



# pystacked: Stacking generalization and machine learning in Stata

Achim Ahrens  
ETH Zürich  
Zürich, Switzerland  
achim.ahrens@gess.ethz.ch

Christian B. Hansen  
University of Chicago  
Chicago, IL  
christian.hansen@chicagobooth.edu

Mark E. Schaffer  
Heriot-Watt University  
Edinburgh, U.K.  
m.e.schaffer@hw.ac.uk

**Abstract.** The `pystacked` command implements stacked generalization (Wolpert, 1992, *Neural Networks* 5: 241–259) for regression and binary classification via Python’s `scikit-learn`. Stacking combines multiple supervised machine learners—the “base” or “level-0” learners—into one learner. The currently supported base learners include regularized regression, random forest, gradient boosted trees, support vector machines, and feed-forward neural nets (multilayer perceptron). `pystacked` can also be used as a “regular” machine learning program to fit one base learner and thus provides an easy-to-use application programming interface for `scikit-learn`’s machine learning algorithms.

**Keywords:** `st0731`, `pystacked`, machine learning, stacked generalization, model averaging, Python, `sci-kit learn`

## 1 Introduction

When faced with a new prediction or classification task, it is rarely obvious which machine learning algorithm is best suited. A common approach is to evaluate the performance of a set of machine learners on a holdout partition of the data or via cross-validation and then select the machine learner that minimizes a chosen loss metric. However, this approach is incomplete because combining multiple learners into one final prediction might lead to superior performance compared with each individual learner. This possibility motivates stacked generalization, or simply “stacking” (see Wolpert [1992] and Breiman [1996]). Stacking is a form of model averaging. Theoretical results in van der Laan, Polley, and Hubbard (2007) support the use of stacking because it performs asymptotically at least as well as the best-performing individual learner as long as the number of base learners is not too large.

In this article, we introduce `pystacked`, a command for stacking regression and binary classification in Stata. `pystacked` allows users to fit multiple machine learning algorithms via Python’s `scikit-learn` (Pedregosa et al. 2011; Buitinck et al. 2013) and combine these into one final prediction as a weighted average of individual pre-

dictions. `pystacked` adds to the growing number of programs for machine learning in Stata, including `lassopack` for regularized regression (Ahrens, Hansen, and Schaffer 2020), `rforest` for random forests (Schonlau and Zou 2020), and `svm` for support vector machines (Guenther and Schonlau 2016, 2018). Similarly to `pystacked`, Cerulli (2022) and Droste (2022) provide an interface to `scikit-learn` in Stata. `mlrtime` allows Stata users to make use of R’s `parsnip` machine learning library (Huntington-Klein 2021). `pystacked` differs from these in that it is, to our knowledge, the first to make stacking available to Stata users. Furthermore, `pystacked` can also be used to fit one machine learner and thus provides an easy-to-use and versatile application programming interface to `scikit-learn`’s machine learning algorithms.

Stacking is widely used in applied predictive modeling in many disciplines—for example, predicting mortality (Hwangbo et al. 2022), bankruptcy filings (Liang et al. 2020; Fedorova et al. 2022), or temperatures (Hooker, Duveiller, and Cescatti 2018). However, the use of stacking as a method and `pystacked` as a program is not restricted only to pure prediction or classification tasks. A growing literature exploits machine learning to facilitate causal inference (for an overview, see Athey and Imbens [2019]). Indeed, a motivation for writing `pystacked` is that it can be used in combination with `ddml` (Ahrens et al. Forthcoming, 2023), which implements the double-debiased machine learning (DDML) methodology of Chernozhukov et al. (2018). DDML uses cross-fitting, a form of iterative sample splitting, which allows leveraging of a wide class of supervised machine learners, including stacking, for the estimation of causal parameters. For instance, in the context of DDML, stacking can be used to estimate the conditional expectation of an outcome with respect to confounders or propensity scores.

We stress that `pystacked` relies on Python’s `scikit-learn` (version 0.24 or higher) and the ongoing work of the `scikit-learn` contributors. Thus, `pystacked` relies on Stata’s Python integration, which was introduced in Stata 16.0. We kindly ask users to cite `scikit-learn` along with this article when using `pystacked`. Throughout, we refer to version 0.7.1 of `pystacked`, which is the latest version at the time of writing. Please check for updates to `pystacked` regularly, and consult the help file to be informed about new features. The `pystacked` help file includes information on how to install a recent Python version and set up Stata’s Python integration.

Section 2 introduces stacking. Section 3 presents the main features of the `pystacked` program. Section 4 demonstrates the use of the program using examples.

## 2 Methodology

In this section, we briefly summarize the stacking approach for regression and binary classification tasks. Van der Laan and Rose (2011) provides a book-length treatment; for a concise textbook treatment, see Hastie, Tibshirani, and Friedman (2009).

We first focus on stacking for regression problems where the aim is to predict the continuous outcome  $y_i$  using predictors  $\mathbf{x}_i$ . The idea of stacking is to combine a set of “base” (or “level-0”) learners using a “final” (or “level-1”) estimator. It is advisable

to include a relatively large and diverse set of base learners to capture different types of patterns in the data. The same algorithm can also be included more than once using different tuning or hypertuning parameters. Typical choices for base learners are regularized regression or ensemble methods, such as random forests or gradient boosting.

In the first step of stacking, we obtain cross-validated predicted values  $\widehat{y}_{(j),i}^{-k(i)}$  for each base learner  $j$  and observation  $i$ . The superscript  $-k(i)$  indicates that we form the cross-validated predicted value for observation  $i$  by fitting the learner to all folds except fold  $k(i)$ , which is the fold that includes observation  $i$ . The use of cross-validation is necessary because stacking would otherwise give more weight to base learners that suffer from overfitting. The second step is to fit a final learner using the observed  $y_i$  as the outcome and the cross-validated predicted values  $\widehat{y}_{(1),i}^{-k(i)}, \dots, \widehat{y}_{(J),i}^{-k(i)}$  as predictors. A typical choice for the final learner is constrained least squares, which enforces the stacking weights to be nonnegative and sum to one. This restriction facilitates the interpretation of stacking as a weighted average of base learners and may lead to better performance (Breiman 1996; Hastie, Tibshirani, and Friedman 2009). Algorithm 1 summarizes the stacking algorithm for regression problems as it is implemented in `pystacked`.

