

LensKit for Python

Next-Generation Software for Recommender Systems Experiments

Michael D. Ekstrand

People & Information Research Team

Dept. of Computer Science, Boise State University

Boise, Idaho, USA

michaelekstrand@boisestate.edu

ABSTRACT

LensKit is an open-source toolkit for building, researching, and learning about recommender systems. First released in 2010 as a Java framework, it has supported diverse published research, small-scale production deployments, and education in both MOOC and traditional classroom settings. In this paper, I present the next generation of the LensKit project, re-envisioning the original tool's objectives as flexible Python package for supporting recommender systems research and development. LensKit for Python (LKPY) enables researchers and students to build robust, flexible, and reproducible experiments that make use of the large and growing PyData and Scientific Python ecosystem, including scikit-learn, and TensorFlow. To that end, it provides classical collaborative filtering implementations, recommender system evaluation metrics, data preparation routines, and tools for efficiently batch running recommendation algorithms, all usable in any combination with each other or with other Python software.

This paper describes the design goals, use cases, and capabilities of LKPY, contextualized in a reflection on the successes and failures of the original LensKit for Java software.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; *Evaluation of retrieval results*; • **General and reference** → *Experimentation*.

KEYWORDS

recommender systems, evaluation, experiments, support software

ACM Reference Format:

Michael D. Ekstrand. 2020. LensKit for Python: Next-Generation Software for Recommender Systems Experiments. In *Proceedings of The 29th ACM International Conference on Information and Knowledge Management (CIKM '20)*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3340531.3412778>

1 INTRODUCTION

The goal of the **LensKit project** is to develop and deploy tools and documentation that enable data processing, evaluation, and

The LensKit code, documentation, and community discussion forums are available from <https://lenskit.org>.

CIKM '20, October 19–23, 2020, Virtual Event, Ireland

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of The 29th ACM International Conference on Information and Knowledge Management (CIKM '20)*, <https://doi.org/10.1145/3340531.3412778>.

algorithms for recommender systems, with a particular focus on empowering research, education, and early-stage development.

For years, the original Java version of the LensKit toolkit [8], the founding software of the LensKit project, has supported a wide array of research, educational, and deployment efforts. Based on the experience my collaborators and I have using LensKit for research and teaching, in addition to feedback reports we received directly and indirectly from researchers and developers in industry and academia, we determined that the design and technology choices of the Java implementation were not a good match for the current and future needs of the recommender systems research community.

When we examined the software landscape to determine what existing tools might support the kinds of evaluation experiments people have carried out with LensKit and might run in coming years, it became clear to us that there is a need for high-quality, well-tested support code for recommender systems experiments in a programming environment better suited for research and teaching. Python data science libraries, including Pandas, scikit-learn, TensorFlow, and PyTorch, are widely used in recommender systems research and data science training, but there are not commonly-used implementations for experimental support code or evaluation metrics. General-purpose machine learning tools, such as scikit-learn, often miss subtle recommender-specific evaluation details, such as splitting data by user instead of rating. Existing recommender systems packages for Python (such as Surprise and PyRecLab) are often self-contained, like our Java LensKit code was, in that they drive the evaluation process and report a metric, instead of providing building blocks that the user can employ to build an experimental protocol that meets their precise needs.

To provide such tools in a maximally flexible fashion, I — with support from my students and collaborators — have developed LensKit for Python, a successor to the original Java code. This new-generation toolkit brings the LensKit vision of *reproducible research* supported by *well-tested code* to a more *widely-used* and *easier-to-learn* computational environment. LensKit code and examples are MIT-licensed, to minimize legal barriers to adoption.

In this paper, I begin by reflecting on the successes and failures of the LensKit project and its Java software to date. I then present the goals, design, and use cases of the LensKit for Python (LKPY) software, along with a comparison to related tools. I invite the community to provide feedback on how this software does or does not meet their needs and to participate in its ongoing development.

2 REFLECTIONS ON LENSKIT FOR JAVA

The Java version of LensKit saw active use from 2011 through 2018, supporting over 50 publications¹, multiple production systems (including MovieLens [16]), and the Recommender Systems MOOC on Coursera [21]. Reflection on my own personal experience and reports my collaborators and I received from members of the community yielded several insights into things that did and did not work for supporting researchers and educators.

2.1 What We Got Right

There are several things that I think we did well in the original LensKit software; some of these carry forward into the Python version, while others seem to have been good ideas in LensKit’s original context but are not as important today.

2.1.1 Testing. In the LensKit development process, we have a strong focus on code testing [7]. This has served us well, and helped ensure the reliability of the LensKit code. It is by no means perfect, and there have been bugs that slipped through, but the effort we spent on testing was a wise investment.

2.1.2 The Java Platform. When we began work on LensKit, there were three platforms we seriously considered: Java, Python, and C++. At that time, the Python data science ecosystem was not what it is today; NumPy and SciPy existed, but Pandas was not yet well-known. Pure Python does not have the performance needed for efficient recommender experiments. Java enabled us to achieve strong computational performance in a widely-taught programming language with well-established standard practices, making it significantly easier for new users (particularly students) to adapt and contribute to it than I expect we would have seen with a corresponding C++ code base. The earliest prototype code was in OCaml, but adoptability demanded a more widely-known language.

