Introduction to Recommender Systems

MICHAEL EKSTRAND

https://md.ekstrandom.net

OCTOBER 27, 2016

BOISE STATE UNIVERSITY

PIReT

People and Information Research Team

20 years of research in 35 minutes

About Me

- Assistant Prof, Boise State CS
- Run *People and Information Research Team* (PIReT) with Dr. Sole Pera
- Ph.D, University of Minnesota (in recommender systems)
- Human-computer interaction researcher
- Involved w/ recommenders since 2009

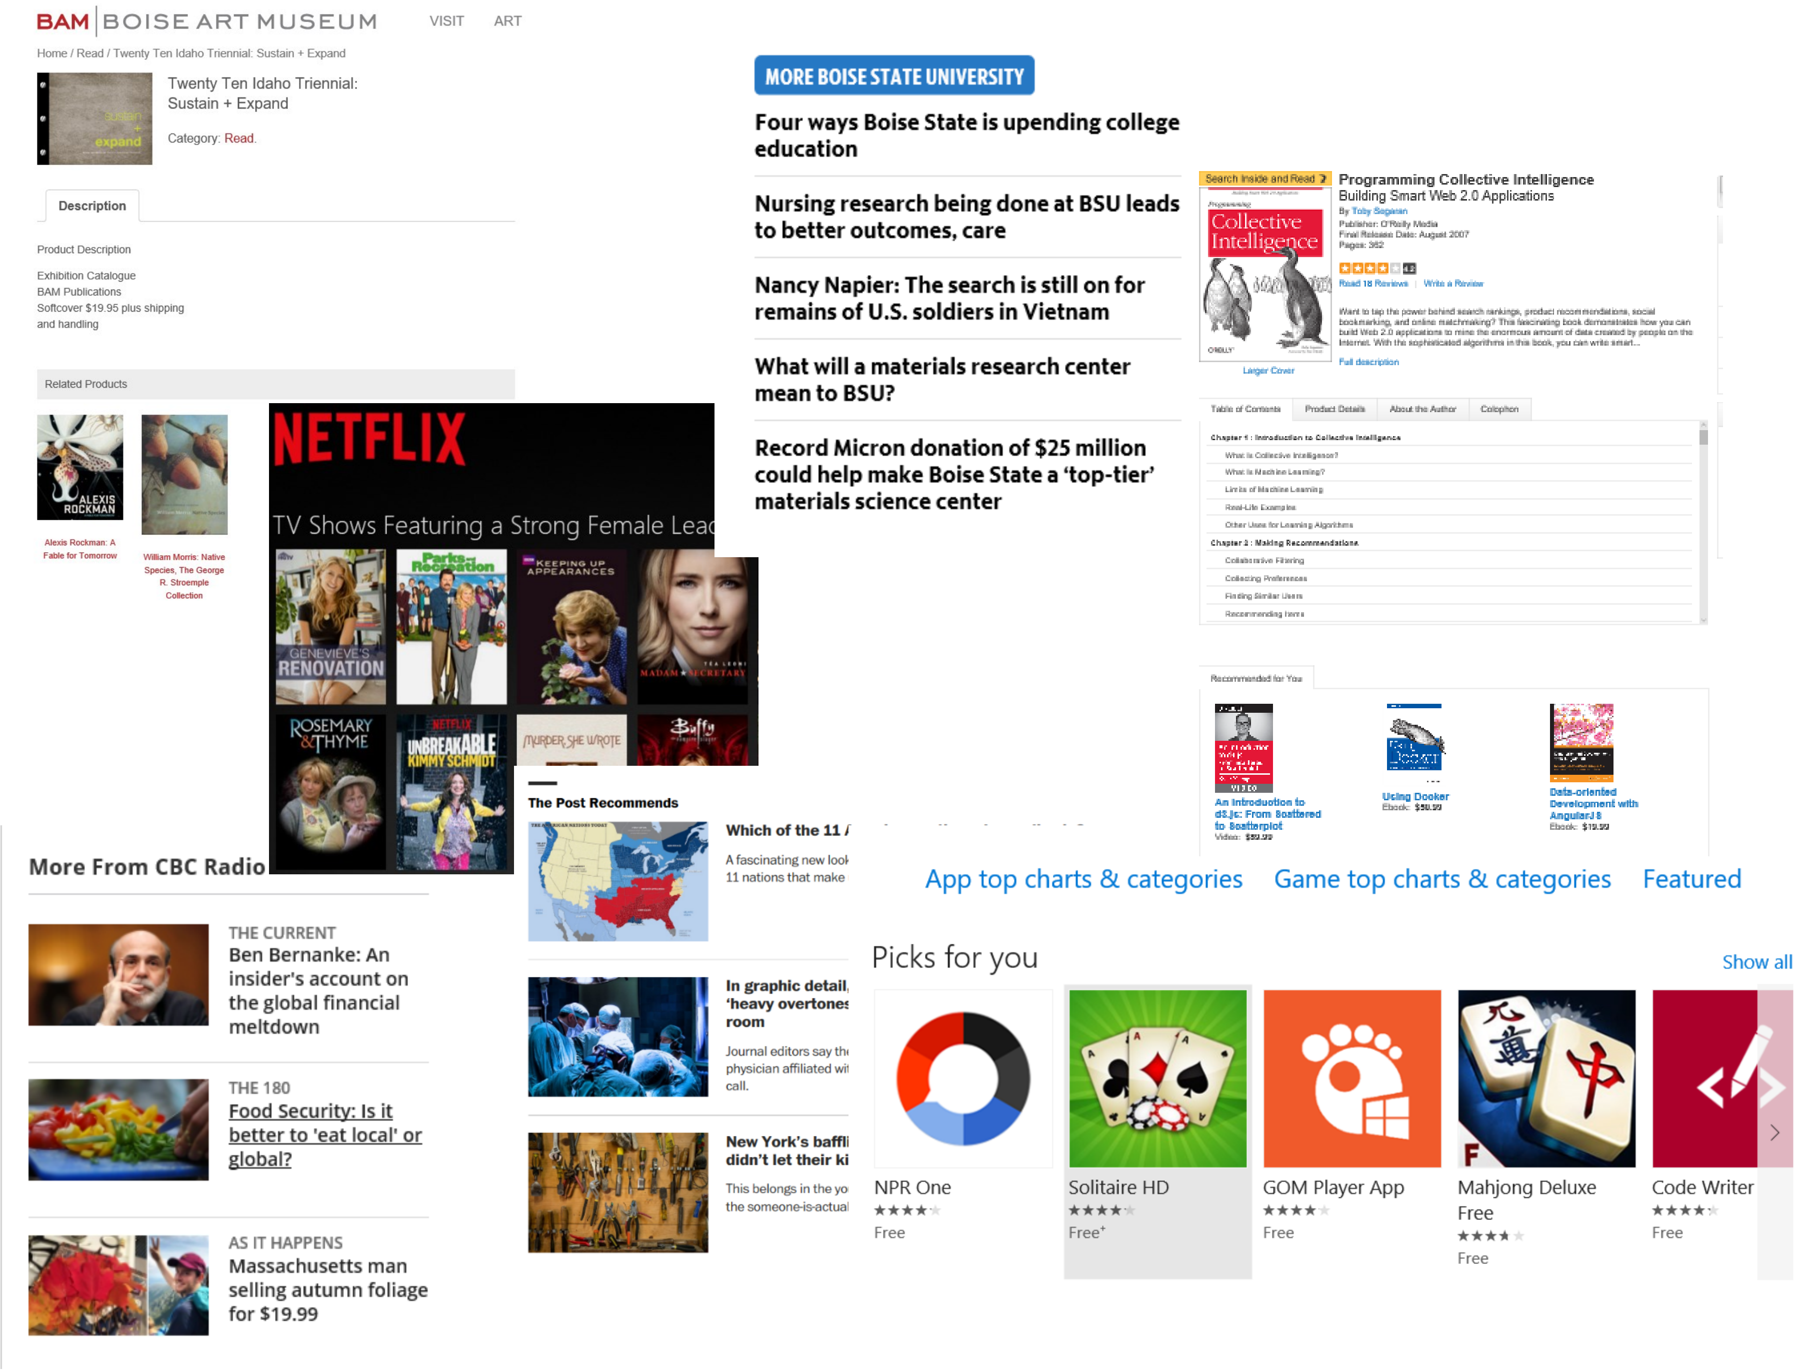
BAM | BOISE ART MUSEUM
VISIT
ART
Home / Read / Twenty Ten Idaho Triennial: Sustain + Expand
Twenty Ten Idaho Triennial: Sustain + Expand
Category: Read.
Description
Product Description
Exhibition Catalogue
BAM Publications
Softcover $19.95 plus shipping and handling
Related Products
Alexis Rockman: A Fable for Tomorrow
William Morris: Native Species, The George R. Stroemple Collection
NETFLIX
TV Shows Featuring a Strong Female Lead
MORE BOISE STATE UNIVERSITY
Four ways Boise State is upending college education
