
LensKit Documentation

Release 0.6.1

Michael D. Ekstrand

Apr 03, 2019

CONTENTS:

1	Installation	3
2	Resources	5
2.1	Getting Started	5
2.2	Algorithm Interfaces	8
2.3	Crossfold preparation	11
2.4	Batch-Running Recommenders	14
2.5	Evaluating Recommender Output	18
2.6	Errors and Diagnostics	22
2.7	Algorithms	23
2.8	Utility Functions	32
2.9	Release Notes	37
3	Indices and tables	39
	Bibliography	41
	Python Module Index	43

LensKit is a set of Python tools for experimenting with and studying recommender systems. It provides support for training, running, and evaluating recommender algorithms in a flexible fashion suitable for research and education.

LensKit for Python (also known as LKPY) is the successor to the Java-based LensKit project.

INSTALLATION

To install the current release with Anaconda (recommended):

```
conda install -c lenskit lenskit
```

Or you can use pip:

```
pip install lenskit
```

To use the latest development version, install directly from GitHub:

```
pip install git+https://github.com/lenskit/lkpy
```

Then see [Getting Started](#).

RESOURCES

- [Mailing list, etc.](#)
- [Source and issues on GitHub](#)

2.1 Getting Started

This notebook gets you started with a brief nDCG evaluation with LensKit for Python.

2.1.1 Setup

We first import the LensKit components we need:

```
[1]: from lenskit import batch, topn, util
      from lenskit import crossfold as xf
      from lenskit.algorithms import Recommender, als, item_knn as knn
      from lenskit import topn
```

And Pandas is very useful:

```
[2]: import pandas as pd
```

```
[3]: %matplotlib inline
```

2.1.2 Loading Data

We're going to use the ML-100K data set:

```
[4]: ratings = pd.read_csv('ml-100k/u.data', sep='\t',
                          names=['user', 'item', 'rating', 'timestamp'])
      ratings.head()
```

```
[4]:
```

	user	item	rating	timestamp
0	196	242	3	881250949
1	186	302	3	891717742
2	22	377	1	878887116
3	244	51	2	880606923
4	166	346	1	886397596

2.1.3 Defining Algorithms

Let's set up two algorithms:

```
[5]: algo_ii = knn.ItemItem(20)
     algo_als = als.BiasedMF(50)
```

2.1.4 Running the Evaluation

In LensKit, our evaluation proceeds in 2 steps:

1. Generate recommendations
2. Measure them

If memory is a concern, we can measure while generating, but we will not do that for now.

We will first define a function to generate recommendations from one algorithm over a single partition of the data set. It will take an algorithm, a train set, and a test set, and return the recommendations.

Note: before fitting the algorithm, we clone it. Some algorithms misbehave when fit multiple times.

Note 2: our algorithms do not necessarily implement the `Recommend` interface, so we adapt them. This fills in a default candidate selector.

The code function looks like this:

```
[6]: def eval(aname, algo, train, test):
     fittable = util.clone(algo)
     fittable = Recommender.adapt(fittable)
     fittable.fit(train)
     users = test.user.unique()
     # now we run the recommender
     recs = batch.recommend(fittable, users, 100)
     # add the algorithm name for analyzability
     recs['Algorithm'] = aname
     return recs
```

Now, we will loop over the data and the algorithms, and generate recommendations:

```
[7]: all_recs = []
     test_data = []
     for train, test in xf.partition_users(ratings[['user', 'item', 'rating']], 5, xf.
     ↪SampleFrac(0.2)):
         test_data.append(test)
         all_recs.append(eval('ItemItem', algo_ii, train, test))
         all_recs.append(eval('ALS', algo_als, train, test))
```

With the results in place, we can concatenate them into a single data frame:

```
[8]: all_recs = pd.concat(all_recs, ignore_index=True)
     all_recs.head()
```

```
[8]:   item  score  user  rank Algorithm
0   285  4.543364    5    1  ItemItem
1  1449  4.532999    5    2  ItemItem
2  1251  4.494639    5    3  ItemItem
3   114  4.479512    5    4  ItemItem
4   166  4.399639    5    5  ItemItem
```

To compute our analysis, we also need to concatenate the test data into a single frame:

```
[9]: test_data = pd.concat(test_data, ignore_index=True)
```

We analyze our recommendation lists with a `RecListAnalysis`. It takes care of the hard work of making sure that the truth data (our test data) and the recommendations line up properly.

We do assume here that each user only appears once per algorithm. Since our crossfold method partitions users, this is fine.

```
[10]: rla = topn.RecListAnalysis()
      rla.add_metric(topn.ndcg)
      results = rla.compute(all_recs, test_data)
      results.head()
```

```
/home/MICHAELEKSTRAND/anaconda3/envs/lkpy-dev/lib/python3.7/site-packages/pandas/core/
↳ indexing.py:1494: PerformanceWarning: indexing past lexsort depth may impact
↳ performance.
      return self._getitem_tuple(key)
```

```
[10]:
```

		ndcg
1	ALS	0.265268
	ItemItem	0.259708
2	ALS	0.148335
	ItemItem	0.081890
3	ALS	0.026615


Now we have nDCG values!

```
[11]: results.groupby('Algorithm').ndcg.mean()
```

```
[11]: Algorithm
      ALS      0.139689
      ItemItem  0.102075
      Name: ndcg, dtype: float64
```

```
[12]: results.groupby('Algorithm').ndcg.mean().plot.bar()
```

```
[12]: <matplotlib.axes._subplots.AxesSubplot at 0x7f03842f8860>
```



[]:

2.2 Algorithm Interfaces

LKPY's batch routines and utility support for managing algorithms expect algorithms to implement consistent interfaces. This page describes those interfaces.

The interfaces are realized as abstract base classes with the Python `abc` module. Implementations must be registered with their interfaces, either by subclassing the interface or by calling `abc.ABCMeta.register()`.

2.2.1 Base Algorithm

Algorithms follow the SciKit fit-predict paradigm for estimators, except they know natively how to work with Pandas objects.

The *Algorithm* interface defines common methods.

class `lenskit.algorithms.Algorithm`

Base class for LensKit algorithms. These algorithms follow the SciKit design pattern for estimators.

fit (*ratings*, **args*, ***kwargs*)

Train a model using the specified ratings (or similar) data.

Parameters

- **ratings** (*pandas.DataFrame*) – The ratings data.
- **args** – Additional training data the algorithm may require.
- **kwargs** – Additional training data the algorithm may require.

Returns The algorithm object.

get_params (*deep=True*)

Get the parameters for this algorithm (as in scikit-learn). Algorithm parameters should match constructor argument names.

The default implementation returns all attributes that match a constructor parameter name. It should be compatible with `scikit.base.BaseEstimator.get_params()` method so that LensKit algorithms can be cloned with `scikit.base.clone()` as well as `lenskit.util.clone()`.

Returns the algorithm parameters.

Return type `dict`

2.2.2 Recommendation

The *Recommender* interface provides an interface to generating recommendations. Not all algorithms implement it; call *Recommender.adapt()* on an algorithm to get a recommender for any algorithm that at least implements *Predictor*. For example:

```
pred = Bias(damping=5)
rec = Recommender.adapt(pred)
```

Note: We are rethinking the ergonomics of this interface, and it may change in LensKit 0.6. We expect keep compatibility in the `lenskit.batch.recommend()` API, though.

class `lenskit.algorithms.Recommender`

Recommends lists of items for users.

classmethod `adapt` (*algo*)

Ensure that an algorithm is a *Recommender*. If it is not a recommender, it is wrapped in a `lenskit.basic.TopN` with a default candidate selector.

Note: Since 0.6.0, since algorithms are fit directly, you should call this method **before** calling *Algorithm.fit()*, unless you will always be passing explicit candidate sets to *recommend()*.

Parameters `algo` (*Predictor*) – the underlying rating predictor.

recommend (*user, n=None, candidates=None, ratings=None*)

Compute recommendations for a user.

Parameters

- **user** – the user ID
- **n** (*int*) – the number of recommendations to produce (*None* for unlimited)
- **candidates** (*array-like*) – The set of valid candidate items; if *None*, a default set will be used. For many algorithms, this is their *CandidateSelector*.
- **ratings** (*pandas.Series*) – the user’s ratings (indexed by item id); if provided, they may be used to override or augment the model’s notion of a user’s preferences.

Returns a frame with an `item` column; if the recommender also produces scores, they will be in a `score` column.

Return type `pandas.DataFrame`

2.2.3 Candidate Selection

Some recommenders use a *candidate selector* to identify possible items to recommend. These are also treated as algorithms, mainly so that they can memorize users' prior ratings to exclude them from recommendation.

class `lenskit.algorithms.CandidateSelector`

Select candidates for recommendation for a user, possibly with some additional ratings.

candidates (*user*, *ratings=None*)

Select candidates for the user.

Parameters

- **user** – The user key or ID.
- **ratings** (*pandas.Series* or *array-like*) – Ratings or items to use instead of whatever ratings were memorized for this user. If a `pandas.Series`, the series index is used; if it is another array-like it is assumed to be an array of items.

static rated_items (*ratings*)

Utility function for converting a series or array into an array of item IDs. Useful in implementations of `candidates()`.