□ **Algorithm 1: Stacking regression**

1. *Cross-validation:*

- a. Split the sample  $I = \{1, \dots, n\}$  randomly into  $K$  partitions of approximately equal size. These partitions are referred to as “folds”. Denote the set of observations in fold  $k = 1, \dots, K$  as  $I_k$  and its complement as  $I_k^c$  such that  $I_k^c = I \setminus I_k$ .  $I_k$  constitutes the step- $k$  validation set and  $I_k^c$  the step- $k$  training sample.
- b. For each fold  $k = 1, \dots, K$  and each base learner  $j = 1, \dots, J$ , fit the supervised machine learner  $j$  to the training data  $I_k^c$  and obtain out-of-sample predicted values  $\widehat{y}_{(j),i}^{-k}$  for  $i \in I_k$ .

2. *Final learner:* Fit the final learner to the full sample. The default choice is non-negative least squares with the additional constraint that coefficients sum to one:

$$\min_{w_1, \dots, w_J} \sum_{i=1}^n \left( y_i - \sum_{j=1}^J w_j \widehat{y}_{(j),i}^{-k(i)} \right)^2 \quad \text{s.t.} \quad w_j \geq 0, \quad \sum_{j=1}^J w_j = 1$$

The stacking predicted values are defined as  $\widehat{y}_i^* = \sum_j \widehat{w}_j \widehat{y}_{(j),i}$  where  $\widehat{w}_j$  is the estimated stacking weight corresponding to learner  $j$  and  $\widehat{y}_{(j),i}$  are the predicted values from refitting learner  $j$  on the full sample  $I$ .

□

It is instructive to compare step 2 with the classical “winner-takes-all” approach that selects one base learner as the one that exhibits the lowest cross-validated loss. Stacking, in contrast, may assign nonzero weights to multiple base learners, thus combining their strengths to produce an overall predictor that can be better than any of the individual

base learners. Other choices for the final learner are possible. In addition to the default final learner, `pystacked` supports, among others, nonnegative least squares without the constraint  $\sum w_j = 1$ , ridge regression, the aforementioned “winner-takes-all” approach that selects the base learner with the smallest cross-validated mean squared error, and unconstrained least squares.

## 2.1 Stacking classification

Stacking can be applied similarly to classification problems. `pystacked` supports stacking for binary classification problems where the outcome  $y_i$  takes the values 0 or 1. The main difference from stacking regression is that  $\widehat{y}_{(j),i}^k$  represents cross-validated predicted probabilities.

## 2.2 Base learners

In the next paragraphs, we briefly describe the base learners supported by `pystacked` and highlight central tuning parameters. We repeat that each of the machine learners discussed below can be fit using `pystacked` as a regular stand-alone machine learner without the stacking layer.

Note that it is beyond the scope of the article to describe each learner in detail. Familiarity with linear regression, logistic regression, and classification and regression trees is assumed. We recommend consulting machine learning textbooks, for example, Hastie, Tibshirani, and Friedman (2009), for more detailed discussion.

**Regularized regression** imposes a penalty on the size of coefficients to control overfitting. The lasso penalizes the absolute size of coefficients, whereas the ridge penalizes the sum of squared coefficients. Both methods shrink coefficients toward zero, but only the lasso yields sparse solutions where some coefficient estimates are set to exactly zero. The elastic net combines lasso- and ridge-type penalties. For classification tasks with a binary outcome, logistic versions of lasso, ridge, and elastic net are available.

The severity of the penalty is most commonly chosen by cross-validation. For lasso only, `pystacked` also supports selecting the penalty by the Akaike information criterion (AIC) or Bayesian information criterion (BIC). The use of AIC or BIC has the advantage that it is computationally less intensive than cross-validation. Ahrens, Hansen, and Schaffer (2020) compare the two approaches.

**Random forests** rely on fitting many regression or decision trees on bootstrap samples of the data. The random forest prediction is obtained as the average across individual trees. A crucial aspect of random forests is that, at each split when growing a tree, one may consider only a random subset of predictors. This restriction aims at decorrelating the individual trees. Central tuning parameters are the number of trees (`n_estimators()`), the maximum depth of individual trees (`max_depth()`), the minimum number of observations per leaf (`min_samples_leaf()`), the number of features to be considered at each split (`max_features()`), and the size of the bootstrap samples (`max_samples()`).

**Gradient boosted trees** also rely on fitting many trees. In contrast to random forests, these trees are fit sequentially to the residuals from the current model. The learning rate determines how much the latest tree contributes to the overall model. Individual trees are usually fit to the whole sample, although subsampling is possible. In addition to tuning parameters relating to the trees, the learning rate (`learnings_rate()`) and the number of trees (`n_estimators()`) are the most important tuning parameters.

**Support vector machines.** Support vector classifiers span a hyperplane that separates observations by their outcome class. The hyperplane is chosen to maximize the distance (`margin`) to correctly classified observations while allowing for some classification errors. The tuning parameter  $C$  (`C()`) controls the frequency and degree of classification mistakes. The hyperplane can be either linear or fit using kernels. The support vector machine algorithm can also be adapted for regression tasks. To this end, the hyperplane is constructed to include as many observations as possible in a tube of size  $2\epsilon$  around the hyperplane. Central tuning parameters for regression are  $\epsilon$  (`epsilon()`) and  $C$  (`C()`), which determine the cost of observations outside of the tube.