2.1.3 Modular Algorithms. LensKit was built around the idea of modular algorithms where individual components can be replaced and reconfigured. In the item-item collaborative filter, for example, users can change the similarity function, the neighborhood weighting function, input rating vector normalizations, item neighborhood normalizations, and the strategy for mapping input data to rating vectors. This configurability has been useful for exploring a range of configuration options for algorithms, at the expense of larger hyperparameter search spaces and increased software complexity.

2.2 What Didn’t Work

Other aspects of LensKit for Java’s design and development do not seem to have met our needs or those of the community as well.

2.2.1 Opinionated Evaluation. While LensKit’s algorithms were highly configurable, its offline evaluation tools are much more opinionated. Users can specify a few types of data splitting strategies, and recommendation candidate strategies, and it has a range of evaluation metrics, but the overall evaluation process and methods for aggregating metric results are fixed. Metrics are also limited in their interface.

As my collaborators and I expanded our research into the recommender evaluation process itself, we repeatedly ran in to limits

¹A full list is at <https://lenskit.org/research>

of this evaluation strategy. We often had to write new evaluation code in a fairly heavy framework or make the evaluator save intermediate files that we would reprocess in R or Python in order to carry out our research. Too often the answer I would have to give to questions on the mailing list or StackOverflow is “we’re sorry, LensKit can’t do that yet”.

One of our goals was to make it difficult to do an evaluation incorrectly: we wanted the defaults to embody best practices for offline evaluation. However, best practices have been sufficiently unknown and fast-moving that we now believe this approach has held our research back more than it has helped the field. It is particularly apparent that a different approach is necessary to support next-generation offline evaluation strategies, such as counterfactual evaluation [4], and to carry out the evaluation research needed to advance understanding of effective, robust, and both internally- and ecologically-valid offline evaluations.

2.2.2 Indirect Configuration. LensKit for Java is built on the dependency injection principle, using a custom dependency injection container [11] to instantiate and connect recommender components. This method gave us quite a few useful capabilities, such as automatically detecting components that could be shared between multiple experiment runs in a parameter tuning experiment. The cost, however, was that it was difficult to configure an algorithm or to understand an algorithm’s configuration. The configuration visualization capabilities afforded by the DI framework were not enough to compensate for the complexity and subtlety of configuring an algorithm through dependency bindings. It was also how to document how to use the system.

We believe this largely stems from the role of inversion of control in working with the LensKit for Java code — users never write code that assembles an algorithm, they simply ask LensKit to instantiate one and LensKit calls their custom components. It is therefore difficult to develop an intuition about how LensKit works and how to realize a desired configuration.

2.2.3 Implicit Features. Beyond indirect configuration, LensKit has a lot of implicit behavior in its algorithms and evaluator. This has at least two downsides: first, it is less clear from reading a configuration precisely what LensKit will do, making it more difficult to review code and experiment scripts; second, if documentation slipped behind the code, understanding the behavior of LensKit experiment scripts required reading the LensKit source code itself.

2.2.4 Living in an Island. LensKit has its own data structures and data access paradigms. Part of this is due to lack of standardized, high-quality scientific data tooling that is not connected to a larger framework such as Spark (while Spark does seem to expose data structures as a separate library, they are not well-documented).

This made it difficult, however, to make LensKit interoperate with other software and data sets. While LensKit for Java is flexible in the data it accepts, users must write data adapters, and common data layers such as Hibernate had unacceptable performance penalties for research experiments. Integrating with other tools such as Spark is difficult. It’s unclear what, given the Java commitment, we could have done differently here, but it is definitely a liability for the future of recommender systems research.

3 DESIGN GOALS

Based on my experience with the Java software and sense of the current landscape of recommender systems research software, I identified several design goals to make LKPY as useful and usable as possible. These goals include:

Build on Standard Tools There are now standard tools, such as Pandas and the surrounding PyData ecosystem [24], that are widely adopted for data science and machine learning research. Using them maximizes interoperability with other software packages and enables code reuse across different types of research. Jupyter notebooks are also a valuable means of promoting reproducibility, and my team uses experiment workflows where the final analysis consists of loading recommender outputs into a Jupyter notebook and computing desired metrics over them. Small experiments can even be driven entirely from Jupyter.

Leverage Existing Software Modern recommender systems are usually machine learning systems; a recommender research toolkit should not try to reinvent that wheel. Scikit-Learn provides many machine learning algorithms suitable for recommender systems research, and PyTorch and TensorFlow provide general-purpose optimization for modular differentiable machine learning models. Recommender research tooling should work seamlessly with algorithms implemented in these kinds of toolkits.

Expose the Data Pipeline Many recommender systems packages, including the Java version of LensKit, hide or at least take ownership of the data pipeline: they control data splitting, algorithm training, recommendation, and evaluation. Outputs of each stage can be examined, but not manipulated, and the pipeline itself cannot be easily changed. Putting the user in control of the pipeline, and providing functions to implement standard versions of each of its stages, has two benefits: the actual pipeline used is clearly documented in the experiment code, and the pipeline can be modified as research needs demand.

Explicit is Better than Implicit Python has a long-standing philosophy of explicitly denoting desired operations instead of depending on implicitly-defined behavior. Applying this to research code would, I expect, make it easier to review experiment designs and improve the reliability and rigor of reproducible research. I also posit that exposing the data pipeline will also make the LKPY code itself easier to read and understand.

Simple Interfaces Interfaces to individual software components should be as simple as possible, so that it is easy to document, test, and reimplement them. They should also be data-agnostic, when feasible, to avoid unnecessary limits the types of data or algorithms that can be used with LKPY.

