



dame-flame: A Python Package Providing Fast Interpretable Matching for Causal Inference

Neha R. Gupta 

Carnegie Mellon University

Vittorio Orlandi 

Duke University

Chia-Rui Chang 

Harvard University

Tianyu Wang 

Fudan University

Marco Morucci 

New York University

Pritam Dey 

Duke University

Thomas J. Howell

Duke University

Xian Sun

Duke University

Angikar Ghosal

Duke University

Sudeepa Roy 

Duke University

Cynthia Rudin 

Duke University

Alexander Volfovsky 

Duke University

Abstract

dame-flame is a Python package for performing *matching* for *observational causal inference* on datasets containing discrete covariates. This package implements the *dynamic almost matching exactly (DAME)* and *fast, large-scale almost matching exactly (FLAME)* algorithms, which match treatment and control units on subsets of the covariates. The resulting matched groups are interpretable, because the matches are made directly on covariates, and high-quality, because machine learning is used to determine which covariates are important to match on instead of human inputs. The package provides several adjustable parameters to adapt the algorithms to specific applications, and can calculate treatment effects after matching. The most recent source code of the implementation is available at <https://github.com/almost-matching-exactly/DAME-FLAME-Python-Package>.

Keywords: causal inference, machine learning, matching, Python.

1. Introduction

The **dame-flame** package for Python (Van Rossum *et al.* 2011) is the first major implementa-

tion of two algorithms, the *dynamic almost matching exactly* (DAME) algorithm (Dieng, Liu, Roy, Rudin, and Volfovsky 2019, published in AISTATS’19), and the *fast, large-scale almost matching exactly* (FLAME) algorithm (Wang *et al.* 2021, published in JMLR’21). These algorithms provide *almost exact* matching of treatment and control units in discrete observational data for causal inference, allowing for accurate causal effect estimation. As discussed in Dieng *et al.* (2019), and Wang *et al.* (2021), the two algorithms produce high-quality interpretable matched groups, by using machine learning on a holdout training set to learn distance metrics. DAME solves an optimization problem that matches units on as many covariates as possible, prioritizing matches on important covariates. FLAME approximates the solution found by DAME via a much faster backward feature selection procedure.

The DAME and FLAME algorithms are discussed in the remainder of this section. We also provide testing and installation details. In Section 2, we discuss the class structure in the **dame-flame** package, detail special features of **dame-flame**, and compare **dame-flame** to other matching packages. In Section 3, we offer examples, including a quickstart example illustrating the goal of the package in Section 3.1, and API documentation is in Section 4.

1.1. Algorithms overview

The advantage of matching is that it allows for accurate treatment effect estimates, and permits interpretable analyses that are easier to troubleshoot than other types of analysis for observational causal studies. However, matching is not trivial; in high dimensional settings, few individuals can be matched exactly on all covariates, so other ways must be found. We offer the Python package **dame-flame** to support *almost exact matching* in a way that focuses on identifying important subsets of the covariates using machine learning and matching units exactly on those subsets.

dame-flame is designed for causal inference problems with a binary treatment variable, an observed outcome variable, and any number of pre-treatment covariates. Several assumptions standard in observational causal inference must be made in order for the DAME and FLAME algorithms to be applicable.

The first assumption is stable unit treatment value assumption (SUTVA), which assumes that treatments applied to one unit do not affect the outcome of other units, and there is only one version of treatment.

A second assumption is unconfoundedness, or ignorability. It is important that the potential outcome is independent of the treatment assignment.

The third is overlap of treatment and control groups. The treatment and control groups are said to not have any overlap at some location in a distribution when the probability of being treated at that location is either exactly 0 or 1. If there is no overlap for all covariates, then FLAME and DAME algorithms would not be able to find any matches. A more moderate issue is when only few treated and control units overlap in covariate values, or partial overlap, where we may not find both treatment and control units with sufficient overlap to match with. In this case, the user’s settings on the **dame-flame** package would determine what *quality of matches* would be acceptable. Match quality is discussed further in Section 1.2. Units that were not able to be matched are not included in treatment effect calculations.

Discrete observed covariate data is a requirement of **dame-flame**. The covariate values do not need to be ordinal if an appropriate learning method is chosen, but need to be categorical. We recommend users only bin continuous variables in a way that is meaningful for the problem

at hand, as shown in Section 3.2. The outcome values must be continuous real values with ordinal meaning, or can be binary if an appropriate learning method is chosen (i.e., logistic regression).

dame-flame is efficient, owing to a combination of fast bit-vector computations, and a backwards feature elimination process (for FLAME) or a type of downwards closure property for systematic feature elimination (for DAME). Therefore, FLAME is faster, but DAME is able to match units on more covariates.

We also support a hybrid execution of FLAME and DAME methods. The combination of FLAME (at earlier iterations) and DAME (at later iterations) permits faster elimination of irrelevant covariates in the earlier iterations and then a more careful elimination of covariates in the later iterations, thereby achieving a trade-off between scalability and quality.

The accuracy of these methods has been established in prior works, [Dieng et al. \(2019\)](#) and [Wang et al. \(2021\)](#), in which the authors run experiments on simulated datasets to show the estimated treatment effects closely match true treatment effects. FLAME is shown to offer more accurate treatment effect estimates than methods that do not distinguish irrelevant variables in [Wang et al. \(2021\)](#), and DAME is shown to perform more accurately than FLAME in [Dieng et al. \(2019\)](#).

1.2. Algorithm methodology

In this section we describe the algorithms implemented. First we discuss the mathematical problem that FLAME and DAME aim to solve, then we describe the steps of each algorithm.

Suppose we have n units, indexed by i , and p covariates. We may interchangeably refer to units as “individuals” or “observations”. Formally, consider a dataframe $\mathcal{S} = [X, Y, T]$, including $n \times p$ matrix $X \in \{0, 1, \dots, k\}^{n \times p}$ where X contains the categorical covariates for all units, $Y \in \mathbb{R}^n$ denotes the outcome vector, and $T \in \{0, 1\}^n$ denotes the treatment indicator vector (1 for treated, 0 for control); \mathbf{x}_i denotes the covariate vector of unit i .

We will use $\theta \in \{0, 1\}^p$ to denote the variable selection indicator vector for a subset of covariates to match on. A unit is a triplet (covariate value \mathbf{x}_i , observed outcome y_i , treatment indicator t_i). Given dataset \mathcal{S} , define the *matched group* for unit i with respect to covariates selected by θ as the units in \mathcal{S} that match i exactly on covariates θ :

$$\text{MG}_i(\theta, \mathcal{S}) = \{i' \in \mathcal{S} : \mathbf{x}_{i'} \circ \theta = \mathbf{x}_i \circ \theta\},$$

where \circ denotes Hadamard product. Under the assumption of no unobserved confounding, the question of the causal effect of T on Y then becomes which covariates θ we should match unit i on.

In FLAME and DAME, the value of a set of covariates θ is determined by how well these covariates can be used together to predict outcomes (and by extension, treatment effects), since it is not possible to match every unit on all covariates. However, we often prefer not to look at the outcomes of our dataset to determine how to match, to avoid risk of biasing the estimates. Thus, we consider a separate training dataset \mathcal{S}^{tr} . Let \mathcal{S}_0^{tr} be the subset (of \mathcal{S}^{tr}) of control units $((X^{tr}, Y^{tr})$ with $T^{tr} = 0$), and let \mathcal{S}_1^{tr} be the subset (of \mathcal{S}^{tr}) of treated units $((X^{tr}, Y^{tr})$ with $T^{tr} = 1$). The empirical prediction error $\widehat{\text{PE}}_{\mathcal{F}_{\|\theta\|_0}}$ is defined with respect to a

class of functions $\mathcal{F}_k := \{f : \{0, 1\}^k \rightarrow [0, 1]\}$ ($1 \leq k \leq d$) as:

$$\widehat{\text{PE}}_{\mathcal{F}_{\|\theta\|_0}}(\theta, \mathcal{S}^{tr}) = \min_{f^{(1)} \in \mathcal{F}_{\|\theta\|_0}} \frac{1}{|\mathcal{S}_1^{tr}|} \sum_{(\mathbf{x}_i, y_i) \in \mathcal{S}_1^{tr}} (f^{(1)}(\mathbf{x}_i \circ \theta) - y_i)^2 + \min_{f^{(0)} \in \mathcal{F}_{\|\theta\|_0}} \frac{1}{|\mathcal{S}_0^{tr}|} \sum_{(\mathbf{x}_i, y_i) \in \mathcal{S}_0^{tr}} (f^{(0)}(\mathbf{x}_i \circ \theta) - y_i)^2.$$