2.2.4 Rating Prediction

class `lenskit.algorithms.Predictor`

Predicts user ratings of items. Predictions are really estimates of the user's like or dislike, and the `Predictor` interface makes no guarantees about their scale or granularity.

predict (*pairs*, *ratings=None*)

Compute predictions for user-item pairs. This method is designed to be compatible with the general SciKit paradigm; applications typically want to use `predict_for_user()`.

Parameters

- **pairs** (*pandas.DataFrame*) – The user-item pairs, as `user` and `item` columns.
- **ratings** (*pandas.DataFrame*) – user-item rating data to replace memorized data.

Returns The predicted scores for each user-item pair.

Return type `pandas.Series`

predict_for_user (*user*, *items*, *ratings=None*)

Compute predictions for a user and items.

Parameters

- **user** – the user ID
- **items** (*array-like*) – the items to predict
- **ratings** (*pandas.Series*) – the user's ratings (indexed by item id); if provided, they may be used to override or augment the model's notion of a user's preferences.

Returns scores for the items, indexed by item id.

Return type `pandas.Series`

2.3 Crossfold preparation

The LKPY *crossfold* module provides support for preparing data sets for cross-validation. Crossfold methods are implemented as functions that operate on data frames and return generators of (*train*, *test*) pairs (*lenskit.crossfold.TTPair* objects). The train and test objects in each pair are also data frames, suitable for evaluation or writing out to a file.

Crossfold methods make minimal assumptions about their input data frames, so the frames can be ratings, purchases, or whatever. They do assume that each row represents a single data point for the purpose of splitting and sampling.

Experiment code should generally use these functions to prepare train-test files for training and evaluating algorithms. For example, the following will perform a user-based 5-fold cross-validation as was the default in the old LensKit:

```
import pandas as pd
import lenskit.crossfold as xf
ratings = pd.read_csv('ml-20m/ratings.csv')
ratings = ratings.rename(columns={'userId': 'user', 'movieId': 'item'})
for i, tp in enumerate(xf.partition_users(ratings, 5, xf.SampleN(5))):
    tp.train.to_csv('ml-20m.exp/train-%d.csv' % (i,))
    tp.train.to_parquet('ml-20m.exp/train-%d.parquet' % (i,))
    tp.test.to_csv('ml-20m.exp/test-%d.csv' % (i,))
    tp.test.to_parquet('ml-20m.exp/test-%d.parquet' % (i,))
```

2.3.1 Row-based splitting

The simplest preparation methods sample or partition the rows in the input frame. A 5-fold *partition_rows()* split will result in 5 splits, each of which extracts 20% of the rows for testing and leaves 80% for training.

`lenskit.crossfold.partition_rows(data, partitions)`

Partition a frame of ratings or other data into train-test partitions. This function does not care what kind of data is in *data*, so long as it is a Pandas DataFrame (or equivalent).

Parameters

- **data** (`pandas.DataFrame` or equivalent) – a data frame containing ratings or other data you wish to partition.
- **partitions** (*integer*) – the number of partitions to produce

Return type

iterator

Returns an iterator of train-test pairs

`lenskit.crossfold.sample_rows(data, partitions, size, disjoint=True)`

Sample train-test a frame of ratings into train-test partitions. This function does not care what kind of data is in *data*, so long as it is a Pandas DataFrame (or equivalent).

We can loop over a sequence of train-test pairs:

```
>>> ratings = util.load_ml_ratings()
>>> for train, test in sample_rows(ratings, 5, 1000):
...     print(len(test))
1000
1000
1000
1000
1000
```

Sometimes for testing, it is useful to just get a single pair:

```
>>> train, test = sample_rows(ratings, None, 1000)
>>> len(test)
1000
>>> len(test) + len(train) - len(ratings)
0
```

Parameters

- **data** (*pandas.DataFrame*) – Data frame containing ratings or other data to partition.
- **partitions** (*int* or *None*) – The number of partitions to produce. If *None*, produce a *_single_* train-test pair instead of an iterator or list.
- **size** (*int*) – The size of each sample.
- **disjoint** (*bool*) – If *True*, force samples to be disjoint.

Returns An iterator of train-test pairs.

Return type iterator

2.3.2 User-based splitting

It’s often desirable to use users, instead of raw rows, as the basis for splitting data. This allows you to control the experimental conditions on a user-by-user basis, e.g. by making sure each user is tested with the same number of ratings. These methods require that the input data frame have a *user* column with the user names or identifiers.

The algorithm used by each is as follows:

1. Sample or partition the set of user IDs into *n* sets of test users.
2. For each set of test users, select a set of that user’s rows to be test rows.
3. **Create a training set for each test set consisting of the non-selected rows from each** of that set’s test users, along with all rows from each non-test user.

```
lenskit.crossfold.partition_users(data, partitions: int, method:
                                lenskit.crossfold.PartitionMethod)
```

Partition a frame of ratings or other data into train-test partitions user-by-user. This function does not care what kind of data is in *data*, so long as it is a Pandas DataFrame (or equivalent) and has a *user* column.

Parameters

- **data** (*pandas.DataFrame* or equivalent) – a data frame containing ratings or other data you wish to partition.
- **partitions** (*integer*) – the number of partitions to produce
- **method** – The method for selecting test rows for each user.

Return type iterator

Returns an iterator of train-test pairs

```
lenskit.crossfold.sample_users(data, partitions: int, size: int, method:
                               lenskit.crossfold.PartitionMethod, disjoint=True)
```

Create train-test partitions by sampling users. This function does not care what kind of data is in *data*, so long as it is a Pandas DataFrame (or equivalent) and has a *user* column.

Parameters

- **data** (*pandas.DataFrame*) – Data frame containing ratings or other data you wish to partition.

- **partitions** (*int*) – The number of partitions.
- **size** (*int*) – The sample size.
- **method** (*PartitionMethod*) – The method for obtaining user test ratings.

Returns An iterator of train-test pairs (as *TTPair* objects).

Return type iterator

Selecting user test rows

These functions each take a *method* to decide how select each user's test rows. The method is a function that takes a data frame (containing just the user's rows) and returns the test rows. This function is expected to preserve the index of the input data frame (which happens by default with common means of implementing samples).

We provide several partition method factories:

`lenskit.crossfold.SampleN(n)`

Randomly select a fixed number of test rows per user/item.

Parameters **n** – The number of test items to select.

`lenskit.crossfold.SampleFrac(frac)`

Randomly select a fraction of test rows per user/item.

Parameters **frac** – the fraction of items to select for testing.

`lenskit.crossfold.LastN(n, col='timestamp')`

Select a fixed number of test rows per user/item, based on ordering by a column.

Parameters

- **n** – The number of test items to select.
- **col** – The column to sort by.

`lenskit.crossfold.LastFrac(frac, col='timestamp')`

Select a fraction of test rows per user/item.

Parameters

- **frac** – the fraction of items to select for testing.
- **col** – The column to sort by.

2.3.3 Utility Classes

class `lenskit.crossfold.PartitionMethod`

Partition methods select test rows for a user or item. Partition methods are callable; when called with a data frame, they return the test rows.

`__call__(udf)`

Subset a data frame.

Parameters **udf** – The input data frame of rows for a user or item.

Returns The data frame of test rows, a subset of *udf*.

class `lenskit.crossfold.TTPair`

Train-test pair (named tuple).

test

Test data for this pair.

train

Train data for this pair.

2.4 Batch-Running Recommenders

The functions in `lenskit.batch` enable you to generate many recommendations or predictions at the same time, useful for evaluations and experiments.

2.4.1 Recommendation

`lenskit.batch.recommend` (*algo*, *users*, *n*, *candidates=None*, *, *nprocs=None*, ***kwargs*)

Batch-recommend for multiple users. The provided algorithm should be a `algorithms.Recommender`.

Parameters

- **algo** – the algorithm
- **users** (*array-like*) – the users to recommend for
- **n** (*int*) – the number of recommendations to generate (None for unlimited)
- **candidates** – the users’ candidate sets. This can be a function, in which case it will be passed each user ID; it can also be a dictionary, in which case user IDs will be looked up in it. Pass None to use the recommender’s built-in candidate selector (usually recommended).
- **nprocs** (*int*) – The number of processes to use for parallel recommendations.

Returns A frame with at least the columns `user`, `rank`, and `item`; possibly also `score`, and any other columns returned by the recommender.

2.4.2 Rating Prediction

`lenskit.batch.predict` (*algo*, *pairs*, *, *nprocs=None*)

Generate predictions for user-item pairs. The provided algorithm should be a `algorithms.Predictor` or a function of two arguments: the user ID and a list of item IDs. It should return a dictionary or a `pandas.Series` mapping item IDs to predictions.