Ease of Development For long-term viability, the development process, not just the resulting software, needs to be approachable for prospective users and contributors. I particularly want to ensure that students are able to contribute to LensKit so it can benefit directly from their research. Therefore, the development tools should be as standard and easy-to-use as possible.

Notably absent from this list is the Java version of LensKit's (in)famous algorithm configurability. That configurability was useful for exploring the space of algorithm configurations, but its particular design is more suitable to heuristic techniques such as

k-NN; machine learning approaches seem better served by a different design. Connecting with existing flexible optimization software such as TensorFlow and PyTorch will provide a great deal of configurability for new algorithms. I think it is more important for recommender-specific software to focus on flexibility in the experiment design to facilitate education and new types of research.

4 LENSKIT USE CASES

LensKit for Python directly supports several use cases, and has a modular and loosely-coupled design to support many more use cases in the future. In this section I describe use cases explicitly supported by LensKit for Python, along with available public work that uses them, as well as some use cases that we are working on supporting in the very near future.

4.1 Recommender System Evaluation

LensKit enables offline evaluation of the effectiveness of new and improved algorithms through several features:

- Support for bulk-running recommendations and predictions for comparison with test data.
- Evaluation code to compute common prediction accuracy and top- N accuracy metrics.
- High-quality, well-tested implementations of classic collaborative filters for use as baselines.
- Easy-to-implement interfaces to enable head-to-head comparisons between new approaches and baseline or prior state-of-the-art techniques.
- Easily usable in and with other tools, such as Jupyter for reporting experiment results and scikit-optimize for hyperparameter search.

These features are loosely coupled, so they can be recombined, adapted, and individually replaced to meet the needs of a wide range of research designs. Unlike the Java version of LensKit and many other recommendation tools, the structure of an evaluation and the relationship of data to metrics is not prescribed by LKPY.

LensKit for Python has seen published use in this capacity to study recommendation techniques for children's books [26], in addition to our demonstration experiment that compares the accuracy of different recommender algorithms on a several public data sets².

4.2 Offline Recommender System Research

The same facilities that make LKPY useful for evaluating new recommendation techniques are also valuable for other experiments that study the behavior of recommender algorithms on available data, the characteristics of experimental protocols, and many other research objectives around recommender systems and their data.

My own research team has used LKPY to study errors in evaluation protocols [36], and to rebuild the experiments from our work on author gender biases [12] for an expanded version currently under review. Narayan [25] used LKPY to study the effect of rating obscure items, and we expect more such projects to use the software in the coming years.

²<https://github.com/lenskit/lk-demo-experiment>

4.3 Education

LensKit is conceived as a toolkit that is meant to be useful in educational settings, as an environment for teaching recommender systems and personalization at either the undergraduate or graduate levels. I have used it successfully to support a graduate-level class (Boise State CS 538 — *Recommender Systems and Online Personalization*), and have received reports of others successfully using it as well.

LensKit is also suitable for online demonstrations. PBS Digital Studios' *Crash Course: Artificial Intelligence* program used it to demonstrate recommender systems in their video and accompanying notebook on Google Collaboratory [29].

4.4 Online Studies (In Progress)

We want to be able to use LensKit for Python in online user studies as well as offline experiments.

Right now, online experiments are possible by embedding the Python code in a web application built with a framework like Flask or Django. However, this does not scale well and requires tedious and error-prone code. To address this limitation, my student and I are currently developing a recommendation server that exposes LKPY's recommender APIs as a web service. When completed (expected by late 2020), using any LKPY algorithm — or any other algorithm that implements the required Python interfaces — will be as easy as making a REST call.

4.5 Small-Scale Production (In Progress)

The same recommender server that will enable online studies (see Section 4.4) will also make it easy to use LensKit for Python in small-to medium-scale production services and prototypes. Large-scale services will likely require other facilities, but startups and research groups will be able to use LKPY to prototype recommendation ideas and move to other software as their needs grow.

5 LENSKIT SOFTWARE DESCRIPTION

To support these design goals and use cases, LKPY provides:

- Data preparation for recommender experiments
- Common APIs for recommendation algorithms
- Batch-processing functions for carrying out experiments
- Evaluation metrics
- A suite of collaborative filtering algorithms

5.1 Example Code

Figure 1 shows an example of using LKPY to compute the normalized discounted cumulative gain of a matrix factorization algorithm with 5-fold cross-validation on the MovieLens 100K data set [16], consisting of 100K movie ratings from just under 1000 users. It is somewhat verbose, but every step of the evaluation process is clear in the code. This allows the experiment structure to be reviewed during the research process and makes the code self-documenting when it is published with the final paper. Further, since the data at each stage is in standard Pandas and NumPy data structures, it can be inspected, analyzed, and transformed at any point to support arbitrary research objectives.

```
import pandas as pd
from lenskit import batch, topn, datasets
from lenskit.metrics.predict import rmse
from lenskit.crossfold import partition_users, SampleN
from lenskit.algorithms import als, Recommender

# Load MovieLens data in expected format
ml = datasets.ML100K('data/ml-100k')

# run the evaluation
all_recs = []
all_preds = []
all_test = []
samp = xf.SampleN(5)
for train, test in partition_users(ml.ratings, 5, samp):
    all_test.append(test)
    # configure the algorithm
    algo = als.BiasedMF(50)
    # make sure it can recommend, not just predict
    algo = Recommender.adapt(algo)
    # train the algorithm
    algo.fit(train)

    # predict ratings
    preds = batch.predict(algo, test)
    all_preds.append(preds)

    # produce recommendations
    users = test['user'].unique()
    recs = batch.recommend(algo, users, 20)
    all_recs.append(recs)

# integrate test data; since splits are disjoint users
# for a single data set, this is sufficient
recs = pd.concat(all_recs, ignore_index=True)
preds = pd.concat(all_preds, ignore_index=True)
test = pd.concat(all_test, ignore_index=True)

# compute RMSE
print('RMSE:', rmse(preds['prediction'],
                    preds['rating']))

# compute nDCG for each user
r1a = topn.RecListAnalysis()
r1a.add_metric(topn.ndcg)
metrics = r1a.compute(recs, test)
# and compute overall system performance
print('NDCG:', metrics['ndcg'].mean())
```