Nursing research being done at BSU leads to better outcomes, care
Nancy Napier: The search is still on for remains of U.S. soldiers in Vietnam
What will a materials research center mean to BSU?
Record Micron donation of $25 million could help make Boise State a 'top-tier' materials science center
Search Inside and Read
Programming Collective Intelligence
Building Smart Web 2.0 Applications
By Toby Segaran
Publisher: O'Reilly Media
Final Release Date: August 2007
Pages: 362
Read 18 Reviews | Write a Review
Want to tap the power behind search rankings, product recommendations, social bookmarking, and online matchmaking? This fascinating book demonstrates how you can build Web 2.0 applications to mine the enormous amount of data created by people on the Internet. With the sophisticated algorithms in this book, you can write smart...
Full description
Larger Cover
Table of Contents
Product Details
About the Author
Colophon
Chapter 1 : Introduction to Collective Intelligence
What is Collective Intelligence?
What is Machine Learning?
Limits of Machine Learning
Real-Life Examples
Other Uses for Learning Algorithms
Chapter 2 : Making Recommendations
Collaborative Filtering
Collecting Preferences
Finding Similar Users
Recommending Items
Recommended for You
An Introduction to d3.js: From Scattered to Scatterplot
Video: $59.99
Using Docker
Ebook: $59.99
Data-oriented Development with AngularJS
Ebook: $19.99
More From CBC Radio
THE CURRENT
Ben Bernanke: An insider's account on the global financial meltdown
THE 180
Food Security: Is it better to 'eat local' or global?
AS IT HAPPENS
Massachusetts man selling autumn foliage for $19.99
The Post Recommends
Which of the 11
A fascinating new look
11 nations that make
In graphic detail, 'heavy overtones room
Journal editors say th physician affiliated wi call.
New York's baffl didn't let their ki
This belongs in the yo the someone-is-actual
App top charts & categories
Game top charts & categories
Featured
Picks for you
Show all
NPR One
Free
Solitaire HD
Free⁺
GOM Player App
Free
Mahjong Deluxe Free
Free
Code Writer
Free