**Feed-forward neural networks** consist of hidden layers that link the predictors (referred to as “input layers”) to the outcome. Each hidden layer is composed of multiple units (nodes) that pass signals to the next layer using an activation function. Central tuning parameters are the choice of the activation function (`activation()`) and the number and size of hidden layers (`hidden_layer_sizes()`). Further tuning choices relate to stochastic gradient descent algorithms, which are typically used to fit neural networks. The default solver is Adam (Kingma and Ba 2017). The option `early_stopping` can be used to set aside a random fraction of the data for validation. The optimization algorithm stops if there is no improvement in performance over a prespecified number of iterations (see related options `n_iter_no_change()` and `tol()`).

### 3 Program

This section introduces the program `pystacked` and its main features. `pystacked` offers two alternative syntaxes to the user (see sections 3.1 and 3.2). The two syntaxes offer the same functionality and are included to accommodate different user preferences.

Sections 3.3 and 3.4 list postestimation commands and general options, respectively. Section 3.5 discusses supported base learners. Section 3.6 is a note on learner-specific predictors. Section 3.7 explains the pipeline feature.

### 3.1 Syntax 1

The first syntax uses `methods()` to select base learners, where *string* is a list of base learners. Options are passed on to base learners via `cmdopt1()`, `cmdopt2()`, etc. That is, base learners can be specified, and options are passed on in the order in which they appear in `methods()` (see section 3.5). Likewise, the `pipe*()` option can be used for preprocessing predictors within Python on the fly, where *\** is a placeholder for 1, 2, etc. (see section 3.7). Finally, `xvars*()` allows the specification of a learner-specific variable list of predictors.

```
pystacked depvar predictors [if] [in] [, methods(string) cmdopt1(string)
  [cmdopt2(string) [...]] pipe1(string) [pipe2(string) [...]]
  xvars1(predictors) [xvars2(predictors) [...]] general_options]
```

*general\_option* are discussed in section 3.4 below.

### 3.2 Syntax 2

In the second syntax, base learners are added before the comma using `method()`, along with further learner-specific settings, and separated by `||`.

```
pystacked depvar [indepvars] || method(string) options(string)
  pipeline(string) xvars(predictors) [|| method(string) options(string)
  pipeline(string) xvars(predictors) || ... ||] [if] [in] [, general_options]
```

### 3.3 Postestimation commands

**Predicted values.** To get predicted values:

```
predict type newname [if] [in] [, xb pr]
```

To get fitted values for each base learner:

```
predict type stub [if] [in] [, basexb cvalid]
```

Predicted values (in and out of sample) are calculated when `pystacked` is run and stored in Python memory. `predict` pulls the predicted values from Python memory and saves them in Stata memory. This storage structure means that no changes on the data in Stata memory should be made between the `pystacked` call and the `predict` call. If changes to the dataset are made, `predict` will return an error.

The option `basexb` returns predicted values for each base learner. By default, the predicted values from refitting base learners on the full estimation sample are returned. If combined with `cvalid`, the cross-fitted predicted values are returned for each base learner.

**Tables.** After estimation, `pystacked` can report a table of in-sample (both cross-validated and full-sample) and, optionally, out-of-sample (or holdout sample) performances for both the stacking regression and the base learners. For regression problems, the table reports the root mean squared prediction error (RMSPE). For classification problems, a confusion matrix is reported. The default holdout sample used for out-of-sample performance with the holdout option is all observations not included in the estimation. Alternatively, the user can specify the holdout sample explicitly using `holdout()`. The table can be requested after estimation as a replay command or as part of the `pystacked` estimation.

```
pystacked [ , table [ (options) ] holdout [ (varname) ] ]
```

**Graphs.** `pystacked` can also create graphs of in-sample and, optionally, out-of-sample performance for both the stacking regression and the base learners. For regression problems, the graphs compare predicted and actual values of `depvar`. For classification problems, the default is to generate receiver operator characteristic (ROC) curves. Optionally, histograms of predicted probabilities are reported. As with the `table` option, the default holdout sample used for out-of-sample performance is all observations not included in the estimation, but the user can instead specify the holdout sample explicitly. The table can be requested after estimation or as part of the `pystacked` estimation command. The `graph` option on its own reports the graphs using `pystacked`'s default settings. Because graphs are produced using Stata's `twoway`, `roctab`, and `histogram` commands, the user can control either the combined graph (`graph()`) or the individual learner graphs (`lgraph()`) by passing options to these commands.

```
pystacked [ , graph [ (options) ] lgraph [ (options) ] histogram
  holdout [ (varname) ] ]
```

### 3.4 General options

A full list of general options is provided in the `pystacked` help file. We list only the most important general options here:

`type(string)` allows `regress` for regression problems or `classify` for classification problems. The default is regression.

`finalest(string)` selects the final estimator used to combine base learners. The default is nonnegative least squares without an intercept and the additional constraint that weights sum to 1 (`nnls1`). Alternatives are `nnls0` (nonnegative least squares without an intercept and without the sum-to-one constraint), `singlebest` (use the base learner with the minimum mean squared error), `ls1` (least squares without an intercept and with the sum-to-one constraint), `ols` (ordinary least squares), or `ridge` for ridge regression or logistic ridge, which is the `scikit-learn` default.

`njobs(integer)` sets the number of jobs for parallel computing; the default is `njobs(0)` (no parallelization). `njobs(-1)` uses all available CPUs, and `njobs(-2)` uses all CPUs minus 1.

`backend(string)` is the backend that is used for parallelization. The default is `backend(threading)`.

`fold(integer)` specifies the number of folds used for cross-validation. The default is `fold(5)`.

`foldvar(varname)` is the integer fold variable for cross-validation.

`bfold(integer)` sets the number of folds used for base learners that use cross-validation (for example, cross-validated lasso). The default is `bfold(5)`.

`sparse` converts the predictor matrix to a sparse matrix. This conversion will lead to speed improvements only if the predictor matrix is sufficiently sparse. Not all learners support sparse matrices, and not all learners will benefit from sparse matrices in the same way. One can also use the sparse pipeline to use sparse matrices for some learners but not for others.