That is, $\widehat{\text{PE}}_{\mathcal{F}_{\|\theta\|_0}}$ is the smallest prediction error we can get on both treatment and control populations using the features specified by θ . Thus, given a matching dataset \mathcal{S}^{ma} and a training dataset \mathcal{S}^{tr} , the best selection indicator we could achieve for a nontrivial matched group that contains treatment unit i would be:

$$\theta_{i, \mathcal{S}^{ma}}^* \in \arg \min_{\theta} \widehat{\text{PE}}_{\mathcal{F}_{\|\theta\|_0}}(\theta, \mathcal{S}^{tr}) \text{ s.t. } \exists \ell \in \text{MG}_i(\theta, \mathcal{S}^{ma}) \text{ s.t. } t_{\ell} = 0. \quad (1)$$

This constraint says that the matched group of any treatment unit, here denoted unit i , contains at least one control unit, which here is denoted unit ℓ . It also matches on covariates that together can be used to predict well on the training set. The covariates selected by $\theta_{i, \mathcal{S}^{ma}}^*$ are those that predict the outcome best, provided that at least one control unit has the same exact covariate values as i on the covariates selected by $\theta_{i, \mathcal{S}^{ma}}^*$. Please note that the predictive error is not the sole determinant of a matched group, and the covariates used in a matched group are determined based on an iterative procedure. This is discussed in more detail below.

The *main matched group* for i is then defined as $\text{MG}_i(\theta_{i, \mathcal{S}^{ma}}^*, \mathcal{S}^{ma})$. Users can choose whether units are matched with replacement; that is, whether a previously matched unit can be matched in a subsequent iteration of the algorithm. The first time a unit is matched, that matched set is its *main matched group*, from which its treatment effect estimates are calculated. If units are allowed to be matched with replacement, a unit can become a member of another unit's main matched group. Any additional groups which a unit belongs to other than its *main matched group* is its *auxiliary matched group*.

The goal of FLAME and DAME is to calculate the main matched group $\text{MG}_i(\theta_{i, \mathcal{S}^{ma}}^*, \mathcal{S}^{ma})$ *for as many units i as possible*. Any units without a main matched group are likely outliers that cannot be easily matched. Then, the matched groups can be used to estimate treatment effects.

The implementation of the above mathematical descriptions in both DAME and FLAME algorithms requires us to iterate over two nested loops, shown in Figure 1.

Users must begin by providing a dataset with discrete observational covariates, a binary treatment indicator, and a continuous or discrete outcome column. If users do not supply a separate holdout training dataset with the same covariates as the input dataset, optional parameters of the package, discussed in further detail in Section 4, will allow them to partition the input dataset to create the holdout training set.

Users have a variety of options for handling *missing covariate data*. They can (1) exclude units with missing values from the procedure, (2) impute the missing data via the multiple imputation by chained equations (MICE) method (Van Buuren and Groothuis-Oudshoorn 2010), or (3) specify that matches should not occur on missing values without imputing them. In this third case, matches can still be made for a unit on its covariates that are not missing.

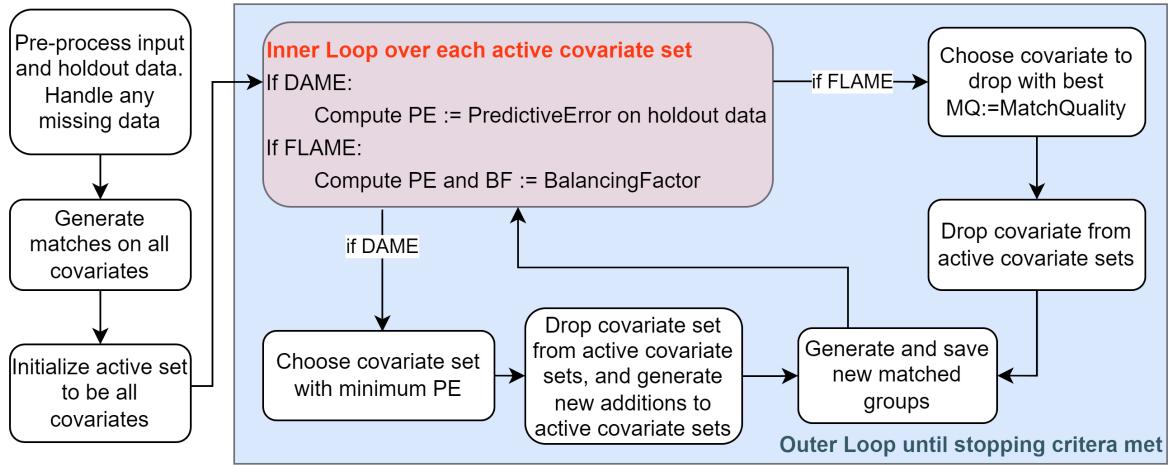


Figure 1: Flow diagram: DAME/FLAME algorithms.

As described above, after handling missing data, the algorithms begin by matching any units that can be matched exactly on all covariates, where at least one treatment unit as well as at least one control unit are contained in each matched group. The algorithms then execute the outer loop: updating sets of covariates to match on, referred to as *active covariate sets*, until a stopping criterion is reached.

In each iteration, the algorithms execute the inner loop, examining each covariate set to select the best one to match on. The *PE* denotes *predictive error*, as defined above, and the inner loop iteratively computes this for each active covariate set. DAME selects the covariate set minimizing *PE* and returns to the outer loop. FLAME determines the best covariate set by selecting the covariates yielding the highest *match quality* *MQ*, defined as $MQ = C \cdot BF - PE$, where C is a user-specified hyperparameter. The *balancing factor*, *BF*, measures the proportion of yet unmatched treatment units that are matched on a covariate set, summed with the proportion of yet unmatched control units that are matched on a covariate set. DAME iterates efficiently over covariate sets, prioritizing matching on large covariate sets if they can be used to effectively predict the outcome on the holdout training set. FLAME approximates this solution for scalability: it consistently matches on a smaller set of covariates than in the previous iteration, while still ensuring that each covariate set can be used to effectively predict the outcome on the holdout training set. It is important to note that *MQ* can be replaced by any non-adaptive, pre-determined, measure of quality that a user might be interested in.

Once all active covariate sets have been iteratively examined in the inner loop, the outer loop updates the active covariate sets and generates matches. Units that have identical values for all of the covariates that are part of the chosen covariate set form a matched group, as long as at least one treatment unit and at least one control unit are present in the group. Units that have identical values for all of the covariates that are part of the chosen covariate set form a matched group, as long as at least one treatment unit and at least one control unit are present in the group. The outer loop has a number of possible stopping criteria. It must stop when all units are placed in matched groups or all covariate sets have been dropped. Additionally, users can enforce stopping based on other criteria, e.g., (1) when there are too

few unmatched (treatment or control) units, (2) after a certain number of iterations, (3) when predictive error rises too much, or (4) when the balancing factor for a given round is not high enough. If the third criterion is chosen (when predictive error rises too much), then any units that have already been matched were matched on a set of covariates that together can be used to predict the outcome well.

1.3. Installation, setup, and testing

The package is designed for Python 3.6 and above. **dame-flame** depends on **scikit-learn** version 0.23.2 and above, **pandas** version 0.11.0 and above, and **numpy** version 1.16.5 (Pedregosa *et al.* 2011; The **pandas** Development Team 2020; Harris *et al.* 2020; Wes McKinney 2010).

dame-flame is available for download on the Python package index (PyPi) and on GitHub at <https://github.com/almost-matching-exactly/DAME-FLAME-Python-Package/>. The public documentation website (<https://almost-matching-exactly.github.io/DAME-FLAME-Python-Package/>) covers the API, installation instructions, a quick-start tutorial, several examples, and a contributing guide. In harmony with the best open source practices, users are invited to report any unexpected bugs, assist with cleaning or maintain code, add details or use-cases to the documentation, and add more test cases. They are welcome to do so via the GitHub bug reporting and pull requesting features, or by directly emailing the core development and research team.

There is also an accompanying R (R Core Team 2025) package **R-FLAME** for the FLAME and DAME algorithms, that can also be found on GitHub at <https://github.com/almost-matching-exactly/R-FLAME>. The R package is also an open source package welcoming user questions and contributions.