Figure 1: Example of a simple evaluation.

5.2 Core Interfaces and Data Expectations

LKPY's algorithm-related code uses the Algorithm interface, based on the patterns used by Scikit-Learn [5], where a fit method estimates model parameters (trains the model) in-place. This makes

the API familiar to users with experience using other packages in the SciKit ecosystem.

Common recommendation operations are defined by additional interfaces that subclass `Algorithm`:

Predictor implements user-item preference prediction; implementations using explicit-feedback data typically return predicted ratings, while other algorithms return purchase probabilities or other scores useful for ranking results. The `predict_for_user` method takes a user ID and a list of items and returns predictions for the specified items. It can optionally take a Pandas series of user ratings; some algorithms are able to use these ratings to provide recommendations that are responsive to user ratings that were not available when the model was trained.

Recommender implements top- N recommendation. It defines the `recommend` method, which takes a user ID and optionally a desired list length and/or list of candidate item IDs and returns a ranked recommendation list, along with scores if applicable. It also can accept a series of ratings.

CandidateSelector is usually not used alone, but selects candidate items for recommendation when no candidate set has been provided to `recommend`.

The `fit` method takes a `ratings` argument containing user-item preference data. This argument is a Pandas `DataFrame` with “user” and “item” columns; explicit-feedback rating data stores the rating value in a “rating” column. LensKit code — including both algorithm implementations and infrastructural code that passes data to them — makes no assumptions about the (non-)existence of other columns, and also allows additional named arguments to `fit` that the built-in algorithm implementations ignore, to allow arbitrary additional data to be passed to algorithm implementations. Therefore, while LensKit’s current algorithm implementations are focused on collaborative filtering, content-based, context-aware, and other types of recommendation techniques can be readily tested. For example, a review-based recommender can require `ratings` to also have a “review” column, or a content-based recommender can take an additional `item_data` parameter with item side information; LensKit code in the call chain between the user’s experimental code and their algorithm’s `fit` method will ignore and pass on data it doesn’t recognize.

A large number of recommendation algorithms, even many learning-to-rank techniques, compute user-personalized scores for individual items and then rank items by score for the final selection (e.g. even Bayesian Personalized Ranking [32] learns a function $s(i|u)$ by optimizing it to correctly rank items). Implementers of such algorithms will typically implement the `Predictor` interface, and use it with LensKit’s top- N wrapper to produce item recommendation lists. Algorithms that directly produce ranked lists, or perturb existing ranked lists (such as diversification strategies [40]), will implement the `Recommender` interface. This separation enables code reuse (top- N logic is the same no matter how the underlying scores are computed), and allows researchers to easily and independently experiment with new ideas for either item scoring or aggregating scored items into recommendation lists.

The LKPY interfaces are designed to support both batch operation — LensKit’s current primary usage mode in an offline experiment — and per-user recommendation for online settings. As I

discuss in §4.4, the online use case is not yet as easy to deploy as we would like, but LKPY’s algorithms and new algorithms others develop with its interfaces will be immediately usable once we complete the required infrastructural development (§6).

5.3 Experiment and Evaluation Support

As we illustrate in Figure 1, a typical offline recommender system experiment consists of three major steps:

- (1) Prepare training and test partitions of user ratings, interaction traces, or other recommender input data.
- (2) Train algorithms and produce recommendations or predictions for test users.
- (3) Compute metrics over the algorithm output

LKPY provides code for each of these steps.

The `crossfold` module provides multiple mechanisms for preparing multiple train-test partitions of ratings or similarly-structured data. Just partitioning the ratings, as a standard machine learning cross-validation package would do, does not result in an accurate simulation of the actual application goals of a typical recommender system [15]. LKPY’s crossfolding support focuses on test *users*, either partitioning the users (as in cross-validation) or producing disjoint sets of test users (for when evaluating on all users is cost-prohibitive). For each test user, it will select a set of test ratings (either randomly or using timestamps) and place the user’s other ratings in the training set, along with the ratings from users not in the test partition, so the model can learn the test user’s preferences. We also provide functions to partition or sample ratings, and to perform item-based sampling as required by certain protocols [3].

The `batch` module provides functions to mass-produce predictions or recommendations, with support for parallelizing over test users, to facilitate the second step of the experiment. It provides both batch prediction, where the input is an algorithm and a data frame of (user,item) pairs (if the data frame also contains ratings, those are included in the prediction output to ease prediction accuracy calculation), and batch recommendation using an algorithm and a list of user IDs. The output of both is a Pandas data frame that can be analyzed directly or serialized for later use.