`pyseed(integer)` sets the Python seed. Note that because `pystacked` uses Python, we also need to set the Python seed to ensure replicability. There are three options:

1. `pyseed(-1)` draws a number between 0 and  $10^8$  in Stata that is then used as a Python seed; this is `pystacked`'s default behavior. This way, one needs to deal only with the Stata seed. For example, `set seed 42` is sufficient because the Python seed is generated automatically.
2. Setting `pyseed(x)` with any positive integer  $x$  allows direct control of the Python seed.
3. `pyseed(0)` sets the seed to `None` in Python.



`printopt` prints the default options for specified learners. Only one learner can be specified. This is for information only; no estimation is done. See section 3.5 for examples.

`showoptions` prints the options passed on to Python.

`showpymessages` prints Python messages.

`voting` selects voting regression or classification, which uses prespecified weights. By default, voting regression uses equal weights; voting classification uses a majority rule.

`voteweights(numlist)` defines positive weights used for voting regression or classification. The length of *numlist* should be the number of base learners  $-1$ . The last weight is calculated to ensure that the sum of weights equals 1.

### 3.5 Base learners

The base learners are chosen using the option `method()` in combination with `type()`. The latter can take the value `regress` for regression and `classify` for classification problems. Table 1 provides an overview of supported base learners and their underlying `scikit-learn` routines.

`cmdopt*()` (syntax 1) and `opt()` (syntax 2) are used to pass options to the base learners. Because of the many options, we do not list all options here. We instead provide a tool that lists options for each base learner. For example, to get the default options for lasso with cross-validated penalty, type

```
. pystacked, type(regress) methods(lassocv) printopt
  (output omitted)
Specified options are translated to:
alphas(None) l1_ratio(1) eps(.001) n_alphas(100) fit_intercept(True)
max_iter(1000) tol(.0001) n_jobs(None) positive(False) selection(cyclic)
random_state(rng)
```

The naming of the options follows `scikit-learn`. Allowed settings for each option can be inferred from the `scikit-learn` documentation. We strongly recommend that the user reads the `scikit-learn` documentation carefully.<sup>1</sup>

---

1. The `scikit-learn` documentation is available at <https://scikit-learn.org/stable/index.html> (last accessed on October 11, 2023).

Table 1. Overview of machine learners available in `pystacked`

<code>method()</code>	<code>type()</code>	Machine learner description	<code>scikit-learn</code> program
<code>ols</code>	<code>regress</code>	Linear regression	<code>linear_model.LinearRegression</code>
<code>logit</code>	<code>classify</code>	Logistic regression	<code>linear_model.LogisticRegression</code>
<code>lassoic</code>	<code>regress</code>	Lasso with AIC/BIC penalty	<code>linear_model.LassoLarsIC</code>
<code>lassocv</code>	<code>regress</code>	Lasso with CV penalty	<code>linear_model.ElasticNetCV</code>
	<code>classify</code>	Logistic lasso with CV penalty	<code>linear_model.LogisticRegressionCV</code>
<code>ridgecv</code>	<code>regress</code>	Ridge with CV penalty	<code>linear_model.ElasticNetCV</code>
	<code>classify</code>	Logistic ridge with CV penalty	<code>linear_model.LogisticRegressionCV</code>
<code>elasticcv</code>	<code>regress</code>	Elastic net with CV penalty	<code>linear_model.ElasticNetCV</code>
	<code>classify</code>	Logistic elastic net with CV penalty	<code>linear_model.LogisticRegressionCV</code>
<code>svm</code>	<code>regress</code>	Support vector regression	<code>svm.SVR</code>
	<code>classify</code>	Support vector classification (SVC)	<code>svm.SVC</code>
<code>gradboost</code>	<code>regress</code>	Gradient boosting regression	<code>ensemble.GradientBoostingRegressor</code>
	<code>classify</code>	Gradient boosting classification	<code>ensemble.GradientBoostingClassifier</code>
<code>rf</code>	<code>regress</code>	Random forest regression	<code>ensemble.RandomForestRegressor</code>
	<code>classify</code>	Random forest classification	<code>ensemble.RandomForestClassifier</code>
<code>linsvm</code>	<code>classify</code>	Linear SVC	<code>svm.LinearSVC</code>
<code>nnet</code>	<code>regress</code>	Neural net regression	<code>sklearn.neural_network.MLPRegressor</code>
	<code>classify</code>	Neural net classification	<code>sklearn.neural_network.MLPClassifier</code>

NOTES: The first two columns list all allowed combinations of `method()` and `type()`, which are used to select base learners. The third column provides a description of each machine learner. The last column lists the underlying `scikit-learn` learn routine. “CV penalty” indicates that the penalty level is chosen to minimize the cross-validated MSPE. “AIC/BIC penalty” indicates that the penalty level minimizes either the AIC or the BIC.

### 3.6 Learner-specific predictors

By default, `pystacked` uses the same set of predictors for each base learner. Using the same predictors for each method is often not desirable because the optimal set of predictors may vary across base learners. For example, when one uses linear machine learners such as the lasso, adding polynomials, interactions, and other transformations of the base set of predictors might greatly improve out-of-sample prediction performance. The inclusion of transformations of base predictors is especially worth considering if the base set of observed predictors is small (relative to the sample size) and the relationship between outcome and predictors is likely nonlinear. Tree-based methods (for example, random forests and boosted trees), on the other hand, can detect certain types of nonlinear patterns automatically. While adding transformations of the base predictors may still lead to performance gains, the added benefit is less striking relative to linear learners and might not justify the additional costs in terms of computational complexity.

There are two approaches to implement learner-specific sets of predictors: Pipelines, discussed in the next section, can be used to create some transformations on the fly for specific base learners. A more flexible approach is the `xvars*`() option, which allows specifying predictors for a particular learner. `xvars*`() supports standard Stata factor-variable notation.