Testing was done to ensure that **R-FLAME** and the **dame-flame** Python package yield consistent results on a range of datasets and parameter options.

For further reliability testing, **dame-flame** offers continuous integration through Travis-CI, and the independent Coveralls API was used to verify the test suite offers an extensive code coverage.

2. Code and its explanation

2.1. Class structure and advanced features

The API consists of standard Pythonic design established by **scikit-learn**. The main feature of the API is the class `matching.MatchingParent`, and two subclasses `DAME` and `FLAME`. These offer standard methods `fit`, where a user provides a holdout training dataset, and `predict`, where a user provides a matching dataset, and the matches of the algorithm are computed.

Continuing to follow **scikit-learn**'s standards, any post-processing is done using `utils` functions, which take as arguments a `matching.DAME` or `matching.FLAME` object, and use these to compute matched groups of units, and estimate treatment effects, including the *average treatment effect* (ATE) for the population, and *conditional average treatment effect* (CATE) of a selected unit. These are mathematically defined in Section 2.2.

In the rest of this section, we proceed to discuss advanced features for computing predictive error, early stopping features, missing data handling options, and FLAME specific options.

For each of these advanced features, we discuss the theory, list the specific parameter for users to select, and the package default for this feature. We conclude this section by comparing these features to features of other popular matching packages.

2.2. Definitions of estimands and estimators

We continue with notation previously introduced: Units are indexed by i , which ranges from 1 to n . We may interchangeably refer to units as “individuals” or “observations”. Also, note that all units which contribute to treatment effect estimates must have been matched.

There are p pre-treatment covariates, and we will refer to the vector of covariates of unit i as \mathbf{x}_i . Let the binary treatment indicator for unit i be denoted by T_i . We let Y_i be the observed outcome for individual i where $Y_i = Y_i(1)T_i + Y_i(0)(1 - T_i)$ and $Y_i(0), Y_i(1)$ are the potential outcomes of unit i under control and treatment, respectively. Lastly, we introduce notation for matched groups, which we index by m , which ranges from 1 to M . The size of a matched group m is $\|m\|$, which is the number of units in the group. The number of treated and control units in the group are denoted $\|m_1\|$ and $\|m_0\|$ respectively.

The conditional average treatment effect, or CATE, is defined as the average treatment effect conditional on particular covariates. Formally, given a set of covariates X_i , the CATE is:

$$\text{CATE}(X_i) = \mathbb{E}[Y(1) - Y(0)|X_i].$$

Our implementation of CATE estimation allows users to input a unit i and receive its CATE estimate, based on its main matched group.

Since our units are matched almost-exactly, all units sharing the same main matched group will have the same CATE estimate. For a unit whose main matched group is group m , we estimate its CATE as:

$$\widehat{\text{CATE}}_m = \frac{1}{\|m_1\|} \sum_{i:T_i=1} [\hat{Y}_i(1)] - \frac{1}{\|m_0\|} \sum_{i:T_i=0} [\hat{Y}_i(0)]$$

where $\hat{Y}_i(0)$ is the observed outcome of units in m with treatment $T_i = 0$, and $\hat{Y}_i(1)$ is the observed outcome of the units in m with treatment $T_i = 1$.

The average treatment effect (ATE) is unconditional on X : $\text{ATE} = \mathbb{E}[Y(1) - Y(0)]$. We offer two estimators of the ATE. For one of the estimators, we will create a weighted combination of CATE estimates from the various matched groups. Let q_i denote the number of matched groups that unit i appears in. Note this quantity can be greater than 1 when matching is done with replacement, via the `repeats` argument. We then define the weight of a matched group m as $w_m = \sum_{i=1}^{\|m\|} q_i$. We estimate the ATE as:

$$\widehat{\text{ATE}} = \frac{\sum_m \widehat{\text{CATE}}_m \times w_m}{\sum_m w_m}.$$

The average treatment effect on treated (ATT) is conditioned on treated units, so $\text{ATT} = \mathbb{E}[Y(1) - Y(0)|T = 1]$. We estimate this as:

$$\widehat{\text{ATT}} = \frac{\sum_i Y_i T_i}{\sum_i T_i} - \frac{\sum_i Y_i q_i (1 - T_i)}{\sum_i q_i (1 - T_i)}.$$

This estimate directly uses the observed outcomes for treated units, while weighting the observed outcomes for control units on whether and how many times a control unit is placed in a matched group.

We offer a second ATE estimator, implementing the simple matching estimator described in [Abadie, Drukker, Herr, and Imbens \(2004\)](#). We also offer the variance estimator of this ATE estimate provided by [Abadie *et al.* \(2004\)](#). Please note that this estimator assumes constant treatment effects and homoscedasticity. It is not asymptotically normal, so users must make the standard implications on confidence intervals or hypothesis tests based on this.

2.3. Predictive error

As discussed in Section 1, the algorithm's decision of best covariates to match on relies on a computation of predictive error, or PE, based on a user-chosen machine learning algorithm run on the holdout dataset. The **dame-flame** Python package offers different options for the machine learning algorithm used, as well as a simplified FLAME and simplified DAME that does not use machine learning, but instead allows the users to input feature importance information for matching. We use **scikit-learn** for the underlying learning algorithms, and refer the reader to their documentation and references to learn more about these popular machine learning algorithms, as well as their specific implementations, applied separately to the treated and control units in the holdout set.

The options for the computation of PE can be chosen as follows.

- Ridge regression. Ridge regression minimizes a residual sum of squares plus a regularization term measuring the ℓ_2 norm of the coefficient vector, multiplied by a shrinkage parameter, α . For this option, a larger α should be chosen if it is believed that there is greater multicollinearity in the data, meaning that many covariates are linearly correlated. This option can be chosen using the parameter `adaptive_weights='ridge'`. The α parameter can also be adjusted using parameter `alpha` when declaring a matching object.
- Cross-validated ridge regression. This is a ridge regression with built-in cross-validation to determine the best regularization parameter. We use the **scikit-learn** `ridgeCV` class, but the default array of regularization strength options that we provide the function to iterate over is larger than the default provided by **scikit-learn**, for greater flexibility. This option is advantageous over the '`ridge`' option without cross-validation in the case when a user is uncertain about the regularization parameter, and a minor speed decrease owing to cross-validation is acceptable. This option can be chosen using the parameter `adaptive_weights='ridgeCV'`.
- Decision tree. Designed on a variation of the CART algorithm ([Breiman, Friedman, Olshen, and Stone 1984](#)), this is the only option that can be used for unordered discrete data. This option can be chosen using `adaptive_weights='decisiontree'` as described in Section 4.
- Cross-validated decision tree. A cross-validated version of the decision tree, iterating over a default array of α options. This option can be chosen using the parameter `adaptive_weights='decisiontreeCV'`.

- Any **scikit-learn** model. The user can choose any model that has a `fit` and `predict` method can be used. Once the model object is declared, and if given the name `object`, it can be chosen using the parameter `adaptive_weights='object'`.

The option a user chooses can be selected using the specified value for the parameter `adaptive_weights` when declaring a matching object, as shown in examples in Section 3. If, instead of allowing the algorithm to select covariates via the `PE` parameter, the user prefers to pre-specify covariate importance, they can do so by specifying `adaptive_weights = False`. The weights to the covariates in `input_data` can be specified using the parameter `weight_array` in the `fit` function. The values in that array must sum to 1.

2.4. Early-stopping options

The FLAME and DAME algorithms will stop after running to completion, or based on user-defined early stopping criteria. The default option is that the algorithm runs until all units are matched, or until there is a large spike in predictive error. If runtime or high accuracy of estimates of treatment effects are important, then we recommend users experiment with their stopping criteria based on their specific needs and dataset size. A large dataset will have a longer runtime, and an early stop will take less time. Without early stopping, the matches could degrade in quality in later iterations, where units that are farther from each other in covariate space would now be matched, leading to worse overall performance of the method. This is illustrated with simulated examples in [Wang *et al.* \(2021\)](#) and [Dieng *et al.* \(2019\)](#).

Below, we define and discuss the early stopping criteria that users can choose. All criteria are controlled via a parameter to the classes defined in Section 4.