Finally, the `metrics` and `topn` modules provide evaluation support. Prediction accuracy is relatively straightforward to support: after merging the prediction data with the rating data, MAE, RMSE, and other metrics can be computed with Pandas aggregate operations. The LKPY’s `rmse` and `mae` functions are usable in this setting, and can be used to compute either per-system or per-user accuracy.

Top- N recommendation list accuracy is more difficult to implement correctly in the general case. Precision, hit rate, average precision, and reciprocal rank can be computed by taking a left join between the recommendation list and the test ratings and then computing the metric; other metrics, such as normalized discounted cumulative gain and recall, require access not only to the recommended items and their relevance information, but also to the test items that were *not* recommended. Top- N evaluation metrics are implemented as functions of two parameters — a recommendation list and the user’s test ratings or other relevance data — and return a metric value, so users can implement new metrics easily.

The `ReclistAnalysis` class groups recommendation results by user, algorithm, data set, and other grouping variables, matches

them with test ratings, and provides both recommendation lists and test ratings to the metric to facilitate correct computation.

5.4 Algorithm Implementations

Non-personalized algorithms are useful as naive baselines, to establish bounds on the performance of more sophisticated algorithms, and as fallbacks or building blocks for other techniques. Non-personalized algorithms in LKPY include:

Popular recommends the most frequently-consumed items a user has not interacted with.

Bias implements the user-item bias rating predictor [9].

Random recommends randomly-selected items.

The k -NN implementations provide classic neighborhood-based collaborative filtering, and work with both explicit ratings and implicit-feedback consumption data:

UserUser implements the user-user algorithm [18], using cosine similarity over mean-centered rating vectors. The similarity function is fixed to facilitate performance and because we have found it to work remarkably well over a wide range of data sets [8].

ItemItem implements the item-item algorithm [6, 33], also with cosine similarity over mean-centered vectors.

Both k -NN algorithms have configurable neighborhood sizes and rating aggregation strategies.

LensKit also provides matrix factorization implementations for both explicit and implicit feedback:

BiasedMF provides biased matrix factorization for explicit feedback learned with alternating least squares [39], optimized with coordinate descent [31].

ImplicitMF provides ALS matrix factorization for implicit feedback data [19] using the conjugate gradient method [35].

FunkSVD implements Simon Funk’s gradient descent biased matrix factorization [13].

svd provides biased matrix factorization using the truncated SVD algorithm from SciKit-Learn [30].

tf implements various matrix factorization models using TensorFlow [1]. This includes biased MF for rating prediction and Bayesian Personalized Ranking (BPR, [32]) for learning-to-rank. Biased MF comes in two versions: one that fits matrices to the residuals of the Bias model, and another that jointly learns biases and latent feature vectors.

The TensorFlow and Scikit-Learn implementations serve both as baselines for evaluating new algorithms and as examples for developing new proposed recommendation techniques with these libraries within the LensKit framework.

Finally, LKPY provides implementations of its algorithm APIs using Implicit³ (Implicit ALS and BPR with GPU acceleration) and HPPREC⁴ (providing Hierarchical Poisson Factorization [14]).

The LensKit project’s primary goal is to make it easy to carry out experiments with algorithms from anywhere, so new algorithms are not our first priority; we do, however, welcome both new implementations and bridges to other packages.

³<https://github.com/benfired/implicit>

⁴<https://github.com/david-cortes/hpprec>

5.5 Dependencies and Environment

LensKit for Python leverages Pandas [24], numpy [27], and scipy [28], along with several other Python modules. We use Numba [23] for native-code acceleration in custom algorithm implementations, as it enables efficient, parallelized NumPy-based code without complex build toolchains. We plan to make more use of TensorFlow for the optimization process of future algorithm implementations. The batch recommendation and predictions routines use Python’s multi-process concurrency support, combined with the shared-memory facilities in the binpickle package, to parallelize over users when generating recommender output for an offline experiment.

LKPY supports recent versions of Python 3 on Windows, Linux, and macOS. Project policy is to support the oldest version of Python 3 available in the most most recent stable release of major Linux distributions (Ubuntu LTS, RHEL/CentOS, and Debian). We optimize for Anaconda-based Python installations (including direct use of certain Intel MKL capabilities, with pure-Python SciPy fallbacks when MKL is not available), but also support vanilla Python on Linux and Windows⁵. We provide binaries via Anaconda Cloud⁶ and source distributions on the Python Package Index⁷.

5.6 Use Across Scales

LensKit supports both research and educational use at multiple scales. Simple experiments and demos can be done entirely in a single Jupyter or Collaboratory notebook, as shown by PBS *Crash Course* and the Quick Start guide in the project documentation.

As a project scales up, the recommendation code can move into scripts that produce results for subsequent analysis. Data pre-processing, model training, output generation, output aggregation, and final performance evaluation can all be split into separate stages with output written to intermediate files. This allows, for example, research teams to use a compute cluster to train models and generate recommendations, and download the resulting files to a local workstation for computing evaluation metrics and producing the final analysis report.

LKPY’s algorithm implementations are focused on single-node performance, and can scale with data that can be processed on a single node (which can be remarkably large when sufficient memory is available). For even larger data sets, users can develop their recommendation techniques in the distributed processing framework of their choice, such as Dask, Spark, or TensorFlow, and use LensKit’s data processing and/or evaluation techniques to prepare and analyze the experiment and its results. LKPY aims to grow with its users, and to progressively get out of their way as they outgrow its built-in capabilities.