Recommender Architecture

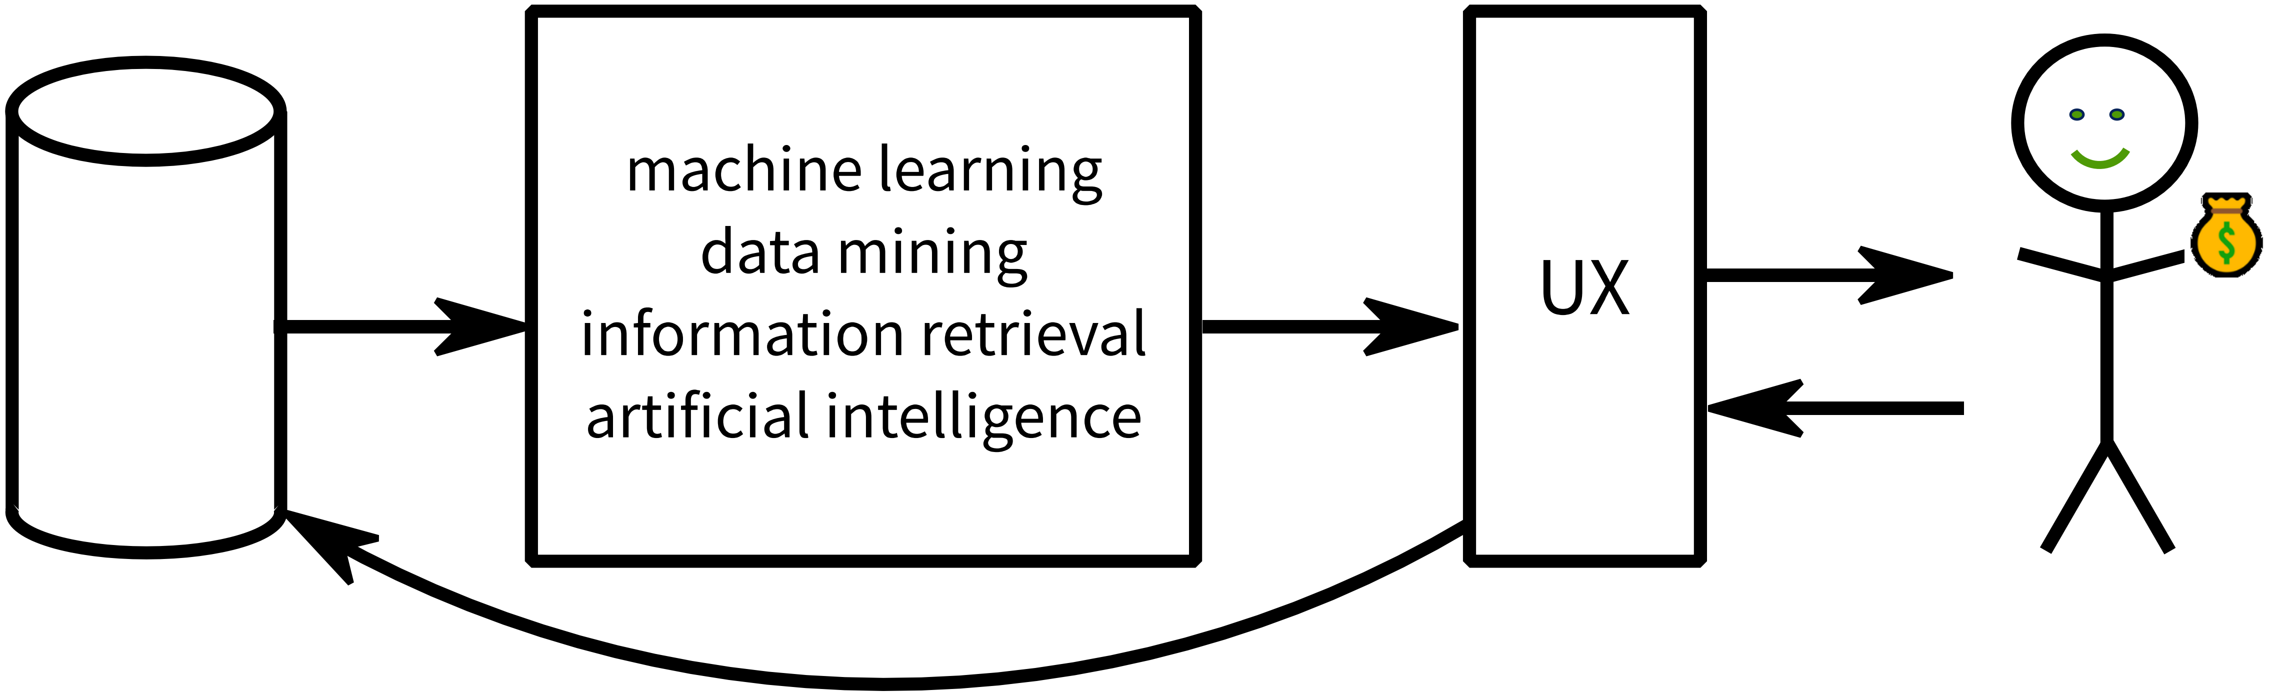

Where We’re Going

1. What does the recommender do?
2. What data do we use?
3. How do we compute recommendations?
4. How do we evaluate effectiveness?
5. What next?