- The maximum number of iterations of the FLAME or DAME algorithm, via the parameter `early_stop_iterations`. Iterations start at 0 so that a value of 0 leads to only exact matches being made. If FLAME is used, then this is the maximum number of covariates that can be dropped, meaning when the total number of covariates is m , no unit will be matched on fewer than $m - \text{early_stop_iterations}$ covariates. This is useful when the user wants only matches of a specific high level of quality, or when the user is concerned about computational time.
- Unmatched units in treatment or control, via the parameters `stop_unmatched_c` and `stop_unmatched_t`. When the algorithm is set with the `repeats=True` parameter, then previously matched units (that is, units whose main matched groups have already been determined) can still be placed in the main matched groups of other units. The algorithm will by default stop iterating when there are no more units that have not been placed in any group.
- Proportion of unmatched units, via the parameters `early_stop_un_c_frac`, and `early_stop_un_t_frac`. This stops the algorithm when the fraction of control units or treatment units that are unmatched goes below a user-defined value. One specific case in which this could be useful is where a user thinks that some percent of the input is unlikely to result in good matches.
- Predictive error, via the parameter `early_stop_pe`. The predictive error measures how important a covariate set is for predicting the outcome on the holdout training dataset,

using a machine learning algorithm. It is the sole determiner of the covariate set to match on for DAME, and one of two factors for FLAME. If FLAME or DAME attempts to drop a covariate set that would raise the predictive error above $(1 + \text{early_stop_pe})$ times the baseline predictive error (the predictive error when using all covariates), the algorithm terminates without dropping this covariate set.

2.5. Missing data handling options

Users are offered a variety of options for handling missing covariate data. Imputing missing values in datasets is possible, but matches become less interpretable when matching on imputed values, in that it is more difficult to discern why a match was recommended by the matching algorithm. Here, we discuss the options we provide in detail and make recommendations. The parameter to select in the **dame-flame** Python package is mentioned here, and more details on usage is provided in Section 4.

There can be missing data in either the input matching data, the holdout training data, or both. The specific character that is used to denote missing value can be selected via the parameter `missing_indicator`, which can be a character, integer, or `numpy.nan`.

For the input dataset, three options exist:

- Omit units with missing values. We recommend using this if missing values indicate data fidelity issues in a unit. The algorithms handle this by ensuring that units in the input dataset that have missing data are dropped from the dataset prior to running the algorithms finding the matches. This option is selected via the parameter `missing_data_replace=1`.
- Match units with missing values, but ignore missing values when considering which units to match to. We recommend this for the majority of cases. The underlying algorithm will handle this when pre-processing the input. This option is selected via the parameter `missing_data_replace=2`.
- Impute missing values with MICE. This is computationally costly and would reduce the interpretability of the matches. The algorithm would create several imputed datasets and iterate over each to find a match according to each dataset. This option is selected via the parameter `missing_data_replace=3`. The number of MICE imputations is selected via the parameter `missing_data_imputations`.

For the holdout dataset, the following two options exist:

- Omit units that have any missing values. We recommend this option only if a missing completely at random assumption is tenable in both holdout and matching datasets (Little 1988). In the underlying algorithm, units in the holdout dataset that have missing data are dropped from the dataset prior to running the DAME or FLAME algorithm to find the matches. This option is chosen by using the parameter `missing_holdout_replace=1`.
- Impute missing values with MICE. In the underlying algorithm, we begin by running MICE to create several imputed training holdout datasets. The DAME or FLAME algorithm is run once, and the best covariate set is chosen based on the predictive error over

all imputed datasets. This option is chosen by using the parameter `missing_holdout_replace=2`.

The underlying MICE implementation is done using **scikit-learn**'s experimental class **IterativeImputer**, and relies on decision tree regressions in the imputation process, to ensure that the data generated is fit for unordered categorical data.

2.6. Additional parameters available

As discussed in Section 2, users can adjust whether they match units with or without replacement. This is controlled via the boolean parameter `repeats`.

Output style can also be controlled by the user, via a range of parameters. All of these parameters are used when declaring a matching object.

- The parameter `verbose`. This is a number that will range from 0 to 3 and higher numbers result in additional information being output. If true, the output of the algorithm will include the predictive error of the covariate sets used for matching in each iteration.
- The boolean parameter `want_pe`. If true, the output will include the predictive error for each iteration.
- The boolean parameter `want_bf`. If true, the output will include the balancing factor for each iteration.

There are two FLAME specific parameters, which users would provide in the final `predict` step. These are:

- `C`, type float. This is the tradeoff parameter between the balancing factor and the predictive error when deciding which covariates to match on.
- `pre_dame`, type {integer}, default = `float('inf')`. The number of iterations to run FLAME before switching to DAME, allowing for a hybrid FLAME-DAME option.

2.7. Comparison to other matching packages

Many other matching methods either produce low-quality matches (leading to potentially poor treatment effect estimates), uninterpretable matches (e.g., in which matches include units with highly dissimilar covariates values), or matches that are manually defined by an analyst. One of the most widely used algorithms is nearest neighbor propensity score matching, provided by the R package **MatchIt** (Ho, Imai, King, and Stuart 2011). Propensity score matching reduces units' covariate information to one dimension, allowing matches to contain units even at extreme ends of the covariate space; such matches are uninterpretable. **MatchIt** allows other matching metrics, such as Mahalanobis distance, but does not allow for learning the proper metric as FLAME and DAME.

Another common matching algorithm is coarsened exact matching (CEM), popularly available in the R package **cem** (Iacus, King, and Porro 2009). CEM requires the user to manually coarsen variables, requiring humans to know detailed information about a high-dimensional

Language: Package	Built-in treatment-effect estimations	Built-in missing data handling	Provide matched groups	Handles continuous covariates
Python: dame-flame	Average, conditional	✓	✓	
Python: DoWhy	Average			✓
Python: PyMatch			✓	✓
R: cem	Average	✓	✓	✓
R: MatchIt	Average	✓	✓	✓
Propensity score				
Python: pymalts2	Average, conditional	✓	✓	✓

Table 1: Features of matching packages.

space in advance, a task at which humans are not naturally adept (Dieng *et al.* 2019). Coarsening covariates via default histogram binning methods fails to take into consideration their impacts on treatment and outcome, resulting in poor matches. Instead of requiring a human to manually input how matches should be constructed (or to use histogram binning), FLAME and DAME use machine learning on a training set to determine this information.

Many of the features described in Section 2.4 and Section 1.1 are unavailable in other matching packages. Table 1 compares the characteristics of **dame-flame** against popular alternatives. Most matching packages are implemented in R. R’s **cem** package only supports average treatment effect on the treated (ATT) treatment effects (Iacus *et al.* 2009). The **MatchIt** package focuses on estimation of average treatment effects and not conditional average treatment effects, both of which are handled in the same coherent manner by **dame-flame**. Users of any propensity score matching algorithm can adjust matched group sizes only by entering a ratio of treatment to control units, forcing all matched groups to be of the same size. Python’s **PyMatch** (Miroglia *et al.* 2017) and **DoWhy** (Sharma and Kiciman 2020) offer propensity score matching, but **DoWhy** does not emphasize matched groups, favoring to present treatment effects and other output. **DoWhy** uses the **EconML** package (Battocchi *et al.* 2019) to provide conditional average treatment effect estimates. R’s **cem** package is a good choice for datasets with multi-level, non-binary treatment variables, whereas the current version of **dame-flame** does not yet offer a multi-level treatment solution.

A further advantage of **dame-flame** is the higher quality of the matched groups generated by DAME and FLAME relative to propensity score matching, as shown by Dieng *et al.* (2019).

A drawback of **dame-flame** is the requirement that covariates be discrete. The packages **DoWhy**, **PyMatch**, **cem** and **MatchIt** do allow users to use continuous covariates without any pre-processing steps or manual binning. Although a user could manually bin continuous covariate values prior to using **dame-flame**, we do not recommend this asides scenarios in which users are confident they are binning variables in a way that is meaningful for their research. A user interested in a matching package that does allow for continuous covariates that is still in the *almost matching exactly* framework may consider exploring **R-MALTS** or **pymalts** (Parikh, Rudin, and Volfovsky 2019). These packages implement the algorithm *matching after learning to stretch* (MALTS), which will use exact matching for discrete variables, and will

learn Mahalanobis distances for continuous variables. Instead of a predetermined distance metric like **MatchIt**, MALTS gives covariates that contribute more towards predicting the outcome higher weights (Parikh, Rudin, and Volfovsky 2022).