5.7 Quality Assurance

LensKit for Python includes an extensive test suite, following the testing practices that we found successful with the Java software [7]. These tests include:

- Unit tests for individual functions and modules.

⁵We only support Anaconda-based Python installs on macOS at this time due to the difficulty of supporting scientific packages on system Python.

⁶<https://anaconda.org/lenskit/lenskit>

⁷<https://pypi.org/project/lenskit/>

- End-to-end model build and serialization tests for all algorithm implementations. These tests train the model on the MovieLens Latest Small data set [16], a snapshot of which is included in the LensKit source tree, and confirm that the algorithm can successfully build, be saved to disk, and reloaded into memory without crashing or exhibiting other unexpected behavior.
- Accuracy tests that cross-validate models on the ML-100K data set and check that their rating prediction and top- N recommendation accuracy is within the expected range for that algorithm on that data set.
- Known-output tests that train the algorithm on a small MovieLens data set and compare its output with known-correct output. This is particularly for algorithms expected to exhibit the same behavior as their Java predecessors.

GitHub Actions run the test suite on all supported OS families and Python versions and measure test coverage. New features and fixes are implemented in feature branches and merged into the main development tree through pull requests, ensuring tests are run on all new code prior to integration. CodeClimate lint helps maintain code standards and catch some potential errors.

5.8 Documentation and Examples

LensKit for Python’s tutorial and reference documentation includes:

- A Getting Started guide that implements a simple recommender accuracy experiment in a self-contained Jupyter notebook, available on Google Collab.
- A user’s guide documenting LKPY’s features APIs.
- An example experiment, published on GitHub, that demonstrates a more realistic comparison of the effectiveness of recommendation algorithms on public data sets.
- Source code for experiments using LKPY [10, 36].

The examples also demonstrate LensKit use across a range of scales, from a single notebook to a full experiment that saves trained models and intermediate data to disk and coordinates experimental runs with Data Version Control [22]. The experiment for our extended author gender paper [10] is completely reproducible end-to-end with a single command (`dvc repro`) and approximately 2 days of compute time. It also demonstrates automatic hyperparameter optimization with `scikit-optimize` [17].

6 FUTURE DIRECTIONS

LensKit for Python development is ongoing, and there is much to add to the software. Two immediate priorities are to:

- Complete the recommendation web service to support online experiments and deployed applications.
- Build example content- and review-based recommendation algorithms to demonstrate working with additional data sources in the LensKit environment.

While LKPY development to date has primarily been by myself with support from the People and Information Research Team and collaborators, we welcome contributions from anyone subject to the project’s licensing requirements (MIT license) and code quality standards. Development and discussion are open to the public on GitHub and Google Groups.

7 COMPARISON TO EXISTING PACKAGES

Before building LKPY, I re-examined the landscape of existing software to determine if the needs we saw would be met by one of the other software packages. I was unable to find existing tooling that supports the LensKit project’s goals of flexible recommender systems experiments that leverage the PyData ecosystem.

Two of the most closely-related Python toolkits are `surprise` [20] and `PyRecLab` [34]. `Surprise` uses `numpy` and `scipy`, but not the fuller PyData ecosystem (such as `Pandas`) and has its own data set management. `PyRecLab` has extensive algorithm support and is an effective toolkit for teaching and rapid prototyping, but it is a C++ library with a Python API, rather than a native Python toolkit of re-configurable pieces, so it doesn’t use PyData tools or integrate with packages like `TensorFlow`, making it less appropriate for supporting the kinds of research I want to enable. Our research requires us to be able to write out the evaluation steps in our own code so that we can experiment with them more readily, and implement new algorithms without developing C++ proficiency.

The LKPY approach is perhaps most similar to that of `mrec`⁸, focusing on discrete steps enabled by separate tools. `mrec`, however, does not leverage contemporary Python data science tools, and does not seem to be under active maintenance. LKPY also focuses on Python as the scripting language for experiment control while `mrec` favors command-line tools orchestrated with shell scripts.

Many Python-based recommender systems research project seem to implement their their own evaluation procedures directly while building the recommendation algorithm in `scikit-learn` or one of the deep learning frameworks. LKPY integrates with such workflows, enabling them to leverage common, well-tested implementations of metrics and other experimental support code while continuing to use their existing data flows for the recommendation process.

8 CONCLUSION

My collaborators and I have learned many lessons developing, maintaining, using, and teaching the Java-based LensKit software. We came to the conclusion that, in its current form, it is not the needs of the recommender systems community well, and resources would be better spent on improving the research being done with the widely-used Python packages that drive much of modern data science. LensKit for Python (LKPY) provides code to make it easy for researchers, students, and developers to build and run experiments on recommender systems and prototype new ideas while leveraging the extensive software, documentation, and tutorials available for data science and machine learning in Python.

Where Java LensKit focused on providing building blocks for recommender *algorithms*, LensKit for Python provides building blocks for recommender *experiments*. Built on the PyData stack and organized around clear, explicit data processing pipelines with a minimum of bespoke concepts, I believe LKPY provides a solid foundation for new experimentation and concepts at all stages of the recommender systems research lifecycle. It is already seeing adoption in both research [2, 25, 38] and instructional settings [29, 37], in addition to supporting my own group’s research and educational efforts. We welcome contributions on GitHub, and invite feedback from the community to guide future development.