To use this function, provide a pre-fit algorithm:

```
>>> from lenskit.algorithms.basic import Bias
>>> from lenskit.metrics.predict import rmse
>>> ratings = util.load_ml_ratings()
>>> bias = Bias()
>>> bias.fit(ratings[: -1000])
<lenskit.algorithms.basic.Bias object at ...>
>>> preds = predict(bias, ratings[-1000:])
>>> preds.head()
   user  item  rating  timestamp  prediction
99004  664  8361    3.0  1393891425    3.288286
99005  664  8528    3.5  1393891047    3.559119
99006  664  8529    4.0  1393891173    3.573008
99007  664  8636    4.0  1393891175    3.846268
99008  664  8641    4.5  1393890852    3.710635
```

(continues on next page)

(continued from previous page)

```
>>> rmse(preds['prediction'], preds['rating'])
0.8326992222...
```

Parameters

- **algo** (`lenskit.algorithms.Predictor`) – A rating predictor function or algorithm.
- **pairs** (`pandas.DataFrame`) – A data frame of (user, item) pairs to predict for. If this frame also contains a *rating* column, it will be included in the result.
- **nprocs** (`int`) – The number of processes to use for parallel batch prediction.

Returns a frame with columns `user`, `item`, and `prediction` containing the prediction results. If `pairs` contains a *rating* column, this result will also contain a *rating* column.

Return type `pandas.DataFrame`

2.4.3 Scripting Evaluation

The `MultiEval` class is useful to build scripts that evaluate multiple algorithms or algorithm variants, simultaneously, across multiple data sets. It can extract parameters from algorithms and include them in the output, useful for hyperparameter search.

For example:

```
from lenskit.batch import MultiEval
from lenskit.crossfold import partition_users, SampleN
from lenskit.algorithms import basic, als
from lenskit.util import load_ml_ratings
from lenskit import topn
import pandas as pd
```

Generate the train-test pairs:

```
pairs = list(partition_users(load_ml_ratings(), 5, SampleN(5)))
```

Set up and run the `MultiEval` experiment:

```
eval = MultiEval('my-eval', recommend=20)
eval.add_datasets(pairs, name='ML-Small')
eval.add_algorithms(basic.Popular(), name='Pop')
eval.add_algorithms([als.BiasedMF(f) for f in [20, 30, 40, 50]],
                    attrs=['features'], name='ALS')
eval.run()
```

Now that the experiment is run, we can read its outputs.

First the run metadata:

```
runs = pd.read_csv('my-eval/runs.csv')
runs.set_index('RunId', inplace=True)
runs.head()
```

Then the recommendations:

```
recs = pd.read_parquet('my-eval/recommendations.parquet')
recs.head()
```

```
D:\Anaconda3\libs\site-packages\pyarrowpandas_compat.py:698: FutureWarning: .
  ↳ labels was deprecated in version 0.24.0. Use .codes instead.
  labels = getattr(columns, 'labels', None) or [
D:\Anaconda3\libs\site-packages\pyarrowpandas_compat.py:725: FutureWarning: the_
  ↳ 'labels' keyword is deprecated, use 'codes' instead
  return pd.MultiIndex(levels=new_levels, labels=labels, names=columns.names)
D:\Anaconda3\libs\site-packages\pyarrowpandas_compat.py:742: FutureWarning: .
  ↳ labels was deprecated in version 0.24.0. Use .codes instead.
  labels, = index.labels
```

In order to evaluate the recommendation list, we need to build a combined set of truth data. Since this is a disjoint partition of users over a single data set, we can just concatenate the individual test frames:

```
truth = pd.concat((p.test for p in pairs), ignore_index=True)
```

Now we can set up an analysis and compute the results.

```
rla = topn.RecListAnalysis()
rla.add_metric(topn.ndcg)
ndcg = rla.compute(recs, truth)
ndcg.head()
```

Next, we need to combine this with our run data, so that we know what algorithms and configurations we are evaluating:

```
ndcg = ndcg.join(runs[['AlgoClass', 'features']], on='RunId')
ndcg.head()
```

The Popular algorithm has NaN feature count, which `groupby` doesn't like; let's fill those in.

```
ndcg.loc[ndcg['AlgoClass'] == 'Popular', 'features'] = 0
```

And finally, we can compute the overall average performance for each algorithm configuration:

```
ndcg.groupby(['AlgoClass', 'features'])['ndcg'].mean()
```

```
AlgoClass  features
BiasedMF   20.0      0.015960
           30.0      0.022558
           40.0      0.025901
           50.0      0.028949
Popular    0.0      0.091814
Name: ndcg, dtype: float64
```

Multi-Eval Class Reference

```
class lenskit.batch.MultiEval(path, predict=True, recommend=100, candidates=<class
                             'lenskit.topn.UnratedCandidates'>, nprocs=None, com-
                             bine=True)
```

A runner for carrying out multiple evaluations, such as parameter sweeps.

Parameters

- **path** (*str* or `pathlib.Path`) – the working directory for this evaluation. It will be created if it does not exist.
- **predict** (*bool*) – whether to generate rating predictions.
- **recommend** (*int*) – the number of recommendations to generate per user. Any false-y value (`None`, `False`, `0`) will disable top-n. The literal value `True` will generate recommendation lists of unlimited size.
- **candidates** (*function*) – the default candidate set generator for recommendations. It should take the training data and return a candidate generator, itself a function mapping user IDs to candidate sets.
- **combine** (*bool*) – whether to combine output; if `False`, output will be left in separate files, if `True`, it will be in a single set of files (runs, recommendations, and predictions).

add_algorithms (*algos*, *parallel=False*, *attrs=[]*, ***kwargs*)

Add one or more algorithms to the run.

Parameters

- **algos** (*algorithm* or *list*) – the algorithm(s) to add.
- **parallel** (*bool*) – if `True`, allow this algorithm to be trained in parallel with others.
- **attrs** (*list of str*) – a list of attributes to extract from the algorithm objects and include in the run descriptions.
- **kwargs** – additional attributes to include in the run descriptions.

add_datasets (*data*, *name=None*, *candidates=None*, ***kwargs*)

Add one or more datasets to the run.

Parameters

- **data** – The input data set(s) to run. Can be one of the following:
 - A tuple of (train, test) data.
 - An iterable of (train, test) pairs, in which case the iterable is not consumed until it is needed.
 - A function yielding either of the above, to defer data load until it is needed.
 Data can be either data frames or paths; paths are loaded after detection using `util.read_df_detect()`.
- **kwargs** – additional attributes pertaining to these data sets.

collect_results ()

Collect the results from non-combined runs into combined output files.

persist_data ()

Persist the data for an experiment, replacing in-memory data sets with file names. Once this has been called, the sweep can be pickled.

run (*runs=None*, ***, *progress=None*)

Run the evaluation.

Parameters

- **runs** (*int* or *set-like*) – If provided, a specific set of runs to run. Useful for splitting an experiment into individual runs. This is a set of 1-based run IDs, not 0-based indexes.
- **progress** – A `tqdm.tqdm()`-compatible progress function.

`run_count()`

Get the number of runs in this evaluation.

2.5 Evaluating Recommender Output

LensKit's evaluation support is based on post-processing the output of recommenders and predictors. The `batch` utilities provide support for generating these outputs.

We generally recommend using `Jupyter` notebooks for evaluation.

2.5.1 Prediction Accuracy Metrics

The `lenskit.metrics.predict` module contains prediction accuracy metrics. These are intended to be used as a part of a Pandas split-apply-combine operation on a data frame that contains both predictions and ratings; for convenience, the `lenskit.batch.predict()` function will include ratings in the prediction frame when its input user-item pairs contains ratings. So you can perform the following to compute per-user RMSE over some predictions:

```
preds = predict(algo, pairs)
user_rmse = preds.groupby('user').apply(lambda df: rmse(df.prediction, df.rating))
```

Metric Functions

Prediction metric functions take two series, *predictions* and *truth*.

`lenskit.metrics.predict.rmse(predictions, truth, missing='error')`

Compute RMSE (root mean squared error).

Parameters

- **predictions** (`pandas.Series`) – the predictions
- **truth** (`pandas.Series`) – the ground truth ratings from data
- **missing** (`string`) – how to handle predictions without truth. Can be one of 'error' or 'ignore'.

Returns the root mean squared approximation error

Return type double

`lenskit.metrics.predict.mae(predictions, truth, missing='error')`

Compute MAE (mean absolute error).

Parameters

- **predictions** (`pandas.Series`) – the predictions
- **truth** (`pandas.Series`) – the ground truth ratings from data
- **missing** (`string`) – how to handle predictions without truth. Can be one of 'error' or 'ignore'.

Returns the mean absolute approximation error

Return type double

Working with Missing Data

LensKit rating predictors do not report predictions when their core model is unable to predict. For example, a nearest-neighbor recommender will not score an item if it cannot find any suitable neighbors. Following the Pandas convention, these items are given a score of NaN (when Pandas implements better missing data handling, it will use that, so use `pandas.Series.isna()/pandas.Series.notna()`, not the `isnan` versions).

However, this causes problems when computing predictive accuracy: recommenders are not being tested on the same set of items. If a recommender only scores the easy items, for example, it could do much better than a recommender that is willing to attempt more difficult items.

A good solution to this is to use a *fallback predictor* so that every item has a prediction. In LensKit, `lenskit.algorithms.basic.Fallback` implements this functionality; it wraps a sequence of recommenders, and for each item, uses the first one that generates a score.