### 3.7 Pipelines

`scikit-learn` uses pipelines to preprocess input data on the fly. In `pystacked`, pipelines can be used to impute missing values, create polynomials and interactions, and standardize predictors. Table 2 lists the pipelines currently supported by `pystacked`.

Table 2. Pipelines supported by `pystacked`

<code>pipe*()</code>	Description	<code>scikit-learn</code> programs
<code>stdscaler</code>	Standardize to mean zero and unit variance (default for regularized regression)	<code>StandardScaler()</code>
<code>nostdscaler</code>	Overwrite default standardization for regularized regression	n/a
<code>stdscaler0</code>	Standardize to unit variance	<code>StandardScaler(with_mean=False)</code>
<code>sparse</code>	Transform to sparse matrix	<code>SparseTransformer()</code>
<code>onehot</code>	Create dummies from categorical variables	<code>OneHotEncoder()</code>
<code>minmaxscaler</code>	Scale to 0–1 range	<code>MinMaxScaler()</code>
<code>medianimputer</code>	Median imputation	<code>SimpleImputer(strategy='median')</code>
<code>knnimputer</code>	KNN imputation	<code>KNNImputer()</code>
<code>poly2</code>	Add second-order polynomials	<code>PolynomialFeatures(degree=2)</code>
<code>poly3</code>	Add third-order polynomials	<code>PolynomialFeatures(degree=3)</code>
<code>interact</code>	Add interactions	<code>PolynomialFeatures(include_bias=False, interaction_only=True)</code>

**Remarks.** First, regularized regressors (the methods `lassoic`, `lassocv`, `ridgecv`, and `elasticcv`) use the `stdscaler` pipeline by default. `pipe*(nostdscaler)` disables this behavior. Second, the `stdscaler0` pipeline is useful in combination with `sparse`, which transforms the predictor matrix into a sparse matrix. `stdscaler0` does not center predictors so that the predictor matrix retains its sparsity property.

## 4 Applications

This section demonstrates how to apply `pystacked` for regression and classification tasks. Before we discuss stacking, we first show how to use `pystacked` as a “regular” machine learning program for fitting one supervised machine learner (see the next subsection). We then illustrate stacking regression and stacking classification in sections 4.2 and 4.3, respectively.

## 4.1 One base learner

We import the Pace and Barry (1997) California house price data and split the sample randomly into training and validation partitions using an approximate 75/25 split. The aim of the prediction task is to predict median house prices (`medhousevalue`) using a set of house price characteristics.<sup>2</sup> We prepare the data for analysis as follows:

```
. use https://statalasso.github.io/dta/cal_housing
. set seed 42
. generate train=runiform()
. replace train = train < .75
(20,640 real changes made)
. replace medh = medh/10e3
variable medhousevalue was long now double
(20,640 real changes made)
```

**Gradient-boosted trees.** As a first example, we use `pystacked` to fit gradient-boosted regression trees and save the out-of-sample predicted values. Because we consider a regression task rather than a classification task, we specify `type(regress)` (but because this is the default, it could also be omitted). The option `methods(gradboost)` selects gradient boosting. We will later see that we can specify more than one learner in `methods()` and that we can also fit gradient-boosted classification trees.

```
. pystacked medh longi-medi if train, type(regress) methods(gradboost)
Single base learner: no stacking done.
Stacking weights:
-----
| Method | Weight |
-----
| gradboost | 1.0000000 |
. predict double yhat_gb1 if !train
```

The output shows the stacking weights associated with each base learner. Because we only consider one method, the output is not particularly informative and simply shows a weight of one for gradient boosting. Yet `pystacked` has fit 100 boosted trees (the default) in the background using `scikit-learn`'s `ensemble.GradientBoostedRegressor`.

Before we tune our gradient boosting learner, we retrieve a list of available options. The default options for gradient boosting can be listed in the console, and the `scikit-learn` documentation provides more detail on the allowed parameters of each option.

---

2. The following predictors are included in the dataset in this order: district longitude (`longitude`), latitude (`latitude`), median house age (`houseage`), average number of rooms per household (`rooms`), average number of bedrooms per household (`bedrooms`), block group population (`population`), average number of household members (`households`), and median income in block group (`medinc`). The dataset was retrieved from the StatLib repository at [https://www.dcc.fc.up.pt/~ltorgo/Regression/cal\\_housing.html](https://www.dcc.fc.up.pt/~ltorgo/Regression/cal_housing.html) (last accessed on February 19, 2022). There are 20,640 observations in the dataset.

```
. pystacked, type(regress) methods(gradboost) printopt
(output omitted)
Specified options are translated to:
loss(squared_error) learning_rate(.1) n_estimators(100) subsample(1)
criterion(friedman_mse) min_samples_split(2) min_samples_leaf(1)
min_weight_fraction_leaf(0) max_depth(3) min_impurity_decrease(0)
init(None) random_state(rng) max_features(None) alpha(.9)
max_leaf_nodes(None) warm_start(False) validation_fraction(.1)
n_iter_no_change(None) tol(.0001) ccp_alpha(0)
```

We consider two additional specifications. Note that we restrict ourselves to a few selected specifications for illustrative purposes. We stress that careful parameter tuning should consider a grid of values across multiple learner parameters. The first specification reduces the learning rate from 0.1 (the default) to 0.01. The second specification reduces the learning rate and increases the number of trees from 100 (the default) to 1,000. We use `cmdopt1()` because gradient boosting is the first (and only) method listed in `methods()`.

```
. pystacked medh longi-medi if train,
> type(regress) methods(gradboost)
> cmdopt1(learning_rate(.01))
Single base learner: no stacking done.
Stacking weights:


| Method    | Weight    |
|-----------|-----------|
| gradboost | 1.0000000 |



```
. predict double yhat_gb2 if !train
. pystacked medh longi-medi if train,
> type(regress) methods(gradboost)
> cmdopt1(learning_rate(.01) n_estimators(1000))
Single base learner: no stacking done.
Stacking weights:


Method	Weight
gradboost	1.0000000


. predict double yhat_gb3 if !train
```


```

We can then compare the performance across the three models using the out-of-sample mean squared prediction error (MSPE):

```
. generate double res_gb1_sq = (medh-yhat_gb1)^2 if !train
(15,448 missing values generated)
. generate double res_gb2_sq = (medh-yhat_gb2)^2 if !train
(15,448 missing values generated)
. generate double res_gb3_sq = (medh-yhat_gb3)^2 if !train
(15,448 missing values generated)
. summarize res_gb* if !train
```

Variable	Obs	Mean	Std. dev.	Min	Max
res_gb1_sq	5,192	29.86727	83.17099	3.92e-06	1424.438
res_gb2_sq	5,192	71.10136	129.1446	4.57e-07	1230.348
res_gb3_sq	5,192	30.37613	82.91207	1.24e-06	1384.924

The initial gradient booster achieves an out-of-sample MSPE of 29.87. The second gradient booster uses a reduced learning rate of 0.01 and performs much worse, with an MSPE of 71.10. The third gradient booster performs only slightly worse than the first, illustrating the tradeoff between the learning rate and the number of trees.

**Pipelines.** We can make use of pipelines to preprocess our predictors. This is especially useful in the context of stacking when we want to, for example, use second-order polynomials of predictors as inputs for one method but only use elementary predictors for another method. Here we compare lasso with and without the `poly2` pipeline:

```
. pystacked medh longi-medi if train, type(regress) methods(lassocv)
Single base learner: no stacking done.
Stacking weights:
```

Method	Weight
lassocv	1.0000000

```
. predict double yhat_lasso1 if !train
. pystacked medh longi-medi if train, type(regress) methods(lassocv)
> pipe1(poly2)
Single base learner: no stacking done.
Stacking weights:
```

Method	Weight
lassocv	1.0000000

```
. predict double yhat_lasso2 if !train
```

We could replace `pipe1(poly2)` with `xvars1(c.(medh longi-medi)##c.(medh longi-medi))`. In fact, the latter is more flexible and allows, for example, creation of interactions for some predictors and not for others.

We again calculate the out-of-sample MSPE:

```
. generate double res_lasso1_sq = (medh-yhat_lasso1)^2 if !train
(15,448 missing values generated)
. generate double res_lasso2_sq = (medh-yhat_lasso2)^2 if !train
(15,448 missing values generated)
. summarize res_lasso1_sq res_lasso2_sq if !train
```

Variable	Obs	Mean	Std. dev.	Min	Max
res_lasso1_sq	5,192	47.00385	108.0759	8.43e-07	2392.572
res_lasso2_sq	5,192	43.91776	109.5823	2.90e-07	2559.051

The `poly2` pipeline improves the performance of the lasso, indicating that squared and interaction terms constitute important predictors. However, the lasso does not perform as well as gradient boosting in this application.

## 4.2 Stacking regression

We now consider a stacking regression application with five base learners: 1) linear regression, 2) lasso with penalty chosen by cross-validation, 3) lasso with second-order polynomials and interactions, 4) random forest with default settings, and 5) gradient boosting with a learning rate of 0.01 and 1,000 trees. That is, we use the lasso twice—once with and once without the `poly2` pipeline. Indeed, nothing keeps us from using the same algorithm multiple times. This way, we can combine the same algorithm with different settings.

Note the numbering of the `pipe*()` and `cmdopt*()` options below. We apply the `poly2` pipe to the third method (`lassocv`). We also change the default learning rate and the number of estimators for gradient boosting (the fifth estimator).

```
. set seed 42
. pystacked medh longi-medi if train,
>   type(regress)
>   methods(ols lasso cv lasso cv rf gradboost)
>   pipe3(poly2) cmdopt5(learning_rate(0.01)
>   n_estimators(1000))
```

Stacking weights:

Method	Weight
ols	0.0000000
lassocv	0.0000000
lassocv	0.0000000
rf	0.8382965
gradboost	0.1617035

The above syntax becomes a bit difficult to read with many methods and many options. `pystacked`'s second syntax is easier to use with many base learners.

```
. set seed 42
. pystacked medh longi-medi          ||
>   method(ols)                    ||
>   method(lassocv)                ||
>   method(lassocv) pipeline(poly2) ||
>   method(rf)                      ||
>   method(gradboost) options(learning_rate(0.01) n_estimators(1000))
>   if train, type(regress)
```

Stacking weights:

Method	Weight
ols	0.0000000
lassocv	0.0000000
lassocv	0.0000000
rf	0.8382965
gradboost	0.1617035

The stacking weights shown in the output determine how much each method contributes to the final stacking predictor. In this example, ordinary least squares and lasso based on both sets of predictors all receive a weight of zero. Random forest receives a weight of 83.8%, and gradient boosting contributes the remaining 16.2% of the weight to the final predictor.

**Predicted values.** In addition to the stacking predicted values, we can also get the predicted values of each base learner using the `basexb` option:

```
. predict double yhat, xb
. predict double ybase, basexb
. list yhat ybase* if _n <= 5
```

	yhat	ybase1	ybase2	ybase3	ybase4	ybase5
1.	42.875278	41.315834	41.193417	40.02005	43.160831	41.394928
2.	38.736568	41.45306	41.421459	44.443344	38.289108	41.05629
3.	40.683399	38.212036	38.154835	37.796288	41.268902	37.648071
4.	33.559074	32.332497	32.266319	32.546425	33.4869	33.933231
5.	24.190563	25.382839	25.369064	25.269325	23.859699	25.905815

**Plotting.** `pystacked` also comes with plotting features. The `graph` option creates a scatterplot of predicted values on the vertical and observed values on the horizontal axis for stacking and each base learner; see figure 1. The black line is a 45-degree line and shown for reference. Because `pystacked` with `graph` can be used as a postestimation command, there is no need to rerun the stacking estimation.



```
. pystacked, graph(scheme(sj)) lgraph(scheme(sj)) holdout
Number of holdout observations: 5192
```




Figure 1. Out-of-sample predicted values and observed values created using the `graph` option after stacking regression

Figure 1 shows the out-of-sample predicted values. To see the in-sample predicted values, simply omit the `holdout` option. Note that the `holdout` option will not work if the estimation was run on the whole sample.