3. Examples

3.1. Basic example

Here we offer an example to illustrate API usage, using a simple, small, 4 unit and 4 covariate simulated dataset to illustrate matched groups easily. An example focused on analysis using a real dataset and its corresponding replication is discussed in Section 3.2. All classes, functions, and parameters used here, as well as additional options for parameters are defined and discussed in Section 4.

The first step is importing the package. We show the dataframe used here as well. The **pandas** dataframe places units in rows and covariates in columns, and requires a column with a boolean variable indicating treatment, and a column for the outcome variable.

```
>>> import pandas as pd
>>> import dame_flame
>>> df = pd.DataFrame([[0,1,1,1,0,5.1], [0,0,1,0,0,5.11], [1,0,1,1,1,6.5],
...     [1,1,1,1,1,6.]], 
...     columns=["x1", "x2", "x3", "x4", "treated", "outcome"])
>>> print(df.head())

```

	x1	x2	x3	x4	treated	outcome
0	0	1	1	1	0	5.10
1	0	0	1	0	0	5.11
2	1	0	1	1	1	6.50
3	1	1	1	1	1	6.0

The first step in the matching procedure is instantiating a **matching** object, with optional parameters that can specify the early stopping criteria, missing data handling methodology, output style, or the machine learning method used to compute the predictive error. All optional parameters are described in more detail in Section 4. Here, we choose the default options, which includes no missing data handling, no early stopping procedures, and computes predictive error with ridge regression. We choose these options because this dataset does not have any missing data that needs to be handled and because it is a small example, we do not need to stop the algorithm early.

```
>>> model = dame_flame.matching.DAME()
```

The next step is to call the **fit** method on the **matching** object created above. Here, users must provide a file location of the holdout training dataset, a **pandas** dataframe, or a fraction of the input dataset to use for matching, in the parameter **holdout_data**.

Additionally, the name of the treatment column and the name of the outcome column can be provided.

```
>>> model.fit(df, "treated", "outcome")
```

At this point, simply calling the `predict` method with the input dataset produces matched results. The return value from the `predict` command contains an output table, which consists of the units that were matched to at least one other unit. For each unit that was matched, the table indicates which of the covariates were used for matching, and the covariate values that each unit was matched on. The covariates that were not used to match the unit are denoted with “*” as their value.

```
>>> result = model.predict(df)
>>> print(result)
```

	x1	x2	x3	x4
0	*	1	1	1
1	*	0	1	*
2	*	0	1	*
3	*	1	1	1

Various result summaries are available, including a printout of all matched groups, and the units belonging to each group. The result of the `predict` function, shown above, can also be retrieved by using the following attribute `df_units_and_covars_matched` of the matching class. The `units_per_group` attribute of the matching class provides an array of arrays. Each sub-array is a matched group, and each item in each sub-array is an integer, indicating the unit in that matched group. If matching is done with the parameter `repeats=False` when defining the matching class, then no unit will appear more than once. If `repeats=True` then the first group in which a unit appears is its main matched group.

```
>>> print(model.units_per_group)
```

```
[[0, 3], [1,2]]
```

This shows us that unit 0 and unit 3 are in a matched group, and that unit 1 and unit 2 are in another matched group.

The `utils` functions offer post-processing. In these functions, users must pass as parameters the matching object declared earlier, and for many of the functions, users must pass in a `unit_ids` parameter, which can be a single unit or a list of unit ids.

The function that provides matched groups of each unit is `MG`. If one unit id was provided, this is a single dataframe containing the main matched group of the unit. If the unit does not have a match, the return will be `numpy.nan`. If multiple unit ids were provided, this will be a list of dataframes with the main matched group of each unit provided.

```
>>> mmg = dame_flame.utils.post_processing.MG(matching_object=model, unit_ids=0)
>>> print(mmg)
```

	x1	x2	x3	x4	treated	outcome
0	*	1	1	1	0	5.1
3	*	1	1	1	1	6.

This shows the main matched group of unit 0 is unit 3, and that covariates that unit 0 and unit 3 matched on are covariates `x2`, `x3`, and `x4`.

The functions in the `utils` library also include treatment effect estimators, as defined in Section 2.2, including an estimate for CATE. If one unit id was provided, the return value will be a single float representing the conditional average treatment effect estimate of the unit. This is equal to the CATE of the group that the unit is in. If the unit does not have a match, the return will be `numpy.nan`. If multiple unit ids were provided, the return value will be a list of floats with the CATE estimate of each unit provided.

```
>>> cate = dame_flame.utils.post_processing.CATE(matching_object=model,
...     unit_ids=0)
>>> print(round(cate,3))
```

0.9

The ATE function, to get the ATE estimate only requires a matching object, but does not require a unit id, and returns a float.

```
>>> ate = dame_flame.utils.post_processing.ATE(matching_object=model)
>>> print(round(ate,3))
```

1.145

As discussed in Section 2.2, the package also offers a second ATE estimator with a corresponding variance estimator. Again, the required parameter is the matching object used earlier.

```
>>> var, ate = dame_flame.utils.post_processing.var_ATE(matching_object=model)
>>> print(round(ate))
>>> print(round(var))
```

1.145

0.03

As is expected, we see that this ATE estimate is the same as or close to the ATE estimate from the other ATE function.

3.2. Example analysis

dame-flame is an interpretable matching package because it allows users to quickly and easily understand which covariates were selected to be important for causal inference. This can be useful for practitioners in determining who benefits from treatment the most and where resources should be spent for future treatment. The package also allows users to view various other aspects of the matching process such as the stopping criteria as they use the package.

Here we demonstrate an experimental use-case for the DAME and FLAME algorithms on the Tennessee's student teachers achievement ratio (STAR) dataset. This dataset originates from

an experiment beginning in 1985, in which elementary school students and their teachers across 79 schools in Tennessee were randomly assigned to classes of small or regular sizes from kindergarten through third grade (Achilles *et al.* 2008). Although data is available for students not participating in the experiment, we limit to the experimental dataset in which treatment was randomized. The results showed that being placed in a small class size led to higher standardized test performance, and long term benefits in increased college entrance exam taking, especially among minority students (Krueger and Whitmore 2001).

Our cleaned dataset has around 5000 students. The treatment group is students in small class sizes, and the control group is students in regular sized classes with or without a teacher's aide. We use a transformed outcome variable by first computing an empirical CDF of math and reading scores for students in regular classes, then computing percentiles of math and reading scores for all students according to this distribution, and finally averaging these two. Our covariates include children's characteristics, teacher's characteristics and school characteristics. The children's characteristics are gender, race (binary, with white and Asian in one group, and all other races in the other group), free lunch status (students who received free lunch at any point in kindergarten through third grade are in one group), and age in months (binned into deciles). Age is typically binned in years, but is here binned in deciles of months since most children in a grade are close to the same age in years. The teacher characteristics include race, gender, and having a higher degree than bachelors. The school characteristics are urbanicity (rural, urban, suburban, and inner city) and a school identification number, with one for each of the 79 schools.

First, we ensure that using DAME and FLAME on this dataset is appropriate, by ensuring that there is a lack of sensitivity to the train/test split, which we take to be a random 80%/20% split. We do so by ensuring that the algorithm matches a sufficient number of units in each case and that the aggregate treatment effect estimates are reasonable. We first run four different trials of DAME on random splits. As we iterate over our four trials, we save the matching class objects for further analysis of treatment effects. When we declare an object of the matching class, we use the early stopping criteria of stopping when there is a significant increase of PE from the baseline PE computed in the first iteration. Then, we run the `fit` and `predict` methods on the matching class to run the match, as follows:

```
>>> models = []
>>> random_seeds = [1111, 2222, 3333, 4444]
>>> for i in range(len(random_seeds)):
...     matching_df, holdout_df = train_test_split(df_trunc, test_size=0.2,
...                                              random_state=random_seeds[i])
...     model_dame = dame_flame.matching.DAME(repeats=False, verbose=0,
...                                            adaptive_weights='decisiontree',
...                                            early_stop_pe=0.33)
...     model_dame.fit(holdout_data=holdout_df,
...                    outcome_column_name='outcome')
...     model_dame.predict(matching_df)
...     models.append(model_dame)
```