⁸<https://github.com/Mendeley/mrec>

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. IIS 17-51278. I thank my collaborators in the People and Information Research Team (PIReT) for their support in this project, particularly Sole Pera for her comments on the manuscript and my students Amifa Raj, Mucun Tian, and Carlos Segura Cerna for pushing the software forward; the GroupLens Center for Social Computing, particularly collaborators John Riedl & Michael Ludwig for supporting the Java software and Joe Konstan & Daniel Kluver for supporting both it and the LKPY re-envisioning; and my students for their use and feedback in CS 538. All errors of fact, judgement, or writing are my own.

REFERENCES

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/>. Software available from tensorflow.org.
- [2] Christian Ricardo Balseca Niñez and Daniel Andrés Fernández Peñañiel. 2019. *Estudio del estado del arte de la ciencia de datos aplicada a la neuroeconomía*. Master's thesis. Universidad Politécnica Salesiana. <http://hdl.handle.net/123456789/17605>
- [3] Alejandro Bellogin. 2012. *Performance Prediction and Evaluation in Recommender Systems: An Information Retrieval Perspective*. Ph.D. Dissertation. Universidad Autónoma de Madrid, Madrid, Spain. <http://ir.ii.uam.es/~alejandro/thesis/>
- [4] Léon Bottou, Jonas Peters, Joaquin Quiñero-Candela, Denis X Charles, D Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. 2013. Counterfactual Reasoning and Learning Systems: The Example of Computational Advertising. *Journal of Machine Learning Research* 14, 1 (2013), 3207–3260. <http://www.jmlr.org/papers/volume14/bottou13a/bottou13a.pdf>
- [5] 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. *CoRR* abs/1309.0238 (2013). <http://arxiv.org/abs/1309.0238>
- [6] Mukund Deshpande and George Karypis. 2004. Item-based Top-N Recommendation Algorithms. *ACM Transactions on Information Systems* 22, 1 (Jan. 2004), 143–177. <https://doi.org/10.1145/963770.963776>
- [7] Michael Ekstrand. 2016. Testing Recommenders. <https://buildingrecommenders.wordpress.com/2016/02/04/testing-recommenders/>
- [8] Michael Ekstrand, Michael Ludwig, Joseph A. Konstan, and John Riedl. 2011. Rethinking the Recommender Research Ecosystem: Reproducibility, Openness, and LensKit. In *Proc. RecSys '11*. ACM, 133–140. <https://doi.org/10.1145/2043932.2043958>
- [9] Michael Ekstrand, John Riedl, and Joseph A Konstan. 2010. Collaborative Filtering Recommender Systems. *Foundations and Trends® in Human-Computer Interaction* 4, 2 (2010), 81–173. <https://doi.org/10.1561/1100000009>
- [10] Michael D Ekstrand and Daniel Kluver. 2019. Exploring Author Gender in Book Rating and Recommendation. <https://md.ekstrand.net/pubs/bag-extended>
- [11] Michael D. Ekstrand and Michael Ludwig. 2016. Dependency Injection with Static Analysis and Context-Aware Policy. *Journal of Object Technology* 15, 1 (Feb. 2016), 1:1. <https://doi.org/10.5381/jot.2016.15.5.a1>
- [12] Michael D Ekstrand, Mucun Tian, Mohammed R Imran Kazi, Hoda Mehrpouyan, and Daniel Kluver. 2018. Exploring Author Gender in Book Rating and Recommendation. In *Proc. RecSys '18*. ACM. <https://doi.org/10.1145/3240323.3240373>
- [13] Simon Funk. 2006. Netflix Update: Try This at Home. <http://sifter.org/~simon/journal/20061211.html>
- [14] Prem Gopalan, Jake M Hofman, and David M Blei. 2013. Scalable Recommendation with Poisson Factorization. *arXiv:1311.1704 [cs, stat]* (Nov. 2013). <http://arxiv.org/abs/1311.1704>
- [15] Asela Gunawardana and Guy Shani. 2009. A Survey of Accuracy Evaluation Metrics of Recommendation Tasks. *Journal of Machine Learning Research* 10 (Dec. 2009), 2935–2962. <http://jmlr.org/papers/v10/gunawardana09a.html>
- [16] F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Transactions on Interactive Intelligent Systems* 5, 4 (Dec. 2015), 19:1–19:19. <https://doi.org/10.1145/2827872>
- [17] Tim Head, MechCoder, Gilles Louppe, Iaroslav Shcherbatyi, fcharras, Zé Vinicius, cmmalone, Christopher Schröder, nel215, Nuno Campos, Todd Young, Stefano Cereda, Thomas Fan, rene rex, Kejia (KJ) Shi, Justus Schwabedal, carlosdanielc-santos, Hvass-Labs, Mikhail Pak, SoManyUsernamesTaken, Fred Callaway, Loïc Estève, Lilian Besson, Mehdi Cherti, Karlson Pfannschmidt, Fabian Linzberger, Christophe Cauet, Anna Gut, Andreas Mueller, and Alexander Fabisch. 2018. *scikit-optimize/scikit-optimize: v0.5.2*. <https://doi.org/10.5281/zenodo.1207017>