**RMSPE table.** The `table` option allows comparing stacking weights with in-sample, cross-validated, and out-of-sample RMSPEs. As with the `graph` option, we can use `table` as a postestimation command:

```
. pystacked, table holdout
Number of holdout observations: 5192
RMSPE: In-Sample, CV, Holdout
```

Method	Weight	In-Sample	CV	Holdout
STACKING	.	2.313	4.980	4.939
ols	0.000	6.986	7.008	6.853
lassocv	0.000	6.987	7.008	6.857
lassocv	0.000	6.696	6.699	6.606
rf	0.838	1.847	5.001	4.963
gradboost	0.162	5.312	5.523	5.511

### 4.3 Stacking classification

`pystacked` can be applied to binary classification problems. For demonstration, we consider the Spambase data of Cranor and LaMacchia (1998), which we retrieve from the UCI Machine Learning Repository. We load the data and split the data into training and validation samples using an approximate 75/25 split.

```
. insheet using
> https://archive.ics.uci.edu/ml/machine-learning-databases/spambase/spambase.data,
> clear comma
(58 vars, 4,601 obs)
. set seed 42
. generate train = runiform()
. replace train = train<.75
(4,601 real changes made)
```

The example below is more complicated. We go through it step by step:

- We use five base learners: logistic regression, two gradient boosters, and two neural nets.
- We apply the `poly2` pipeline to logistic regression, which creates squares and interaction terms of the predictors, but not to other methods.
- We use gradient boosting with 600 and 1,000 classification trees.
- We consider two specifications for the neural nets: one neural net with two hidden layers of five nodes each and another neural net with one hidden layer of five nodes.
- Finally, we use `type(classify)` to specify that we are considering a classification task and `njobs(8)` switches parallelization on using 8 cores.

```
. pystacked v58 v1-v57 ||
> method(logit) pipeline(poly2) ||
> method(gradboost) options(n_estimators(600)) ||
> method(gradboost) options(n_estimators(1000)) ||
> method(nnet) options(hidden_layer_sizes(5 5)) ||
> method(nnet) options(hidden_layer_sizes(5)) ||
> if train, type(classify) njobs(8)
```

Stacking weights:

Method	Weight
logit	0.0000000
gradboost	0.4917594
gradboost	0.3345601
nnet	0.1171906
nnet	0.0564899

As in the previous regression example, gradient boosting receives the largest stacking weight and thus contributes most to the final stacking prediction.

**Confusion matrix.** Confusion matrices allow comparing actual and predicted outcomes in a  $2 \times 2$  matrix. `pystacked` provides a compact table format that combines confusion matrices for each base learner and for the final stacking classifier for both the training and the validation partitions.

```
. pystacked, table holdout
Number of holdout observations: 1133
Confusion matrix: In-Sample, CV, Holdout
```

Method		Weight	In-Sample		CV		Holdout	
			0	1	0	1	0	1
STACKING	0	.	2079	6	2004	80	678	29
STACKING	1	.	2	1381	77	1307	29	397
logit	0	0.000	1266	88	0	1	413	32
logit	1	0.000	815	1299	0	21	294	394
gradboost	0	0.492	2078	9	0	0	679	30
gradboost	1	0.492	3	1378	0	0	28	396
gradboost	0	0.335	2079	1	0	0	680	29
gradboost	1	0.335	2	1386	0	0	27	397
nnet	0	0.117	1834	55	0	0	617	28
nnet	1	0.117	247	1332	0	6	90	398
nnet	0	0.056	2005	146	0	0	671	51
nnet	1	0.056	76	1241	0	2	36	375

For example, the table shows 29 false positives for stacking and 294 for logistic regression in the validation partition, while the number of false negatives is 29 and 32, respectively. The accuracy of stacking and logistic regression is thus given by  $(678 + 397)/1133 = 94.9\%$  and  $(413 + 394)/1133 = 71.2\%$ .

**Plotting.** `pystacked` supports ROC curves, which allow assessment of the classification performance for varying discrimination thresholds. The  $y$  axis in a ROC plot corresponds to sensitivity (true-positive rate), and the  $x$  axis corresponds to  $1 - \text{specificity}$  (false-positive rate). The area under the curve displayed below each ROC plot is a common evaluation metric for classification problems.

```
. pystacked,
> graph(subtitle(Spam data) scheme(sj))
> lgraph(plotopts(msymbol(i))
> ylabel(0 1, format(%3.1f))) scheme(sj))
> holdout
Number of holdout observations: 1133
```




Figure 2. Out-of-sample ROC curves created using the `graph` option after stacking regression

## 5 Conclusion

In this article, we introduced the `pystacked` command. The command not only makes a range of popular supervised machine learners available, such as regularized regression, random forests, and gradient-boosted trees, but also facilitates combining multiple learners for stacking regression or classification. `pystacked` comes with several practically relevant features: for example, sparse matrix support, learner-specific predictors, pipelines for predictor transformations, and two alternative syntaxes. Nevertheless, we also see the potential for improvement in at least two directions. First, `pystacked` offers few diagnostics for individual learners. Features such as variable importance plots and coefficient estimates for parametric learners could be added in later versions.<sup>3</sup> Second, the deep learning features of `pystacked` are currently relatively limited and could be improved by adding support for other deep learning algorithms.

## 6 Acknowledgments

We thank Jan Ditzen, Ben Jann, Blaise Melly, Alessandro Oliveira, Matthias Schonlau, and Thomas Wiemann for their helpful feedback. All remaining errors are our own.