You set it up like this:

```
cf = ItemItem(20)
base = Bias(damping=5)
algo = Fallback(cf, base)
```

2.5.2 Top-N Evaluation

LensKit's support for top-*N* evaluation is in two parts, because there are some subtle complexities that make it more difficult to get the right data in the right place for computing metrics correctly.

Top-N Analysis

The `lenskit.topn` module contains the utilities for carrying out top-*N* analysis, in conjunction with `lenskit.batch.recommend()` and its wrapper in `lenskit.batch.MultiEval`.

The entry point to this is `RecListAnalysis`. This class encapsulates an analysis with one or more metrics, and can apply it to data frames of recommendations. An analysis requires two data frames: the recommendation frame contains the recommendations themselves, and the truth frame contains the ground truth data for the users. The analysis is flexible with regards to the columns that identify individual recommendation lists; usually these will consist of a user ID, data set identifier, and algorithm identifier(s), but the analysis is configurable and its defaults make minimal assumptions. The recommendation frame does need an `item` column with the recommended item IDs, and it should be in order within a single recommendation list.

The truth frame should contain (a subset of) the columns identifying recommendation lists, along with `item` and, if available, `rating` (if no rating is provided, the metrics that need a rating value will assume a rating of 1 for every item present). It can contain other items that custom metrics may find useful as well.

For example, a recommendation frame may contain:

- DataSet
- Partition
- Algorithm
- user
- item
- rank
- score

And the truth frame:

- DataSet
- user
- item
- rating

The analysis will use this truth as the relevant item data for measuring the accuracy of the recommendation lists. Recommendations will be matched to test ratings by data set, user, and item, using *RecListAnalysis* defaults.

class `lenskit.topn.RecListAnalysis` (*group_cols=None*)
Compute one or more top-N metrics over recommendation lists.

This method groups the recommendations by the specified columns, and computes the metric over each group. The default set of grouping columns is all columns *except* the following:

- item
- rank
- score
- rating

The truth frame, `truth`, is expected to match over (a subset of) the grouping columns, and contain at least an `item` column. If it also contains a `rating` column, that is used as the users' rating for metrics that require it; otherwise, a rating value of 1 is assumed.

Parameters `group_cols` (*list*) – The columns to group by, or `None` to use the default.

add_metric (*metric*, *, *name=None*, ***kwargs*)
Add a metric to the analysis.

A metric is a function of two arguments: the a single group of the recommendation frame, and the corresponding truth frame. The truth frame will be indexed by item ID. Many metrics are defined in `lenskit.metrics.topn`; they are re-exported from `lenskit.topn` for convenience.

Parameters

- **metric** – The metric to compute.
- **name** – The name to assign the metric. If not provided, the function name is used.
- ****kwargs** – Additional arguments to pass to the metric.

compute (*recs*, *truth*, *, *progress=<function RecListAnalysis.<lambda>>*)
Run the analysis. Neither data frame should be meaningfully indexed.

Parameters

- **recs** (*pandas.DataFrame*) – A data frame of recommendations.
- **truth** (*pandas.DataFrame*) – A data frame of ground truth (test) data.

Returns The results of the analysis.

Return type `pandas.DataFrame`

Metrics

The `lenskit.metrics.topn` module contains metrics for evaluating top-*N* recommendation lists.

Classification Metrics

These metrics treat the recommendation list as a classification of relevant items.

`lenskit.metrics.topn.precision` (*recs*, *truth*)
 Compute recommendation precision.

`lenskit.metrics.topn.recall` (*recs*, *truth*)
 Compute recommendation recall.

Ranked List Metrics

These metrics treat the recommendation list as a ranked list of items that may or may not be relevant.

`lenskit.metrics.topn.recip_rank` (*recs*, *truth*)
 Compute the reciprocal rank of the first relevant item in a list of recommendations.
 If no elements are relevant, the reciprocal rank is 0.

Utility Metrics

The NDCG function estimates a utility score for a ranked list of recommendations.

`lenskit.metrics.topn.ndcg` (*recs*, *truth*, *discount*=<*ufunc* 'log2'>)
 Compute the normalized discounted cumulative gain.

Discounted cumulative gain is computed as:

$$\text{DCG}(L, u) = \sum_{i=1}^{|L|} \frac{r_{ui}}{d(i)}$$

This is then normalized as follows:

$$\text{nDCG}(L, u) = \frac{\text{DCG}(L, u)}{\text{DCG}(L_{\text{ideal}}, u)}$$

Parameters

- **recs** – The recommendation list.
- **truth** – The user's test data.
- **discount** (*ufunc*) – The rank discount function. Each item's score will be divided the discount of its rank, if the discount is greater than 1.

We also expose the internal DCG computation directly.

`lenskit.metrics.topn._dcg` (*scores*, *discount*=<*ufunc* 'log2'>)
 Compute the Discounted Cumulative Gain of a series of recommended items with rating scores. These should be relevance scores; they can be 0, 1 for binary relevance data.

This is not a true top-N metric, but is a utility function for other metrics.

Parameters

- **scores** (*array-like*) – The utility scores of a list of recommendations, in recommendation order.
- **discount** (*ufunc*) – the rank discount function. Each item's score will be divided the discount of its rank, if the discount is greater than 1.

Returns the DCG of the scored items.

Return type double

2.5.3 Loading Outputs

We typically store the output of recommendation runs in LensKit experiments in CSV or Parquet files. The `lenskit.batch.MultiEval` class arranges to run a set of algorithms over a set of data sets, and store the results in a collection of Parquet files in a specified output directory.

There are several files:

runs.parquet The `_runs_`, algorithm-dataset combinations. This file contains the names & any associated properties of each algorithm and data set run, such as a feature count.

recommendations.parquet The recommendations, with columns `RunId`, `user`, `rank`, `item`, and `rating`.

predictions.parquet The rating predictions, if the test data includes ratings.

For example, if you want to examine nDCG by neighborhood count for a set of runs on a single data set, you can do:

```
import pandas as pd
from lenskit.metrics import topn as lm

runs = pd.read_parquet('eval-dir/runs.parquet')
recs = pd.read_parquet('eval-dir/recs.parquet')
meta = runs.loc[:, ['RunId', 'max_neighbors']]

# compute each user's nDCG
user_ndcg = recs.groupby(['RunId', 'user']).rating.apply(lm.ndcg)
user_ndcg = user_ndcg.reset_index(name='nDCG')
# combine with metadata for feature count
user_ndcg = pd.merge(user_ndcg, meta)
# group and aggregate
nbr_ndcg = user_ndcg.groupby('max_neighbors').nDCG.mean()
nbr_ndcg.plot()
```

2.6 Errors and Diagnostics

2.6.1 Logging

LensKit algorithms and evaluation routines report diagnostic data using the standard Python `logging` framework. Loggers are named after the corresponding Python module, and all live under the `lenskit` namespace.

Algorithms usually report erroneous or anomalous conditions using Python exceptions and warnings. **Evaluation code**, such as that in `lenskit.batch`, typically reports such conditions using the logger, as the common use case is to be running them in a script.

2.6.2 Warnings

In addition to Python standard warning types such as `warnings.DeprecationWarning`, LensKit uses the following warning classes to report anomalous problems in use of LensKit.

class `lenskit.DataWarning`

Warning raised for detectable problems with input data.

2.7 Algorithms

LKPY provides general algorithmic concepts, along with implementations of several algorithms. These algorithm interfaces are based on the SciKit design patterns [SKAPI], adapted for Pandas-based data structures.

2.7.1 Basic and Utility Algorithms

The `lenskit.algorithms.basic` module contains baseline and utility algorithms for nonpersonalized recommendation and testing.

Personalized Mean Rating Prediction

class `lenskit.algorithms.basic.Bias` (*items=True, users=True, damping=0.0*)

Bases: `lenskit.algorithms.Predictor`

A user-item bias rating prediction algorithm. This implements the following predictor algorithm:

$$s(u, i) = \mu + b_i + b_u$$

where μ is the global mean rating, b_i is item bias, and b_u is the user bias. With the provided damping values β_u and β_i , they are computed as follows:

$$\mu = \frac{\sum_{r_{ui} \in R} r_{ui}}{|R|} \quad b_i = \frac{\sum_{r_{ui} \in R_i} (r_{ui} - \mu)}{|R_i| + \beta_i} \quad b_u = \frac{\sum_{r_{ui} \in R_u} (r_{ui} - \mu - b_i)}{|R_u| + \beta_u}$$

The damping values can be interpreted as the number of default (mean) ratings to assume *a priori* for each user or item, damping low-information users and items towards a mean instead of permitting them to take on extreme values based on few ratings.

Parameters

- **items** – whether to compute item biases
- **users** – whether to compute user biases
- **damping** (*number or tuple*) – Bayesian damping to apply to computed biases. Either a number, to damp both user and item biases the same amount, or a (user,item) tuple providing separate damping values.

mean_

The global mean rating.

Type `double`

item_offsets_

The item offsets (b_i values)

Type `pandas.Series`

user_offsets_

The user offsets (b_u values)

Type `pandas.Series`

fit (*data*)

Train the bias model on some rating data.