- [18] Jon Herlocker, Joseph A. Konstan, and John Riedl. 2002. An Empirical Analysis of Design Choices in Neighborhood-Based Collaborative Filtering Algorithms. *Information Retrieval* 5, 4 (Oct. 2002), 287–310. <https://doi.org/10.1023/A:1020443909834>
- [19] Y Hu, Y Koren, and C Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *Proc. IEEE ICDM '08*. ieeexplore.ieee.org, 263–272. <https://doi.org/10.1109/ICDM.2008.22>
- [20] Nicolas Hug. 2017. Surprise, a Python Library for Recommender Systems. <http://surpriselib.com>
- [21] Joseph A Konstan, J D Walker, D Christopher Brooks, Keith Brown, and Michael D Ekstrand. 2015. Teaching Recommender Systems at Large Scale: Evaluation and Lessons Learned from a Hybrid MOOC. *ACM Transactions on Computer-Human Interaction* 22, 2 (April 2015), 10:1–10:23. <https://doi.org/10.1145/2728171>
- [22] Ruslan Kupriev, Dmitry Petrov, Ramón Valles, Paweł Redzyński, Casper da Costa-Luis, Alexander Schepanovski, Ivan Shcheklein, Saugat Pachhai, Jorge Orpínel, Fábio Santos, Aman Sharma, Zhanibek, Dani Hodovic, Earl, Andrew Grigorev, Nabanita Dash, George Vyshnya, maykulkarni, Vera, Max Hora, xliiv, Peter Rowlands, Wojciech Baranowski, Sanidhya Mangal, and Clemens Wolff. 2020. DVC: Data Version Control - Git for Data & Models. <https://doi.org/10.5281/zenodo.3813759>
- [23] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. 2015. Numba: A LLVM-based Python JIT Compiler. In *Proc. 2nd Workshop on LLVM Compiler Infrastructure in HPC (LLVM '15)*. ACM, 7:1–7:6. <https://doi.org/10.1145/2833157.2833162>
- [24] Wes McKinney. 2018. *Python for Data Analysis: Data Wrangling with pandas, NumPy, and IPython*. O'Reilly.
- [25] A Narayan. 2019. *What is the Value of Rating Obscure Items? An Analysis of the Effect of Less-Popular Items on Recommendation Quality*. Master's thesis. University of Minnesota. <https://hdl.handle.net/11299/206148>
- [26] Yiu-Kai Ng. 2020. CBRec: A Book Recommendation System for Children Using the Matrix Factorisation and Content-Based Filtering Approaches. *International Journal of Business Intelligence and Data Mining* 16, 2 (Jan. 2020), 129–149. <https://doi.org/10.1504/IJBIDM.2020.104738>
- [27] Travis E. Oliphant. 2006. *A Guide to NumPy*. Trelgol Publishing.
- [28] Travis E. Oliphant. 2007. Python for Scientific Computing. *Computing in Science Engineering* 9, 3 (May 2007), 10–20. <https://doi.org/10.1109/MCSE.2007.58>
- [29] PBS Digital Studios. 2019. Let's Make a Movie Recommendation System (LAB) #16. <https://www.pbs.org/video/3036002175/>
- [30] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Courville, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [31] István Pilászy, Dávid Zibriczky, and Dávid Tikk. 2010. Fast ALS-based Matrix Factorization for Explicit and Implicit Feedback Datasets. In *Proc. RecSys '10*. ACM, 71–78. <https://doi.org/10.1145/1864708.1864726>
- [32] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proc. UAI '09*. AAAI Press, 452–461. <https://doi.org/10.1145/1795114.1795167>
- [33] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based Collaborative Filtering Recommendation Algorithms. In *Proc. WWW '01*. ACM, 285–295. <https://doi.org/10.1145/371920.372071>
- [34] Gabriel Sepulveda, Vicente Dominguez, and Denis Parra. 2017. pyRecLab: A Software Library for Quick Prototyping of Recommender Systems. In *RecSys '17 Poster Proceedings*. <https://arxiv.org/abs/1706.06291>
- [35] Gábor Takács, István Pilászy, and Dávid Tikk. 2011. Applications of the conjugate gradient method for implicit feedback collaborative filtering. In *Proc. RecSys '11*. ACM, 297–300. <https://doi.org/10.1145/2043932.2043987>
- [36] Mucun Tian and Michael D Ekstrand. 2020. Estimating Error and Bias in Offline Evaluation Results. In *Proc. CHIIR '20*. ACM, 392–396. <https://doi.org/10.1145/3343413.3378004>
- [37] Ervin Varga. 2019. Recommender Systems. In *Practical Data Science with Python* 3. Apress, 317–339. <https://doi.org/10.1007/978-1-4842-4859-1>
- [38] X Wang and S Kadioglu. 2019. Bayesian Deep Learning Based Exploration-Exploitation for Personalized Recommendations. In *Proc. ICTAI '19*. IEEE, 1715–1719. <https://doi.org/10.1109/ICTAI.2019.00253>
- [39] Yunhong Zhou, Dennis Wilkinson, Robert Schreiber, and Rong Pan. 2008. Large-Scale Parallel Collaborative Filtering for the Netflix Prize. In *Proc. AAIM '08*. Springer, 337–348. https://doi.org/10.1007/978-3-540-68880-8_32
- [40] Cai-Nicolas Ziegler, Sean McNee, Joseph A Konstan, and Georg Lausen. 2005. Improving Recommendation Lists through Topic Diversification. In *Proc. WWW '05*. ACM, Chiba, Japan, 22–32. <https://doi.org/10.1145/1060745.1060754>