## 7 Programs and supplemental material

To install the software files as they existed at the time of the publication of this article, type

```
. net sj 23-4
. net install st0731      (to install program files, if available)
. net get st0731         (to install ancillary files, if available)
```

To install the latest version of `pystacked`:

```
. net install pystacked,
> from(https://raw.githubusercontent.com/aahrens1/pystacked/main) replace
```

## 8 References

- Ahrens, A., C. B. Hansen, and M. E. Schaffer. 2020. lasso-pack: Model selection and prediction with regularized regression in Stata. *Stata Journal* 20: 176–235. <https://doi.org/10.1177/1536867X20909697>.
- Ahrens, A., C. B. Hansen, M. E. Schaffer, and T. Wiemann. 2023. ddml: Stata module for double/debiased machine learning. Statistical Software Components S459175, Department of Economics, Boston College. <https://ideas.repec.org/c/boc/bocode/s459175.html>.

---

3. We introduced the `showcoef()` option in version 0.7.3 of `pystacked`. This option allows to display coefficient estimates.

- . Forthcoming. ddml: Double/debiased machine learning in Stata. *Stata Journal*.
- Athey, S., and G. W. Imbens. 2019. Machine learning methods that economists should know about. *Annual Review of Economics* 11: 685–725. <https://doi.org/10.1146/annurev-economics-080217-053433>.
- Breiman, L. 1996. Stacked regressions. *Machine Learning* 24: 49–64. <https://doi.org/10.1007/BF00117832>.
- Buitinck, L., G. Louppe, M. Blondel, F. Pedregosa, A. Müller, O. Grisel, V. Niculae, et al. 2013. API design for machine learning software: Experiences from the scikit-learn project. European Conference on Machine Learning and Principles and Practices of Knowledge Discovery in Databases Workshop: Languages for Data Mining and Machine Learning.
- Cerulli, G. 2022. Machine learning using Stata/Python. *Stata Journal* 22: 772–810. <https://doi.org/10.1177/1536867X221140944>.
- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins. 2018. Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal* 21: C1–C68. <https://doi.org/10.1111/ectj.12097>.
- Cranor, L. F., and B. A. LaMacchia. 1998. Spam! *Communications of the ACM* 41: 74–83. <https://doi.org/10.1145/280324.280336>.
- Droste, M. 2022. stata-pylearn. GitHub. <https://github.com/mdroste/stata-pylearn>.
- Fedorova, E., S. Ledyeva, P. Drogovoz, and A. Nevredinov. 2022. Economic policy uncertainty and bankruptcy filings. *International Review of Financial Analysis* 82: 102174. <https://doi.org/10.1016/j.irfa.2022.102174>.
- Guenther, N., and M. Schonlau. 2016. Support vector machines. *Stata Journal* 16: 917–937. <https://doi.org/10.1177/1536867X1601600407>.
- . 2018. svmachines: Stata module providing support vector machines for both classification and regression. Statistical Software Components S458564, Department of Economics, Boston College. <https://ideas.repec.org/c/boc/bocode/s458564.html>.
- Hastie, T., R. Tibshirani, and J. Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd ed. New York: Springer. <https://doi.org/10.1007/978-0-387-84858-7>.
- Hooker, J., G. Duveiller, and A. Cescatti. 2018. A global dataset of air temperature derived from satellite remote sensing and weather stations. *Scientific Data* 5: 180246. <https://doi.org/10.1038/sdata.2018.246>.
- Huntington-Klein, N. C. 2021. mlrtime. GitHub. <https://github.com/NickCH-K/MLRtime/>.

- Hwangbo, L., Y. J. Kang, H. Kwon, J. I. Lee, H.-J. Cho, J.-K. Ko, S. M. Sung, and T. H. Lee. 2022. Stacking ensemble learning model to predict 6-month mortality in ischemic stroke patients. *Scientific Reports* 12: 17389. <https://doi.org/10.1038/s41598-022-22323-9>.
- Kingma, D. P., and J. Ba. 2017. Adam: A method for stochastic optimization. arXiv:1412.6980 [cs.LG]. <https://doi.org/10.48550/arXiv.1412.6980>.
- Liang, D., C.-F. Tsai, H.-Y. Lu, and L.-S. Chang. 2020. Combining corporate governance indicators with stacking ensembles for financial distress prediction. *Journal of Business Research* 120: 137–146. <https://doi.org/10.1016/j.jbusres.2020.07.052>.
- Pace, R. K., and R. Barry. 1997. Sparse spatial autoregressions. *Statistics and Probability Letters* 33: 291–297. [https://doi.org/10.1016/S0167-7152\(96\)00140-X](https://doi.org/10.1016/S0167-7152(96)00140-X).
- Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, et al. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12: 2825–2830. <https://doi.org/10.48550/arXiv.1201.0490>.
- Schonlau, M., and R. Y. Zou. 2020. The random forest algorithm for statistical learning. *Stata Journal* 20: 3–29. <https://doi.org/10.1177/1536867X20909688>.
- van der Laan, M. J., E. C. Polley, and A. E. Hubbard. 2007. Super learner. *Statistical Applications in Genetics and Molecular Biology* 6. <https://doi.org/10.2202/1544-6115.1309>.
- van der Laan, M. J., and S. Rose. 2011. *Targeted Learning: Causal Inference for Observational and Experimental Data*. New York: Springer. <https://doi.org/10.1007/978-1-4419-9782-1>.
- Wolpert, D. H. 1992. Stacked generalization. *Neural Networks* 5: 241–259. [https://doi.org/10.1016/S0893-6080\(05\)80023-1](https://doi.org/10.1016/S0893-6080(05)80023-1).

#### About the authors

Achim Ahrens is a postdoctoral researcher and senior data scientist at the Public Policy Group and Immigration Policy Lab, ETH Zürich.

Christian B. Hansen is the Wallace W. Booth Professor of Econometrics and Statistics at the University of Chicago Booth School of Business.

Mark E. Schaffer is a professor of economics in the Edinburgh Business School at Heriot-Watt University, Edinburgh, U.K., and a research fellow at the Institute of Labor Economics (IZA), Bonn.