Parameters **data** (`DataFrame`) – a data frame of ratings. Must have at least *user*, *item*, and *rating* columns.

Returns the fit bias object.

Return type *Bias*

predict_for_user (*user, items, ratings=None*)

Compute predictions for a user and items. Unknown users and items are assumed to have zero bias.

Parameters

- **user** – the user ID
- **items** (*array-like*) – the items to predict
- **ratings** (*pandas.Series*) – the user’s ratings (indexed by item id); if provided, will be used to recompute the user’s bias at prediction time.

Returns scores for the items, indexed by item id.

Return type *pandas.Series*

Fallback Predictor

The `Fallback` rating predictor is a simple hybrid that takes a list of composite algorithms, and uses the first one to return a result to predict the rating for each item.

A common case is to fill in with *Bias* when a primary predictor cannot score an item.

class `lenskit.algorithms.basic.Fallback` (*algorithms, *others*)

Bases: *lenskit.algorithms.Predictor*

The `Fallback` algorithm predicts with its first component, uses the second to fill in missing values, and so forth.

fit (*ratings, *args, **kwargs*)

Train a model using the specified ratings (or similar) data.

Parameters

- **ratings** (*pandas.DataFrame*) – The ratings data.
- **args** – Additional training data the algorithm may require.
- **kwargs** – Additional training data the algorithm may require.

Returns The algorithm object.

predict_for_user (*user, items, ratings=None*)

Compute predictions for a user and items.

Parameters

- **user** – the user ID
- **items** (*array-like*) – the items to predict
- **ratings** (*pandas.Series*) – the user’s ratings (indexed by item id); if provided, they may be used to override or augment the model’s notion of a user’s preferences.

Returns scores for the items, indexed by item id.

Return type *pandas.Series*

Memorized Predictor

The Memorized recommender is primarily useful for test cases. It memorizes a set of rating predictions and returns them.

```
class lenskit.algorithms.basic.Memorized(scores)
```

Bases: *lenskit.algorithms.Predictor*

The memorized algorithm memorizes scores provided at construction time.

```
fit (*args, **kwargs)
```

Train a model using the specified ratings (or similar) data.

Parameters

- **ratings** (*pandas.DataFrame*) – The ratings data.
- **args** – Additional training data the algorithm may require.
- **kwargs** – Additional training data the algorithm may require.

Returns The algorithm object.

```
predict_for_user (user, items, ratings=None)
```

Compute predictions for a user and items.

Parameters

- **user** – the user ID
- **items** (*array-like*) – the items to predict
- **ratings** (*pandas.Series*) – the user’s ratings (indexed by item id); if provided, they may be used to override or augment the model’s notion of a user’s preferences.

Returns scores for the items, indexed by item id.

Return type *pandas.Series*

2.7.2 k-NN Collaborative Filtering

LKPY provides user- and item-based classical k-NN collaborative Filtering implementations. These lightly-configurable implementations are intended to capture the behavior of the Java-based LensKit implementations to provide a good upgrade path and enable basic experiments out of the box.

Item-based k-NN

```
class lenskit.algorithms.item_knn.ItemItem(nnbrs,      min_nbrs=1,      min_sim=1e-
                                          06,      save_nbrs=None,      center=True,
                                          aggregate='weighted-average')
```

Bases: *lenskit.algorithms.Predictor*

Item-item nearest-neighbor collaborative filtering with ratings. This item-item implementation is not terribly configurable; it hard-codes design decisions found to work well in the previous Java-based LensKit code.

```
item_index_
```

the index of item IDs.

Type *pandas.Index*

```
item_means_
```

the mean rating for each known item.

Type `numpy.ndarray`

item_counts_
the number of saved neighbors for each item.

Type `numpy.ndarray`

sim_matrix_
the similarity matrix.

Type `matrix.CSR`

user_index_
the index of known user IDs for the rating matrix.

Type `pandas.Index`

rating_matrix_
the user-item rating matrix for looking up users' ratings.

Type `matrix.CSR`

fit (*ratings*)
Train a model.

The model-training process depends on `save_nbrs` and `min_sim`, but *not* on other algorithm parameters.

Parameters **ratings** (`pandas.DataFrame`) – (user,item,rating) data for computing item similarities.

predict_for_user (*user, items, ratings=None*)
Compute predictions for a user and items.

Parameters

- **user** – the user ID
- **items** (*array-like*) – the items to predict
- **ratings** (`pandas.Series`) – the user's ratings (indexed by item id); if provided, they may be used to override or augment the model's notion of a user's preferences.

Returns scores for the items, indexed by item id.

Return type `pandas.Series`

User-based k-NN

class `lenskit.algorithms.user_knn.UserUser` (*nbrs, min_nbrs=1, min_sim=0, center=True, aggregate='weighted-average'*)

Bases: `lenskit.algorithms.Predictor`

User-user nearest-neighbor collaborative filtering with ratings. This user-user implementation is not terribly configurable; it hard-codes design decisions found to work well in the previous Java-based LensKit code.

user_index_
User index.

Type `pandas.Index`

item_index_
Item index.

Type `pandas.Index`

user_means_

User mean ratings.

Type `numpy.ndarray`**rating_matrix_**

Normalized user-item rating matrix.

Type `matrix.CSR`**transpose_matrix_**

Transposed un-normalized rating matrix.

Type `matrix.CSR`**fit** (*ratings*)

“Train” a user-user CF model. This memorizes the rating data in a format that is usable for future computations.

Parameters **ratings** (`pandas.DataFrame`) – (user, item, rating) data for collaborative filtering.**Returns** a memorized model for efficient user-based CF computation.**Return type** `UUModel`**predict_for_user** (*user, items, ratings=None*)

Compute predictions for a user and items.

Parameters

- **user** – the user ID
- **items** (*array-like*) – the items to predict
- **ratings** (`pandas.Series`) – the user’s ratings (indexed by item id); if provided, will be used to recompute the user’s bias at prediction time.

Returns scores for the items, indexed by item id.**Return type** `pandas.Series`

2.7.3 Classic Matrix Factorization

LKPYPY provides classical matrix factorization implementations.

- *Common Support*
- *Alternating Least Squares*
- *FunkSVD*

Common Support

The `mf_common` module contains common support code for matrix factorization algorithms.

class `lenskit.algorithms.mf_common.MFPredictor`

Common predictor for matrix factorization.

user_index_

Users in the model (length=:math:m).

Type `pandas.Index`

item_index_

Items in the model (length=:math:n).

Type `pandas.Index`

user_features_

The $m \times k$ user-feature matrix.

Type `numpy.ndarray`

item_features_

The $n \times k$ item-feature matrix.

Type `numpy.ndarray`

lookup_items (*items*)

Look up the indices for a set of items.

Parameters **items** (*array-like*) – the item IDs to look up.

Returns the item indices. Unknown items will have negative indices.

Return type `numpy.ndarray`

lookup_user (*user*)

Look up the index for a user.

Parameters **user** – the user ID to look up

Returns the user index.

Return type `int`

n_features

The number of features.

n_items

The number of items.

n_users

The number of users.

score (*user, items*)

Score a set of items for a user. User and item parameters must be indices into the matrices.

Parameters

- **user** (*int*) – the user index
- **items** (*array-like of int*) – the item indices
- **raw** (*bool*) – if True, do return raw scores without biases added back.

Returns the scores for the items.

Return type `numpy.ndarray`

class `lenskit.algorithms.mf_common.BiasMFPredictor`

Common model for biased matrix factorization.

user_index_

Users in the model (length=:math:m).

Type `pandas.Index`

item_index_

Items in the model (length=:math:n).

Type `pandas.Index`

global_bias_

The global bias term.

Type `double`

user_bias_

The user bias terms.

Type `numpy.ndarray`

item_bias_

The item bias terms.

Type `numpy.ndarray`

user_features_

The $m \times k$ user-feature matrix.

Type `numpy.ndarray`

item_features_

The $n \times k$ item-feature matrix.

Type `numpy.ndarray`

score (*user*, *items*, *raw=False*)

Score a set of items for a user. User and item parameters must be indices into the matrices.

Parameters

- **user** (*int*) – the user index
- **items** (*array-like of int*) – the item indices
- **raw** (*bool*) – if True, do return raw scores without biases added back.

Returns the scores for the items.

Return type `numpy.ndarray`

Alternating Least Squares

LensKit provides alternating least squares implementations of matrix factorization suitable for explicit feedback data. These implementations are parallelized with Numba, and perform best with the MKL from Conda.

class `lenskit.algorithms.als.BiasedMF` (*features*, *, *iterations=20*, *reg=0.1*, *damping=5*, *bias=True*)

Bases: `lenskit.algorithms.mf_common.BiasMFPredictor`

Biased matrix factorization trained with alternating least squares [ZWSP2008]. This is a prediction-oriented algorithm suitable for explicit feedback data.

Parameters

- **features** (*int*) – the number of features to train
- **iterations** (*int*) – the number of iterations to train
- **reg** (*double*) – the regularization factor
- **damping** (*double*) – damping factor for the underlying mean

fit (*ratings*)

Run ALS to train a model.

