



LENSKIT IN ACTION

User-facing recommender applications

- MovieLens (~3k users/month)
- BookLens (built in to Twin Cities public libraries)
- Confer system for CSCW

Education

- Coursera MOOC
- Recommender classes @ UMN, Boise State

Research

- Used in 30+ published
- Offline evaluation, online user studies

REFLECTIONS

What Went Well

- Testing
- Java
- Modularity

What Didn't Work

- Opinionated Evaluation
- Indirect Configuration
- Implicit Behavior
- Technical Isolation

What Aged Poorly

- Java
- Modularity

WHAT WE NEED

Build on standard tools

Leverage existing software

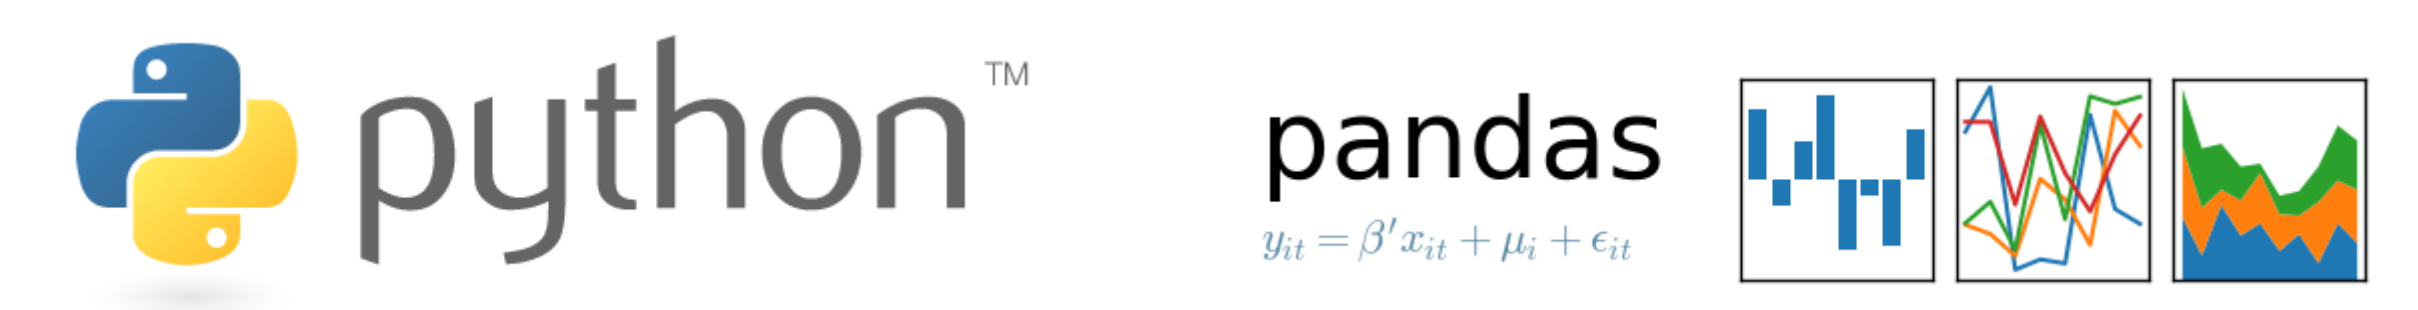
Expose the data pipeline

Explicit is better than implicit

Simple interfaces

Easy-to-use environment

THE NEXT GENERATION



LKPY FACILITIES

Data Preparation

- Prepare train-test splits for cross-validation
- Consumes and produces Pandas data frames

Algorithm APIs

- Predict
- Recommend
- Train, save, and load models

Evaluation Metrics

- Prediction accuracy (RMSE, MAE)
- Top-N (nDCG, Precision, Recall, MRR, etc.)
- Operate over Pandas data frames
- Works well with Jupyter

Classical CF Algorithms

- Optimized with Cython and OpenMP

Batch Utilities for Experiments

- Run over multiple train-test sets
- Combine results for analysis



LensKit provides building blocks for recommender algorithms.
LKPY is building blocks for recommender *experiments*.