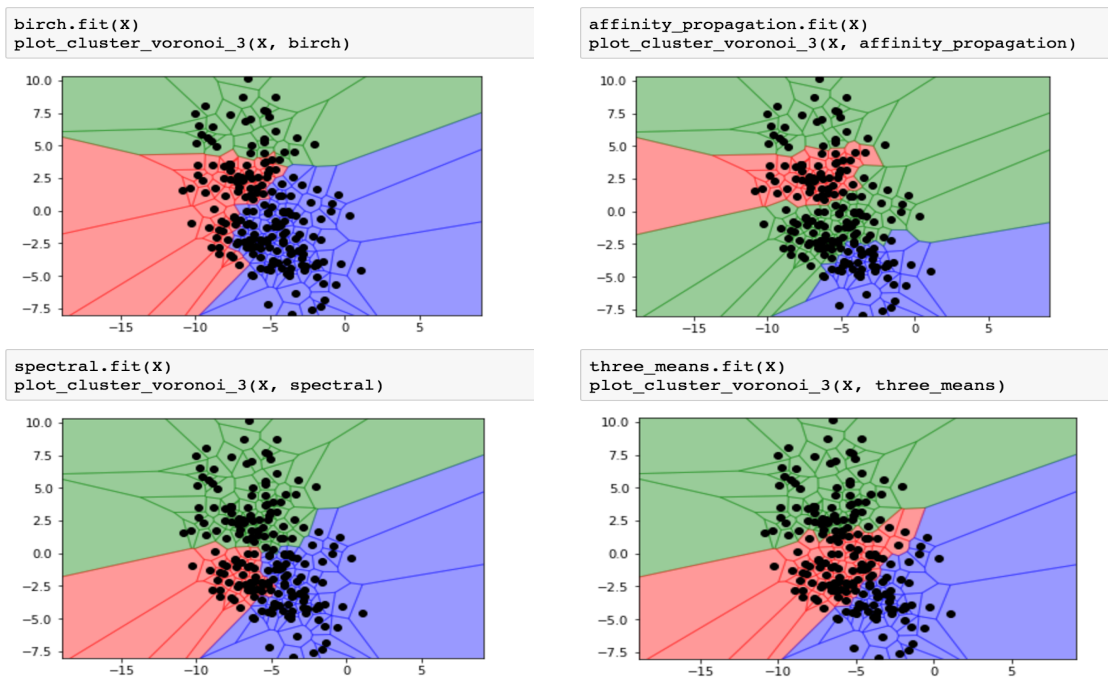


Figure 2: Visualising clustering through Voronoi tessellations.

Documentation and Jupyter Notebooks. A consistent API documentation is provided, along with an additional installation guide and examples. The documentation website is [here](#). The [notebooks directory](#) contains Jupyter notebooks which serve as both a documentation tool and a tutorial resource. These notebooks cover use-cases of `Pycobra`, from solving regression and classification problems to using Voronoi tessellations.

Ease of use and quality assurance. Ensemble learning with `pycobra` is as simple as loading trained `scikit-learn` machines – or any machine which has a `predict` method. Visualising involves little to no parameters, and after loading machines it is straightforward to analyse their performance on a particular dataset. In order to ensure code quality, a set of unit tests is provided for all classes in `pycobra`, and continuous integration via [Travis CI](#) ensures all commits are tested. All the code is PEP8 compliant, keeping the code easy to read and contribute to.

4. Conclusion and Future Work

The future of `pycobra` would be to grow its user base by adding new ways to add predictors, as well as further implement ensemble learning techniques and visualisation tools. Statistical aggregation and ensemble learning are an important part of the machine learning literature and are widely used by practitioners, yet it seems under-represented in the machine learning open source community. By creating `pycobra` and releasing it to the community, we intend to enrich the existing ecosystem with ensemble algorithms, a generic toolbox for ensemble learning on-the-go, along with analysis and visualisation tools.

References

- Robert M. Bell and Yehuda Koren. Lessons from the Netflix prize challenge. *ACM SIGKDD Explorations Newsletter*, 9(2):75–79, 2007.
- G erard Biau, Aur elie Fischer, Benjamin Guedj, and James D. Malley. COBRA: A combined regression strategy. *Journal of Multivariate Analysis*, 146:18–28, 2016.
- Arnak Dalalyan and Alexandre Tsybakov. Aggregation by exponential weighting and sharp oracle inequalities. In *Proceedings of the Twentieth Conference on Learning theory, COLT 2007*, pages 97–111. Springer, 2007.
- John D. Hunter. Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.
- Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001. URL <http://www.scipy.org/>.
- Majid Mojrirsheibani. Combining classifiers via discretization. *Journal of the American Statistical Association*, 94(446):600–609, 1999.
- Fabian Pedregosa, Ga el Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12(Oct):2825–2830, 2011.
- Fernando P erez and Brian E. Granger. IPython: a system for interactive scientific computing. *Computing in Science & Engineering*, 9(3):21–29, May 2007. ISSN 1521-9615. doi: 10.1109/MCSE.2007.53. URL <http://ipython.org>.
- Volodimir G. Vovk. Aggregating strategies. In *Proceedings of the Third Annual Workshop on Computational Learning theory, COLT 1990*, pages 371–386. Morgan Kaufmann Publishers Inc., 1990.
- St efan van der Walt, S. Chris Colbert, and Ga el Varoquaux. The NumPy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22–30, 2011.