Parameters **ratings** – the ratings data frame.

Returns The algorithm (for chaining).

predict_for_user (*user, items, ratings=None*)

Compute predictions for a user and items.

Parameters

- **user** – the user ID
- **items** (*array-like*) – the items to predict
- **ratings** (*pandas.Series*) – the user’s ratings (indexed by item id); if provided, they may be used to override or augment the model’s notion of a user’s preferences.

Returns scores for the items, indexed by item id.

Return type `pandas.Series`

class `lenskit.algorithms.als.ImplicitMF` (*features, *, iterations=20, reg=0.1, weight=40*)

Bases: `lenskit.algorithms.mf_common.MFPredictor`

Implicit matrix factorization trained with alternating least squares [HKV2008]. This algorithm outputs ‘predictions’, but they are not on a meaningful scale. If its input data contains `rating` values, these will be used as the ‘confidence’ values; otherwise, confidence will be 1 for every rated item.

Parameters

- **features** (*int*) – the number of features to train
- **iterations** (*int*) – the number of iterations to train
- **reg** (*double*) – the regularization factor
- **weight** (*double*) – the scaling weight for positive samples (α in [HKV2008]).

fit (*ratings*)

Train a model using the specified ratings (or similar) data.

Parameters

- **ratings** (*pandas.DataFrame*) – The ratings data.
- **args** – Additional training data the algorithm may require.
- **kwargs** – Additional training data the algorithm may require.

Returns The algorithm object.

predict_for_user (*user, items, ratings=None*)

Compute predictions for a user and items.

Parameters

- **user** – the user ID
- **items** (*array-like*) – the items to predict
- **ratings** (*pandas.Series*) – the user’s ratings (indexed by item id); if provided, they may be used to override or augment the model’s notion of a user’s preferences.

Returns scores for the items, indexed by item id.

Return type `pandas.Series`

FunkSVD

FunkSVD is an SVD-like matrix factorization that uses stochastic gradient descent, configured much like coordinate descent, to train the user-feature and item-feature matrices.

```
class lenskit.algorithms.funksvd.FunkSVD (features, iterations=100, *, lrate=0.001,
                                         reg=0.015, damping=5, range=None,
                                         bias=True)
```

Bases: `lenskit.algorithms.mf_common.BiasMFPredictor`

Algorithm class implementing FunkSVD matrix factorization.

Parameters

- **features** (*int*) – the number of features to train
- **iterations** (*int*) – the number of iterations to train each feature
- **lrate** (*double*) – the learning rate
- **reg** (*double*) – the regularization factor
- **damping** (*double*) – damping factor for the underlying mean
- **bias** (*Predictor*) – the underlying bias model to fit. If `True`, then a `basic.Bias` model is fit with damping.
- **range** (*tuple*) – the (min, max) rating values to clamp ratings, or `None` to leave predictions unclamped.

fit (*ratings*)

Train a FunkSVD model.

Parameters **ratings** – the ratings data frame.

predict_for_user (*user, items, ratings=None*)

Compute predictions for a user and items.

Parameters

- **user** – the user ID
- **items** (*array-like*) – the items to predict
- **ratings** (*pandas.Series*) – the user’s ratings (indexed by item id); if provided, they may be used to override or augment the model’s notion of a user’s preferences.

Returns scores for the items, indexed by item id.

Return type `pandas.Series`

2.7.4 Hierarchical Poisson Factorization

This module provides a LensKit bridge to the `hpfrec` library implementing hierarchical Poisson factorization [GHB2013].

```
class lenskit.algorithms.hpfc.HPF (features, **kwargs)
    Hierarchical Poisson factorization, provided by hpfrec.
```

Parameters

- **features** (*int*) – the number of features
- ****kwargs** – arguments passed to `hpfrec.HPF`.

fit (*ratings*)

Train a model using the specified ratings (or similar) data.

Parameters

- **ratings** (*pandas.DataFrame*) – The ratings data.
- **args** – Additional training data the algorithm may require.
- **kwargs** – Additional training data the algorithm may require.

Returns The algorithm object.

predict_for_user (*user, items, ratings=None*)

Compute predictions for a user and items.

Parameters

- **user** – the user ID
- **items** (*array-like*) – the items to predict
- **ratings** (*pandas.Series*) – the user’s ratings (indexed by item id); if provided, they may be used to override or augment the model’s notion of a user’s preferences.

Returns scores for the items, indexed by item id.

Return type *pandas.Series*

2.7.5 Implicit

This module provides a LensKit bridge to Ben Frederickson’s [implicit](#) library implementing some implicit-feedback recommender algorithms, with an emphasis on matrix factorization.

class `lenskit.algorithms.implicit.ALS` (**args, **kwargs*)
LensKit interface to `implicit.als`.

class `lenskit.algorithms.implicit.BPR` (**args, **kwargs*)
LensKit interface to `implicit.bpr`.

2.8 Utility Functions

2.8.1 Matrix Utilities

We have some matrix-related utilities, since matrices are used so heavily in recommendation algorithms.

Building Ratings Matrices

`lenskit.matrix.sparse_ratings` (*ratings, scipy=False*)
Convert a rating table to a sparse matrix of ratings.

Parameters

- **ratings** (*pandas.DataFrame*) – a data table of (user, item, rating) triples.
- **scipy** – if True, return a SciPy matrix instead of *CSR*.

Returns a named tuple containing the sparse matrix, user index, and item index.

Return type *RatingMatrix*

class `lenskit.matrix.RatingMatrix`

A rating matrix with associated indices.

matrix

The rating matrix, with users on rows and items on columns.

Type `CSR` or `scipy.sparse.csr_matrix`

users

mapping from user IDs to row numbers.

Type `pandas.Index`

items

mapping from item IDs to column numbers.

Type `pandas.Index`

Compressed Sparse Row Matrices

We use CSR-format sparse matrices in quite a few places. Since SciPy’s sparse matrices are not directly usable from Numba, we have implemented a Numba-compiled CSR representation that can be used from accelerated algorithm implementations.

class `lenskit.matrix.CSR` (*nrows=None, ncols=None, nnz=None, ptrs=None, inds=None, vals=None, N=None*)

Simple compressed sparse row matrix. This is like `scipy.sparse.csr_matrix`, with a couple of useful differences:

- It is backed by a Numba jitclass, so it can be directly used from Numba-optimized functions.
- The value array is optional, for cases in which only the matrix structure is required.
- The value array, if present, is always double-precision.

You generally don’t want to create this class yourself with the constructor. Instead, use one of its class methods.

If you need to pass an instance off to a Numba-compiled function, use *N*:

```
_some_numba_fun(csr.N)
```

We use the indirection between this and the Numba jitclass so that the main CSR implementation can be pickled, and so that we can have class and instance methods that are not compatible with jitclass but which are useful from interpreted code.

N

the Numba jitclass backing (has the same attributes and most methods).

Type `_CSR`

nrows

the number of rows.

Type `int`

ncols

the number of columns.

Type `int`

nnz

the number of entries.

Type `int`

rowptrs

the row pointers.

Type `numpy.ndarray`

colinds

the column indices.

Type `numpy.ndarray`

values

the values

Type `numpy.ndarray`

classmethod from_coo (*rows, cols, vals, shape=None*)

Create a CSR matrix from data in COO format.

Parameters

- **rows** (*array-like*) – the row indices.
- **cols** (*array-like*) – the column indices.
- **vals** (*array-like*) – the data values; can be None.
- **shape** (*tuple*) – the array shape, or None to infer from row & column indices.

classmethod from_scipy (*mat, copy=True*)

Convert a scipy sparse matrix to an internal CSR.

Parameters

- **mat** (*scipy.sparse.spmatrix*) – a SciPy sparse matrix.
- **copy** (*bool*) – if False, reuse the SciPy storage if possible.

Returns a CSR matrix.

Return type `CSR`

row (*row*)

Return a row of this matrix as a dense ndarray.

Parameters **row** (*int*) – the row index.

Returns the row, with 0s in the place of missing values.

Return type `numpy.ndarray`

row_cs (*row*)

Get the column indices for the stored values of a row.

row_extent (*row*)

Get the extent of a row in the underlying column index and value arrays.

Parameters **row** (*int*) – the row index.

Returns (*s, e*), where the row occupies positions [*s, e*) in the CSR data.

Return type `tuple`

row_nnz ()

Get a vector of the number of nonzero entries in each row.

Note: This method is not available from Numba.

Returns the number of nonzero entries in each row.

Return type `numpy.ndarray`

row_vs (*row*)

Get the stored values of a row.

rowinds () → `numpy.ndarray`

Get the row indices from this array. Combined with `colinds` and `values`, this can form a COO-format sparse matrix.

Note: This method is not available from Numba.

sort_values ()

Sort CSR rows in nonincreasing order by value.

Note: This method is not available from Numba.

to_scipy ()

Convert a CSR matrix to a SciPy `scipy.sparse.csr_matrix`.

Parameters **self** (*CSR*) – A CSR matrix.

Returns A SciPy sparse matrix with the same data.