The output is:

```
Matching stopped while attempting iteration 28 due to the PE fraction early
stopping criterion.
```

```
Predictive error of covariate set would have been 92.12026755312186
Matching stopped while attempting iteration 20 due to the PE fraction early
stopping criterion.
```

```
Predictive error of covariate set would have been 105.55615864739742
Matching stopped while attempting iteration 16 due to the PE fraction early
stopping criterion.
```

```
Predictive error of covariate set would have been 126.2251752950491
Matching stopped while attempting iteration 20 due to the PE fraction early
stopping criterion.
```

```
Predictive error of covariate set would have been 64.27750359184846
```

Next, we compute the ATE estimate on each of the trials. We do this by iterating over the matching class objects declared above, and calling the `utils` function `var_ATE`, which is defined further in Section 4.

```
>>> for i in range(len(models)):
...     var, ate = dame_flame.utils.post_processing.var_ATE(models[i])
...     print("Trial", i, "matched", len(models[i].df_units_and_covars_matched),
...           "units with an ATE of", round(ate,2), "and a variance of ATE of",
...           round(var,2))
```

```
Trial 0 matched 1882 units with an ATE of 5.11 and a variance of ATE of 1.35
Trial 1 matched 1904 units with an ATE of 5.95 and a variance of ATE of 1.34
Trial 2 matched 1814 units with an ATE of 4.98 and a variance of ATE of 1.44
Trial 3 matched 1790 units with an ATE of 5.3 and a variance of ATE of 1.41
```

It is a good sign that each of the random test/train partitions behaved similarly in that they all matched a similar number of units and have similar ATE estimates. This indicates that the conclusions are not sensitive to the random partitioning. The ATE estimate and its variance estimate indicate that placing students in smaller class sizes caused those students to achieve higher kindergarten test scores by a few percentile points. We conclude this is a reasonable ATE estimate because in the analysis done in Krueger and Whitmore (2001), the estimate for the impact of small class sizes provided by Figure 1 in their work, which is based on a linear regression on many similar covariates, is between 5 and 6 for average percentile of math and reading scores.

Next, we consider whether DAME or FLAME is a better matching method for this dataset. We run four trials of FLAME with the same random test/train split. Again, we define an object of the matching class, this time of the FLAME subclass, and we run the `fit` and `predict` methods.

```
>>> flame_models = []
>>> random_seeds = [1111, 2222, 3333, 4444]
>>> for i in range(4):
...     matching_df, holdout_df = train_test_split(df_trunc, test_size=0.2,
...                                               random_state=random_seeds[i])
...     model_flame = dame_flame.matching.FLAME(
...         repeats=False, verbose=3, adaptive_weights='decisiontree',
...         early_stop_pe=0.33)
```

Trial 1	Trial 2	Trial 3	Trial 4
All covariates, 1079	All covariates, 1053	All covariates, 1047	All covariates, 1056
Teacher gender, 0	Teacher gender, 0	Teacher gender, 0	Teacher gender, 0
Urbanicity, 0	Urbanicity, 0	Urbanicity, 0	Urbanicity, 0
Teacher race, 198	Student race, 53	Teacher race, 209	Student race, 54
Student race, 269	Teacher race, 265	Student race, 282	Teacher race, 252

Table 2: Summary of covariate dropping order from FLAME.

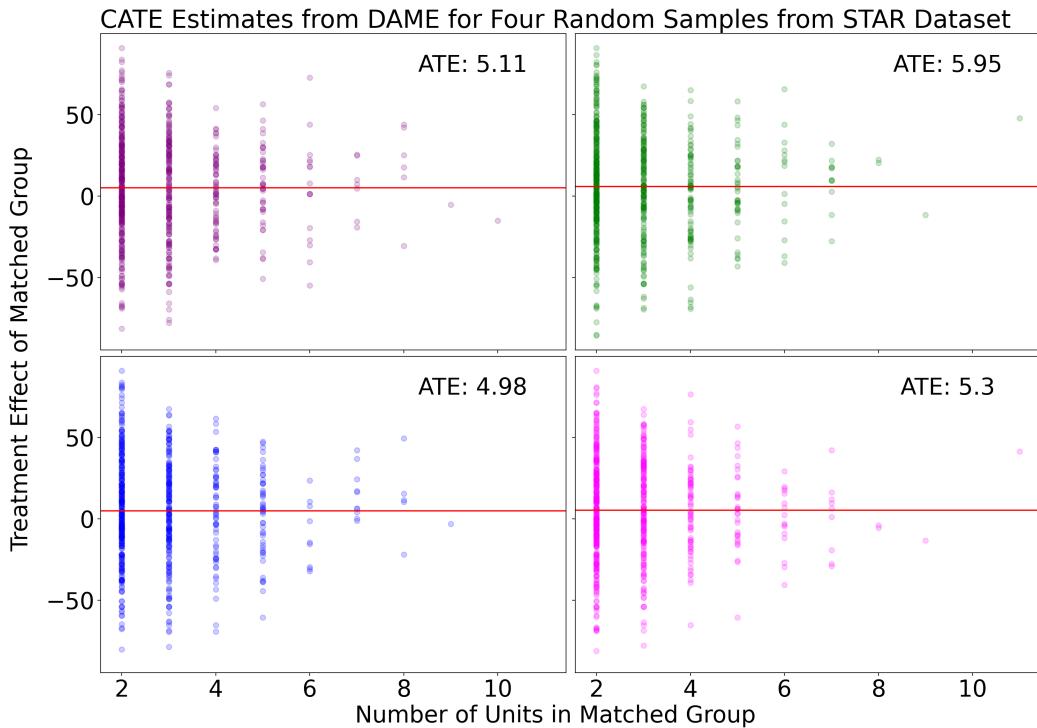


Figure 2: CATE estimates of matched groups, across four random splits of the dataset into matching and holdout training datasets. The vertical line is the ATE estimate.

```
...     model_flame.fit(holdout_data=holdout_df, outcome_column_name='outcome')
...     result_flame = model_flame.predict(matching_df)
...     flame_models.append(model_flame)
```

We omit the output of this matching result, which is available with the full replication code in the supplementary materials or on our GitHub repository. We observe that the FLAME algorithm iterates fewer times than the DAME algorithm and that both stopped iterating according to the stopping criterion `early_stop_pe`. Using the `var_ATE` function on the FLAME matching objects yields similar estimates as when using DAME.

We summarize the dropping criteria from the full verbose output of FLAME, generated using the parameter `verbose=3` in Table 2. Each entry is a iteration of the algorithm, listing the

covariate dropped in that iteration, and the number of units matched after that covariate is dropped. The fact that teacher gender is dropped first in each trial makes sense, since all teachers are female in the cleaned data set. In Krueger and Whitmore (2001), teacher gender and teacher race covariates were not used in their linear regression estimating the impact of small class sizes on percentile test scores for kindergarten students.

We plot the CATE estimates of the matched groups from DAME against the number of units of each group in Figure 2. We omit the code to create this plot (available on the GitHub repository and in the replication materials accompanying this paper) as it highlights Python semantics not specific to **dame-flame**. This plot highlights heterogeneity in the treatment effects of groups. Across the trials, when we examine the matched groups that correspond to the largest group sizes, or the rightmost points on the graphs, we notice that the large matched group contains different units in each trial, so it is sensitive to the test/train split.

Overall, because the ATE estimates using the DAME and the FLAME algorithms correspond to previous findings, **dame-flame** proves itself in this situation to be a valuable robustness check for researchers who wish to verify measurable impacts of this well known experiment. We hope that **dame-flame** can additionally be of use to researchers interested in matching and exploring matched groups and CATE estimates of matched groups.

4. API documentation

In this section, we provide details on the matching class definitions and function parameters. The API consists of standard Pythonic design established by **scikit-learn** (Buitinck *et al.* 2013). To begin matching, the user declares an object of type **DAME** or type **FLAME**. Both of these inherit from the base class **MatchingParent**. Following this, the user calls the method **fit**, providing a holdout training dataset, and finally **predict**, where a user provides a matching dataset, and the matches of the algorithm are computed.

The user provides the input data as a data frame or file in the function **predict**, which must contain an outcome column, and a treatment column. FLAME and DAME by default produce an output table consisting of the units that were matched to at least one other unit. For each unit that was matched, the table indicates which of the covariates were used for matching, and the covariate values that each unit was matched on. The covariates that were not used to match the unit are denoted with “*” as their values.