Return type `scipy.sparse.csr_matrix`

ttranspose (*values=True*)

Transpose a CSR matrix.

Note: This method is not available from Numba.

Parameters **values** (*bool*) – whether to include the values in the transpose.

Returns the transpose of this matrix (or, equivalently, this matrix in CSC format).

Return type *CSR*

class `lenskit.matrix._CSR` (*nrows, ncols, nnz, ptrs, inds, vals*)

Internal implementation class for *CSR*. If you work with CSRs from Numba, you will use this.

2.8.2 Math utilities

Solvers

`lenskit.math.solve.dposv` (*A, b, lower=False*)

Interface to the BLAS `dposv` function. A Numba-accessible version without error checking is exposed as `_dposv()`.

`lenskit.math.solve.solve_tri` (*A, b, transpose=False, lower=True*)

Solve the system $Ax = b$, where A is triangular. This is equivalent to `scipy.linalg.solve_triangular()`, but does *not* check for non-singularity. It is a thin wrapper around the BLAS `dtrsv` function.

Parameters

- **A** (*ndarray*) – the matrix.
- **b** (*ndarray*) – the target vector.
- **transpose** (*bool*) – whether to solve $Ax = b$ or $A^T x = b$.
- **lower** (*bool*) – whether A is lower- or upper-triangular.

Numba-accessible internals

`lenskit.math.solve._dposv()`

`lenskit.math.solve._dtrsv()`

2.8.3 Miscellaneous

Miscellaneous utility functions.

`lenskit.util.clone` (*algo*)

Clone an algorithm, but not its fitted data. This is like `scikit.base.clone()`, but may not work on arbitrary SciKit estimators. LensKit algorithms are compatible with SciKit clone, however, so feel free to use that if you need more general capabilities.

This function is somewhat derived from the SciKit one.

```
>>> from lenskit.algorithms.basic import Bias
>>> orig = Bias()
>>> copy = clone(orig)
>>> copy is orig
False
>>> copy.damping == orig.damping
True
```

`lenskit.util.fspath` (*path*)

Backport of `os.fspath()` function for Python 3.5.

`lenskit.util.load_ml_ratings` (*path='ml-latest-small'*)

Load the ratings from a modern MovieLens data set (ML-20M or one of the ‘latest’ data sets).

```
>>> load_ml_ratings().head()
  user item rating  timestamp
0    1    31    2.5  1260759144
1    1  1029    3.0  1260759179
2    1  1061    3.0  1260759182
3    1  1129    2.0  1260759185
4    1  1172    4.0  1260759205
```

Parameters *path* – The path where the MovieLens data is unpacked.

Returns The rating data, with user and item columns named properly for LensKit.

Return type `pandas.DataFrame`

`lenskit.util.read_df_detect` (*path*)

Read a Pandas data frame, auto-detecting the file format based on filename suffix. The following file types are supported:

CSV File has suffix `.csv`, read with `pandas.read_csv()`.

Parquet File has suffix `.parquet`, `.parq`, or `.pq`, read with `pandas.read_parquet()`.

`lenskit.util.write_parquet(path, frame, append=False)`

Write a Parquet file.

Parameters

- **path** (`pathlib.Path`) – The path of the Parquet file to write.
- **frame** (`pandas.DataFrame`) – The data to write.
- **append** (`bool`) – Whether to append to the file or overwrite it.

2.9 Release Notes

2.9.1 0.6.1

See the [GitHub milestone](#) for full change list.

- Fix inconsistency in both code and docs for recommend list sizes for top- N evaluation.
- Fix user-user to correctly use `sum` aggregate.
- Improve performance and documentation

2.9.2 0.6.0

See the [GitHub milestone](#) for a summary of what's happening!

- The `save` and `load` methods on algorithms have been removed. Just pickle fitted models to save their data. This is what SciKit does, we see no need to deviate.
- The APIs and model structures for top- N recommendation is reworked to enable algorithms to produce recommendations more automatically. The `Recommender` interfaces now take a `CandidateSelector` to determine default candidates, so client code does not need to compute candidates on their own. One effect of this is that the `batch.recommend` function no longer requires a candidate selector, and there can be problems if you call `Recommender.adapt` before fitting a model.
- Top- N evaluation has been completely revamped to make it easier to correctly implement and run evaluation metrics. Batch recommend no longer attaches ratings to recommendations. See [Top- \$N\$ evaluation](#) for details.
- Batch recommend & predict functions now take `nprocs` as a keyword-only argument.
- Several bug fixes and testing improvements.

Internal Changes

These changes should not affect you if you are only consuming LensKit's algorithm and evaluation capabilities.

- Rewrite the `CSR` class to be more ergonomic from Python, at the expense of making the NumPy jitclass indirect. It is available in the `.N` attribute. Big improvement: it is now picklable.

2.9.3 0.5.0

LensKit 0.5.0 modifies the algorithm APIs to follow the SciKit design patterns instead of our previous custom patterns. Highlights of this change:

- Algorithms are trained in-place — we no longer have distinct model objects.
- Model data is stored as attributes on the algorithm object that end in `_`.
- Instead of writing `model = algo.train_model(ratings)`, call `algo.fit(ratings)`.

We also have some new capabilities:

- Ben Frederickson's Implicit library

2.9.4 0.3.0

A number of improvements, including replacing Cython/OpenMP with Numba and adding ALS.

2.9.5 0.2.0

A lot of fixes to get ready for RecSys.

2.9.6 0.1.0

Hello, world!

INDICES AND TABLES

- genindex
- modindex
- search

BIBLIOGRAPHY

- [SKAPI] Lars Buitinck, Gilles Louppe, Mathieu Blondel, Fabian Pedregosa, Andreas Mueller, Olivier Grisel, Vlad Niculae, Peter Prettenhofer, Alexandre Gramfort, Jaques Grobler, Robert Layton, Jake Vanderplas, Arnaud Joly, Brian Holt, and Gaël Varoquaux. 2013. API design for machine learning software: experiences from the scikit-learn project. arXiv:1309.0238 [cs.LG].
- [ZWSP2008] Yunhong Zhou, Dennis Wilkinson, Robert Schreiber, and Rong Pan. 2008. Large-Scale Parallel Collaborative Filtering for the Netflix Prize. In *Algorithmic Aspects in Information and Management*, LNCS 5034, 337–348. DOI 10.1007/978-3-540-68880-8_32.
- [HKV2008] Y. Hu, Y. Koren, and C. Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining*, 263–272. DOI 10.1109/ICDM.2008.22
- [GHB2013] Prem Gopalan, Jake M. Hofman, and David M. Blei. 2013. Scalable Recommendation with Poisson Factorization. arXiv:1311.1704 [cs, stat] (November 2013). Retrieved February 9, 2017 from <http://arxiv.org/abs/1311.1704>.

PYTHON MODULE INDEX

|

lenskit.algorithms, 8
lenskit.algorithms.als, 29
lenskit.algorithms.basic, 23
lenskit.algorithms.funksvd, 31
lenskit.algorithms.hpf, 31
lenskit.algorithms.implicit, 32
lenskit.algorithms.item_knn, 25
lenskit.algorithms.mf_common, 27
lenskit.algorithms.user_knn, 26
lenskit.batch, 14
lenskit.crossfold, 11
lenskit.math.solve, 35
lenskit.matrix, 32
lenskit.metrics.predict, 18
lenskit.metrics.topn, 20
lenskit.topn, 19
lenskit.util, 36

Symbols

`_CSR` (class in `lenskit.matrix`), 35
`__call__` () (lenskit.crossfold.PartitionMethod method), 13
`_dcg` () (in module `lenskit.metrics.topn`), 21
`_dposv` () (in module `lenskit.math.solve`), 36
`_dtrsv` () (in module `lenskit.math.solve`), 36

A

`adapt` () (lenskit.algorithms.Recommender class method), 9
`add_algorithms` () (lenskit.batch.MultiEval method), 17
`add_datasets` () (lenskit.batch.MultiEval method), 17
`add_metric` () (lenskit.topn.RecListAnalysis method), 20
`Algorithm` (class in `lenskit.algorithms`), 8
`ALS` (class in `lenskit.algorithms.implicit`), 32

B

`Bias` (class in `lenskit.algorithms.basic`), 23
`BiasedMF` (class in `lenskit.algorithms.als`), 29
`BiasMFPredictor` (class in `lenskit.algorithms.mf_common`), 28
`BPR` (class in `lenskit.algorithms.implicit`), 32

C

`candidates` () (lenskit.algorithms.CandidateSelector method), 10
`CandidateSelector` (class in `lenskit.algorithms`), 10
`clone` () (in module `lenskit.util`), 36
`colinds` (lenskit.matrix.CSR attribute), 34
`collect_results` () (lenskit.batch.MultiEval method), 17
`compute` () (lenskit.topn.RecListAnalysis method), 20
`CSR` (class in `lenskit.matrix`), 33

D

`DataWarning` (class in `lenskit`), 22
`dposv` () (in module `lenskit.math.solve`), 35

F

`Fallback` (class in `lenskit.algorithms.basic`), 24
`fit` () (lenskit.algorithms.Algorithm method), 8
`fit` () (lenskit.algorithms.als.BiasedMF method), 29
`fit` () (lenskit.algorithms.als.ImplicitMF method), 30
`fit` () (lenskit.algorithms.basic.Bias method), 23
`fit` () (lenskit.algorithms.basic.Fallback method), 24
`fit` () (lenskit.algorithms.basic.Memorized method), 25
`fit` () (lenskit.algorithms.funksvd.FunkSVD method), 31
`fit` () (lenskit.algorithms.hpf.HPF method), 31
`fit` () (lenskit.algorithms.item_knn.ItemItem method), 26
`fit` () (lenskit.algorithms.user_knn.UserUser method), 27
`from_coo` () (lenskit.matrix.CSR class method), 34
`from_scipy` () (lenskit.matrix.CSR class method), 34
`fspath` () (in module `lenskit.util`), 36
`FunkSVD` (class in `lenskit.algorithms.funksvd`), 31

G

`get_params` () (lenskit.algorithms.Algorithm method), 8
in `global_bias_` (lenskit.algorithms.mf_common.BiasMFPredictor attribute), 29

H

`HPF` (class in `lenskit.algorithms.hpf`), 31

I

`ImplicitMF` (class in `lenskit.algorithms.als`), 30
`item_bias_` (lenskit.algorithms.mf_common.BiasMFPredictor attribute), 29
`item_counts_` (lenskit.algorithms.item_knn.ItemItem attribute), 26
`item_features_` (lenskit.algorithms.mf_common.BiasMFPredictor attribute), 29
`item_features_` (lenskit.algorithms.mf_common.MFPredictor attribute), 28
`item_index_` (lenskit.algorithms.item_knn.ItemItem attribute), 25

item_index_ (*lenskit.algorithms.mf_common.BiasMFPredictor* attribute), 28
 item_index_ (*lenskit.algorithms.mf_common.MFPredictor* attribute), 28
 item_index_ (*lenskit.algorithms.user_knn.UserUser* attribute), 26
 item_means_ (*lenskit.algorithms.item_knn.ItemItem* attribute), 25
 item_offsets_ (*lenskit.algorithms.basic.Bias* attribute), 23
 ItemItem (class in *lenskit.algorithms.item_knn*), 25
 items (*lenskit.matrix.RatingMatrix* attribute), 33

L

LastFrac() (in module *lenskit.crossfold*), 13
 LastN() (in module *lenskit.crossfold*), 13
 lenskit.algorithms (module), 8
 lenskit.algorithms.als (module), 29
 lenskit.algorithms.basic (module), 23
 lenskit.algorithms.funksvd (module), 31
 lenskit.algorithms.hpf (module), 31
 lenskit.algorithms.implicit (module), 32
 lenskit.algorithms.item_knn (module), 25
 lenskit.algorithms.mf_common (module), 27
 lenskit.algorithms.user_knn (module), 26
 lenskit.batch (module), 14
 lenskit.crossfold (module), 11
 lenskit.math.solve (module), 35
 lenskit.matrix (module), 32
 lenskit.metrics.predict (module), 18
 lenskit.metrics.topn (module), 20
 lenskit.topn (module), 19
 lenskit.util (module), 36
 load_ml_ratings() (in module *lenskit.util*), 36
 lookup_items() (*lenskit.algorithms.mf_common.MFPredictor* method), 28
 lookup_user() (*lenskit.algorithms.mf_common.MFPredictor* method), 28

M

mae() (in module *lenskit.metrics.predict*), 18
 matrix (*lenskit.matrix.RatingMatrix* attribute), 33
 mean_ (*lenskit.algorithms.basic.Bias* attribute), 23
 Memorized (class in *lenskit.algorithms.basic*), 25
 MFPredictor (class in *lenskit.algorithms.mf_common*), 27
 MultiEval (class in *lenskit.batch*), 16

N

N (*lenskit.matrix.CSR* attribute), 33
 n_features (*lenskit.algorithms.mf_common.MFPredictor* attribute), 28
 n_items (*lenskit.algorithms.mf_common.MFPredictor* attribute), 28

ndcg() (in module *lenskit.metrics.topn*), 21
 nnz (*lenskit.matrix.CSR* attribute), 33
 nrows (*lenskit.matrix.CSR* attribute), 33

P

partition_rows() (in module *lenskit.crossfold*), 11
 partition_users() (in module *lenskit.crossfold*), 12
 PartitionMethod (class in *lenskit.crossfold*), 13
 persist_data() (*lenskit.batch.MultiEval* method), 17
 precision() (in module *lenskit.metrics.topn*), 21
 predict() (in module *lenskit.batch*), 14
 predict() (*lenskit.algorithms.Predictor* method), 10
 predict_for_user() (*lenskit.algorithms.als.BiasedMF* method), 30
 predict_for_user() (*lenskit.algorithms.als.ImplicitMF* method), 30
 predict_for_user() (*lenskit.algorithms.basic.Bias* method), 24
 predict_for_user() (*lenskit.algorithms.basic.Fallback* method), 24
 predict_for_user() (*lenskit.algorithms.basic.Memorized* method), 25
 predict_for_user() (*lenskit.algorithms.funksvd.FunkSVD* method), 31
 predict_for_user() (*lenskit.algorithms.hpf.HPF* method), 32
 predict_for_user() (*lenskit.algorithms.item_knn.ItemItem* method), 26
 predict_for_user() (*lenskit.algorithms.Predictor* method), 10
 predict_for_user() (*lenskit.algorithms.user_knn.UserUser* method), 27
 Predictor (class in *lenskit.algorithms*), 10

R

rated_items() (*lenskit.algorithms.CandidateSelector* static method), 10
 rating_matrix_ (*lenskit.algorithms.item_knn.ItemItem* attribute), 26
 rating_matrix_ (*lenskit.algorithms.user_knn.UserUser* attribute), 27
 RatingMatrix (class in *lenskit.matrix*), 32
 read_df_detect() (in module *lenskit.util*), 36

recall() (in module *lenskit.metrics.topn*), 21
 recip_rank() (in module *lenskit.metrics.topn*), 21
 RecListAnalysis (class in *lenskit.topn*), 20
 recommend() (in module *lenskit.batch*), 14
 recommend() (*lenskit.algorithms.Recommender* method), 9
 Recommender (class in *lenskit.algorithms*), 9
 rmse() (in module *lenskit.metrics.predict*), 18
 row() (*lenskit.matrix.CSR* method), 34
 row_cs() (*lenskit.matrix.CSR* method), 34
 row_extent() (*lenskit.matrix.CSR* method), 34
 row_nnz() (*lenskit.matrix.CSR* method), 34
 row_vs() (*lenskit.matrix.CSR* method), 35
 rowinds() (*lenskit.matrix.CSR* method), 35
 rowptrs (*lenskit.matrix.CSR* attribute), 33
 run() (*lenskit.batch.MultiEval* method), 17
 run_count() (*lenskit.batch.MultiEval* method), 17

S

sample_rows() (in module *lenskit.crossfold*), 11
 sample_users() (in module *lenskit.crossfold*), 12
 SampleFrac() (in module *lenskit.crossfold*), 13
 SampleN() (in module *lenskit.crossfold*), 13
 score() (*lenskit.algorithms.mf_common.BiasMFPredictor* method), 29
 score() (*lenskit.algorithms.mf_common.MFPredictor* method), 28
 sim_matrix_ (*lenskit.algorithms.item_knn.ItemItem* attribute), 26
 solve_tri() (in module *lenskit.math.solve*), 35
 sort_values() (*lenskit.matrix.CSR* method), 35
 sparse_ratings() (in module *lenskit.matrix*), 32

T

test (*lenskit.crossfold.TTPair* attribute), 13
 to_scipy() (*lenskit.matrix.CSR* method), 35
 train (*lenskit.crossfold.TTPair* attribute), 14
 transpose() (*lenskit.matrix.CSR* method), 35
 transpose_matrix_ (*lenskit.algorithms.user_knn.UserUser* attribute), 27
 TTPair (class in *lenskit.crossfold*), 13

U

user_bias_ (*lenskit.algorithms.mf_common.BiasMFPredictor* attribute), 29
 user_features_ (*lenskit.algorithms.mf_common.BiasMFPredictor* attribute), 29
 user_features_ (*lenskit.algorithms.mf_common.MFPredictor* attribute), 28
 user_index_ (*lenskit.algorithms.item_knn.ItemItem* attribute), 26
 user_index_ (*lenskit.algorithms.mf_common.BiasMFPredictor* attribute), 28

user_index_ (*lenskit.algorithms.mf_common.MFPredictor* attribute), 27
 user_index_ (*lenskit.algorithms.user_knn.UserUser* attribute), 26
 user_means_ (*lenskit.algorithms.user_knn.UserUser* attribute), 26
 user_offsets_ (*lenskit.algorithms.basic.Bias* attribute), 23
 users (*lenskit.matrix.RatingMatrix* attribute), 33
 UserUser (class in *lenskit.algorithms.user_knn*), 26

V

values (*lenskit.matrix.CSR* attribute), 34

W

write_parquet() (in module *lenskit.util*), 37