Class definition

```
dame_flame.matching.DAME(adaptive_weights='ridge', alpha=0.1,
    repeats=True, verbose=2, early_stop_iterations=float('inf'),
    stop_unmatched_c=False, early_stop_un_c_frac=False,
    stop_unmatched_t=False, early_stop_un_t_frac=False,
    early_stop_pe=0.05, missing_indicator=numpy.nan,
    missing_data_replace=0, missing_holdout_replace=0,
    missing_holdout_imputations=10,
    missing_data_imputations=1, want_pe=False, want_bf=False)

dame_flame.matching.FLAME(adaptive_weights='ridge', alpha=0.1,
    repeats=True, verbose=2, early_stop_iterations=float('inf'),
```

```
stop_unmatched_c=False, early_stop_un_c_frac=False,
stop_unmatched_t=False, early_stop_un_t_frac=False,
early_stop_pe=0.05, missing_indicator=numpy.nan,
missing_data_replace=0, missing_holdout_replace=0,
missing_holdout_imputations=10,
missing_data_imputations=1, want_pe=False, want_bf=False)
```

Arguments:

- **adaptive_weights**, type {bool, 'ridge', 'decisiontree', 'ridgeCV', 'decisiontreeCV', **scikit-learn** model}, default = 'ridge'. The method used to decide what covariate set should be dropped next.
- **alpha**, type float, default = 0.1. If **adaptive_weights** is set to 'ridge', this is the α for ridge regression.
- **repeats**, type bool, default = True. Whether or not units for whom a main matched has been found can be used again, and placed in an auxiliary matched group.
- **verbose**, type int: {0, 1, 2, 3}, default = 2. Style of printout while algorithm runs. If 0, no output. If 1, provides iteration number. If 2, provides iteration number and additional information on the progress of the matching at every 10th iteration. If 3, provides iteration number and additional information on the progress of the matching at every iteration.
- **early_stop_iterations**, type float, int, default = `float('inf')`. A number of iterations after which to hard stop the algorithm.
- **stop_unmatched_c**, type bool, default = False. If True, then the algorithm terminates when there are no more control units to match.
- **stop_unmatched_t**, type bool, default = False. If True, then the algorithm terminates when there are no more treatment units to match.
- **early_stop_un_c_frac**, type float, default = 0.1. Must be between 0.0 and 1.0. This provides a fraction of unmatched control units. When the threshold is met, the algorithm will stop iterating. For example, using an input dataset with 100 control units, the algorithm will stop when 10 control units are unmatched and 90 are matched (or earlier, depending on other stopping conditions).
- **early_stop_un_t_frac**, type float, default = 0.1. Must be between 0.0 and 1.0. This provides a fraction of unmatched treatment units. When the threshold is met, the algorithm will stop iterating. For example, using an input dataset with 100 treatment units, the algorithm will stop when 10 control units are unmatched and 90 are matched (or earlier, depending on other stopping conditions).
- **early_stop_pe**, type float, default = 0.05. If the algorithm attempts to drop a covariate set that would lead to a predictive error above $(1 + \text{early_stop_pe})$ times the baseline predictive error (the predictive error when using all covariates to predict), then the algorithm terminates before dropping this covariate set.

- **want_pe**, type bool, default = `False`. If `True`, the output of the algorithm will include the predictive error of the covariate sets used for matching in each iteration.
- **want_bf**, type bool, default = `False`. If `True`, the output will include the balancing factor for each iteration.
- **missing_indicator**, type {character, integer, `numpy.nan`}, default = `numpy.nan`. This is the indicator for missing data in the dataset.
- **missing_holdout_replace**, type int: {0, 1, 2}, default = 0. If 0, assume no missing holdout training data and proceed. If 1, the algorithm excludes units with missing values from the holdout dataset. If 2, do MICE on holdout dataset. If this option is selected, it will be done for a number of iterations equal to `missing_holdout_imputations`.
- **missing_data_replace**, type int: {0, 1, 2, 3}, default = 0. If 0, assume no missing data in matching data and proceed. If 1, the algorithm does not match on units that have missing values. If 2, prevent all `missing_indicator` values from being matched on. If 3, do MICE on matching dataset. This is not recommended. If this option is selected, it will be done for a number of iterations equal to `missing_data_imputations`.
- **missing_holdout_imputations**, type int, default = 10. If `missing_holdout_replace`=2, this is the number of imputations.
- **missing_data_imputations**, type int, default = 1. If `missing_data_replace`=3, this is the number of imputations.

fit function

```
fit(self, holdout_data=False, treatment_column_name='treated',
     outcome_column_name='outcome', weight_array=False)
```

Arguments:

- **holdout_data**, type {string, dataframe, float, `False`}, default = `False`. This is the holdout training dataset. If a string is given, that should be the location of a CSV file to input. If a float between 0.0 and 1.0 is given, that corresponds the percent of the input dataset to randomly select for holdout data. If `False`, the holdout data is equal to the entire input data. If users choose to use units repeatedly in both the holdout and training dataset, they should be careful that the data do not have a special situation that needs to be respected in subsampling such as a hierarchy.
- **treatment_column_name**, type string, default = "treated". This is the name of the column with a binary indicator for whether a row is a treatment or control unit.
- **outcome_column_name**, type string, default = "outcome". This is the name of the column with the outcome variable of each unit.
- **weight_array**, type array, optional. If `adaptive_weights=False`, these are the weights to the covariates in `input_data`, for the non adaptive version of DAME. Must sum to 1. In this case, we do not use machine learning for the weights, they are manually entered as `weight_array`.

predict function

```
predict(self, input_data)
```

Argument for both FLAME and DAME objects:

- `input_data`, type {string, dataframe}, required parameter. The dataframe on which to perform the matching, or the location of the CSV with the dataframe.

Arguments for FLAME object only:

- `C`, type float, default = 0.1. The tradeoff parameter between the balancing factor and the predictive error when deciding which covariates to match on.
- `pre_dame`, type float, integer, default = `float('inf')`. The number of iterations to run FLAME before switching to DAME, allowing for a hybrid FLAME-DAME option.

Return values:

- **Result.** `pandas` dataframe of matched units and covariates matched on, with a “*” at each covariate that a unit did not use in matching.

Matching class attributes

- `units_per_group`, type array. This is an array of arrays. Each sub-array is a matched group, and each item in each sub-array is an `int`, indicating the unit in that matched group. If matching is done with `repeats=False` when defining the `DAME` or `FLAME` matching object, then no unit will appear more than once. If `repeats=True` then the first group in which a unit appears is its main matched group.
- `df_units_and_covars_matched`, type datafarme. This are the resulting matches of DAME. Each matched unit is in this array, and the covariates they were matched on have the value used to match. The covariates units that were not matched on are indicated with a “*”.
- `groups_per_unit`, type array. The length of this is equal to the number of units in the input array. Each item in this array corresponds to the number of times that each item was matched. If matching is done with `repeats=False` when defining the matching objects, then this number will be either 0 or 1.
- `bf_each_iter`, type array. If argument `want_bf=True` when defining the matching class, this will contain the balancing factor of the chosen covariate set at each iteration.
- `pe_each_iter`, type array. If argument `want_pe=True` when defining the matching class, this will contain the predictive error of the chosen covariate set at each iteration.

Post-processing utils

```
dame_flame.utils.post_processing.MG(matching_object, unit_ids,
    output_style=1, mice_iter=0)
dame_flame.utils.post_processing.all_MGs(matching_object)
dame_flame.utils.post_processing.ATE(matching_object, mice_iter=0)
dame_flame.utils.post_processing.CATE(matching_object, unit_ids,
    mice_iter=0)
dame_flame.utils.post_processing.ATT(matching_object, mice_iter=0)
dame_flame.utils.post_processing.var_ATE(matching_object)
```

Arguments:

- **matching_object**, type `dame_flame.matching.DAME`, `dame_flame.matching.FLAME`. The matching object used to run DAME and FLAME, after the `fit` and `predict` methods have been called to create the matches. If the `matching_object`'s parameter for `verbose` is not 0, then, as units without matches appear, the function will print that those units did not have matches.
- **unit_ids**, type {int, list}. A unit id or list of unit ids.
- **output_style**, type int: {0,1}. Default = 1. If 1, the covariates which were not used in matching for the group of the unit will have a “*” rather than the covariate value. Otherwise, it will output all covariate values.
- **mice_iter**, type int. Default = 0. If matching was done using MICE, this is the iteration of MICE for which the treatment effects or matched groups will be found.

Return values and corresponding functions:

- **MG**: type {list, dataframe, `numpy.nan`}. Returns matched groups corresponding to inputs' unit ids. If one unit id was provided, this is a single dataframe containing the main matched group of the unit. If the unit does not have a match, the return will be `numpy.nan`. If multiple unit ids were provided, this will be a list of dataframes with the main matched group of each unit provided. If any unit does not have a match, rather than a dataframe, at its place will be `numpy.nan`.
- **all_MGs**: type dictionary. Returns all matched groups, in a dictionary mapping a unit id that was matched to a list of unit ids that are in the main matched group of that unit.
- **ATE**: type {float, `numpy.nan`}. Returns the average treatment effect (ATE) estimate in a float. If no units were matched, then the output will be `numpy.nan`.
- **CATE**: type {list, float, `numpy.nan`}. Returns the conditional average treatment effects (CATE) estimates. If one unit id was provided, this is a single float representing the estimate of the conditional average treatment effect of the unit. This is equal to the CATE estimate of the group that the unit is in. If the unit does not have a match, the return will be `numpy.nan`. If multiple unit ids were provided, this will be a list of floats with the CATE of each unit provided. If any unit does not have a match, rather than a float within the list, at its place will be `numpy.nan`.

- **ATT**: type {float, `numpy.nan`}. Returns the average treatment effect on treated (ATT) estimate in a float. If no units were matched, then the output will be `numpy.nan`.
- **var_ATE**: Tuple with two items. The first is the variance of ATE estimate in type float. The second is the ATE estimate in type float. These estimates are computed using the methodology in [Abadie *et al.* \(2004\)](#).

5. Conclusions

The FLAME and DAME algorithms for matching of observational data with discrete covariates provide interpretable and high-quality matches. The **dame-flame** open-source Python package offers efficient, easy-to-use implementations of these algorithms. The package is easily accessible, and here, we provide detailed documentation, with concrete examples. The package is written in a highly modular manner, facilitating the introduction of new features and variations of the DAME and FLAME algorithms. It is available at <https://github.com/almost-matching-exactly/DAME-FLAME-Python-Package>. We believe this package will be a useful tool for social science researchers, health researchers, and other scientists that use matching.

Acknowledgments

This work was supported in part by awards NIH R01EB025021, NSF IIS-1703431, and the Duke University Energy Initiative Energy Research Seed Fund.

References

Abadie A, Drukker D, Herr JL, Imbens GW (2004). “Implementing Matching Estimators for Average Treatment Effects in Stata.” *The Stata Journal*, 4(3), 290–311. [doi:10.1177/1536867X0400400307](https://doi.org/10.1177/1536867X0400400307).

Achilles CM, Bain HP, Bellott F, Boyd-Zaharias J, Finn J, Folger J, Johnston J, Word E (2008). “Tennessee’s Student Teacher Achievement Ratio (STAR) Project.” [doi:10.7910/DVN/SIWH9F](https://doi.org/10.7910/DVN/SIWH9F).

Battocchi K, Dillon E, Hei M, Lewis G, Oka P, Oprescu M, Syrgkanis V (2019). “**EconML**: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation.” Python package version 0.15.1, URL <https://github.com/microsoft/EconML>.

Breiman L, Friedman J, Olshen RA, Stone CJ (1984). *Classification and Regression Trees*. 1st edition. Chapman & Hall/CRC. [doi:10.1201/9781315139470](https://doi.org/10.1201/9781315139470).

Buitinck L, Louppe G, Blondel M, Pedregosa F, Mueller A, Grisel O, Niculae V, Prettenhofer P, Gramfort A, Grobler J, *et al.* (2013). “API Design for Machine Learning Software: Experiences From the **scikit-learn** Project.” *arXiv 1309.0238*, arXiv.org E-Print Archive. [doi:10.48550/arXiv.1309.0238](https://doi.org/10.48550/arXiv.1309.0238).

Dieng A, Liu Y, Roy S, Rudin C, Volfovsky A (2019). “Interpretable Almost-Exact Matching for Causal Inference.” In *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019*, pp. 2445–2453. Naha, Okinawa. URL <http://proceedings.mlr.press/v89/dieng19b/dieng19b.pdf>.

Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, Wieser E, Taylor J, Berg S, Smith NJ, Kern R, Picus M, Hoyer S, van Kerkwijk MH, Brett M, Haldane A, del Río JF, Wiebe M, Peterson P, Gérard-Marchant P, Sheppard K, Reddy T, Weckesser W, Abbasi H, Gohlke C, Oliphant TE (2020). “Array Programming with **NumPy**.” *Nature*, **585**(7825), 357–362. doi:10.1038/s41586-020-2649-2.

Ho D, Imai K, King G, Stuart EA (2011). “**MatchIt**: Nonparametric Preprocessing for Parametric Causal Inference.” *Journal of Statistical Software*, **42**, 1–28. doi:10.18637/jss.v042.i08.

Iacus S, King G, Porro G (2009). “**cem**: Software for Coarsened Exact Matching.” *Journal of Statistical Software*, **30**, 1–27. doi:10.18637/jss.v030.i09.

Krueger AB, Whitmore DM (2001). “The Effect of Attending a Small Class in the Early Grades on College-Test Taking and Middle School Test Results: Evidence from Project STAR.” *The Economic Journal*, **111**(468), 1–28. doi:10.1111/1468-0297.00586.

Little RJA (1988). “A Test of Missing Completely at Random for Multivariate Data with Missing Values.” *Journal of the American Statistical Association*, **83**(404), 1198–1202. doi:10.2307/2290157.

Miroglio B, *et al.* (2017). “**pymatch**.” URL <https://github.com/benmiroglio/pymatch>.

Parikh H, Rudin C, Volfovsky A (2019). “**PyMALTS** and MALTS Implementation.” URL <https://github.com/almost-matching-exactly/MALTS>.

Parikh H, Rudin C, Volfovsky A (2022). “MALTS: Matching after Learning to Stretch.” *Journal of Machine Learning Research*, **23**(1), 10952–10993. URL <https://www.jmlr.org/papers/volume23/21-0053/21-0053.pdf>.

Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay E (2011). “**scikit-learn**: Machine Learning in Python.” *Journal of Machine Learning Research*, **12**, 2825–2830. URL <https://www.jmlr.org/papers/v12/pedregosa11a.html>.

R Core Team (2025). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.

Sharma A, Kiciman E (2020). “**DoWhy**: An End-to-End Library for Causal Inference.” *arXiv 2011.04216*, arXiv.org E-Print Archive. doi:10.48550/arXiv.2011.04216.

The **pandas** Development Team (2020). “pandas-dev/pandas: **Pandas**.” doi:10.5281/zenodo.3509134.

Van Buuren S, Groothuis-Oudshoorn K (2010). “**mice**: Multivariate Imputation by Chained Equations in R.” *Journal of Statistical Software*, **45**, 1–67. doi:10.18637/jss.v045.i03.

Van Rossum G, *et al.* (2011). *Python Programming Language*. URL <http://www.python.org/>.

Wang T, Morucci M, Awan MU, Liu Y, Roy S, Rudin C, Volfovsky A (2021). “FLAME: A Fast Large-Scale Almost Matching Exactly Approach to Causal Inference.” *Journal of Machine Learning Research*, **22**(1), 1477–1517. URL <https://www.jmlr.org/papers/v22/19-853.html>.

Wes McKinney (2010). “Data Structures for Statistical Computing in Python.” In S Van der Walt, J Millman (eds.), *Proceedings of the 9th Python in Science Conference*, pp. 56–61. doi:10.25080/Majora-92bf1922-00a.

Affiliation:

Neha R. Gupta
 Carnegie Mellon University
 5000 Forbes Avenue, Pittsburgh, PA 15213, United States of America
 E-mail: nehargupta@cmu.edu

Sudeepa Roy
 Department of Computer Science
 Duke University
 Durham, North Carolina, United States of America
 E-mail: sudeepa@cs.duke.edu

Cynthia Rudin
 Department of Computer Science, Department of Electrical and Computer Engineering
 Duke University
 Durham, North Carolina, United States of America
 E-mail: cynthia@cs.duke.edu

Alexander Volfovsky
 Department of Statistical Science
 Duke University
 Durham, North Carolina, United States of America
 E-mail: alexander.volfovsky@duke.edu