Overview

1. **What does the recommender do?**
2. What data do we use?
3. How do we compute recommendations?
4. How do we evaluate effectiveness?
5. What next?

Recommender Vocabulary

Items

are the things we recommend

Users

receive & act on recommendations

provide input used for recommendation

Ratings

encode user preference for items

Recommender Tasks

Predict

estimate how much 👽 likes 🎁

can be predicted rating, purchase probability, score

Recommend

identify items that 👽 may like

maybe a 🚀?

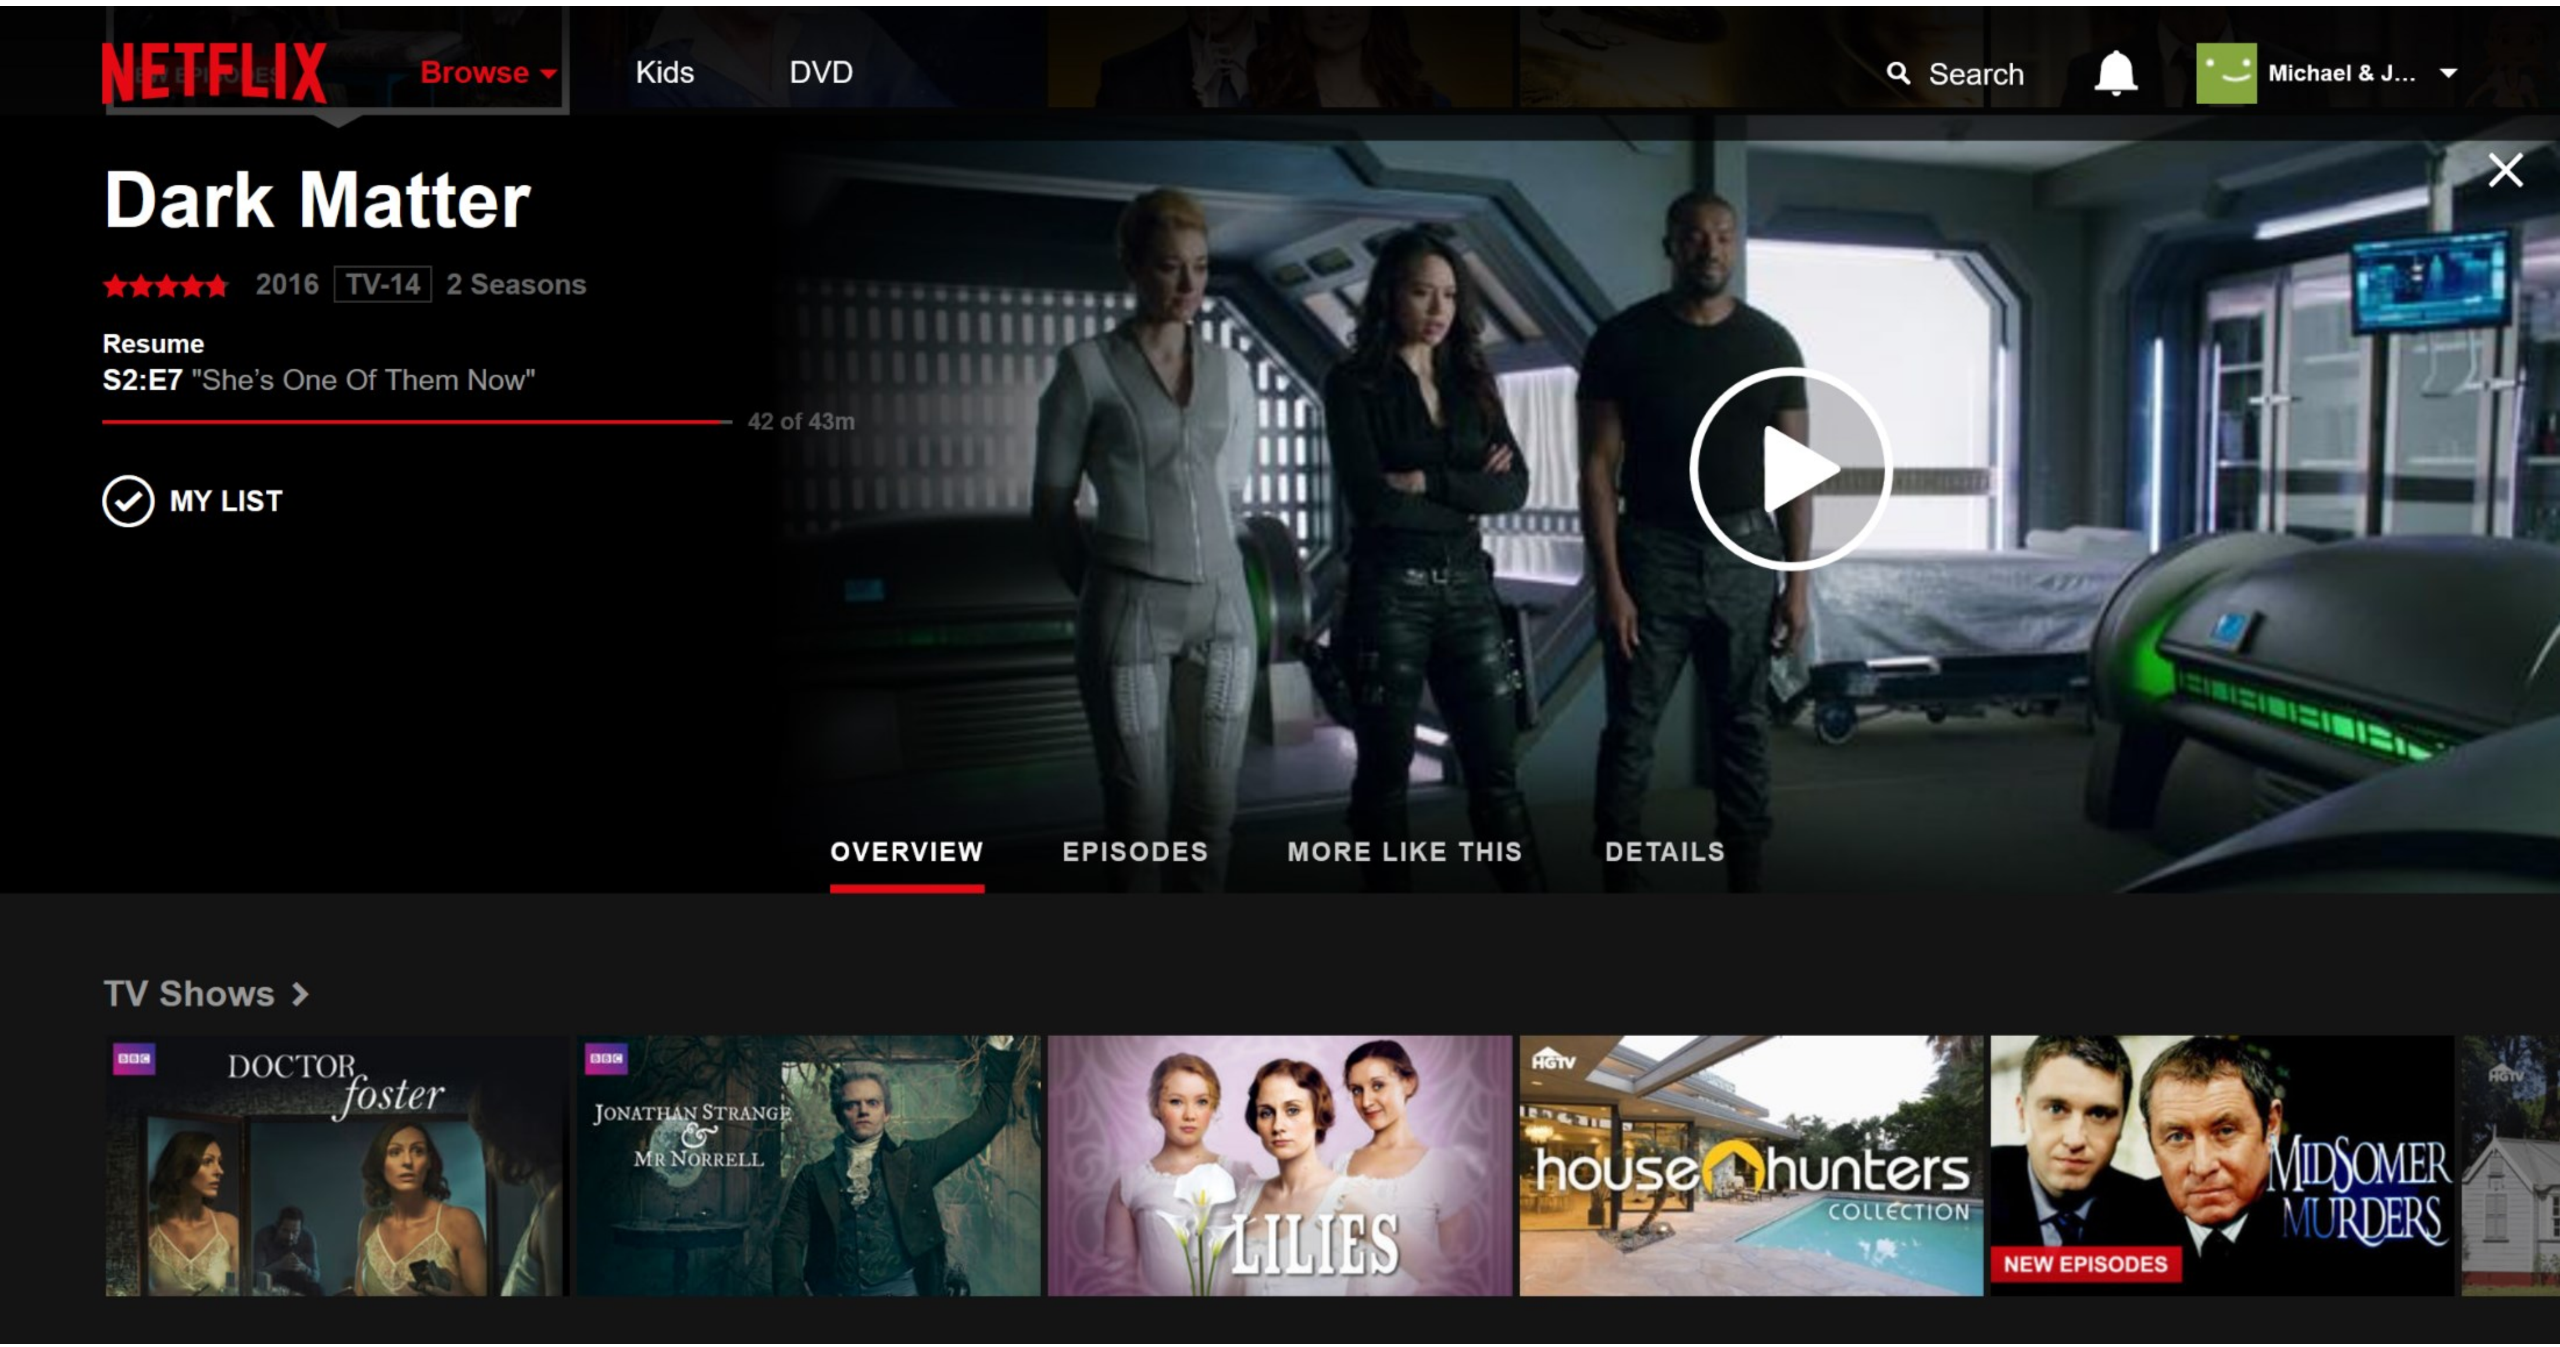
NETFLIX
Browse
Kids
DVD
Search
Michael & J...
Dark Matter
2016
TV-14
2 Seasons
Resume
S2:E7 "She's One Of Them Now"
42 of 43m
MY LIST
OVERVIEW
EPISODES
MORE LIKE THIS
DETAILS
TV Shows
DOCTOR foster
JONATHAN STRANGE & MR NORRELL
LILIES
house hunters COLLECTION
MIDSOMER MURDERS
NEW EPISODES

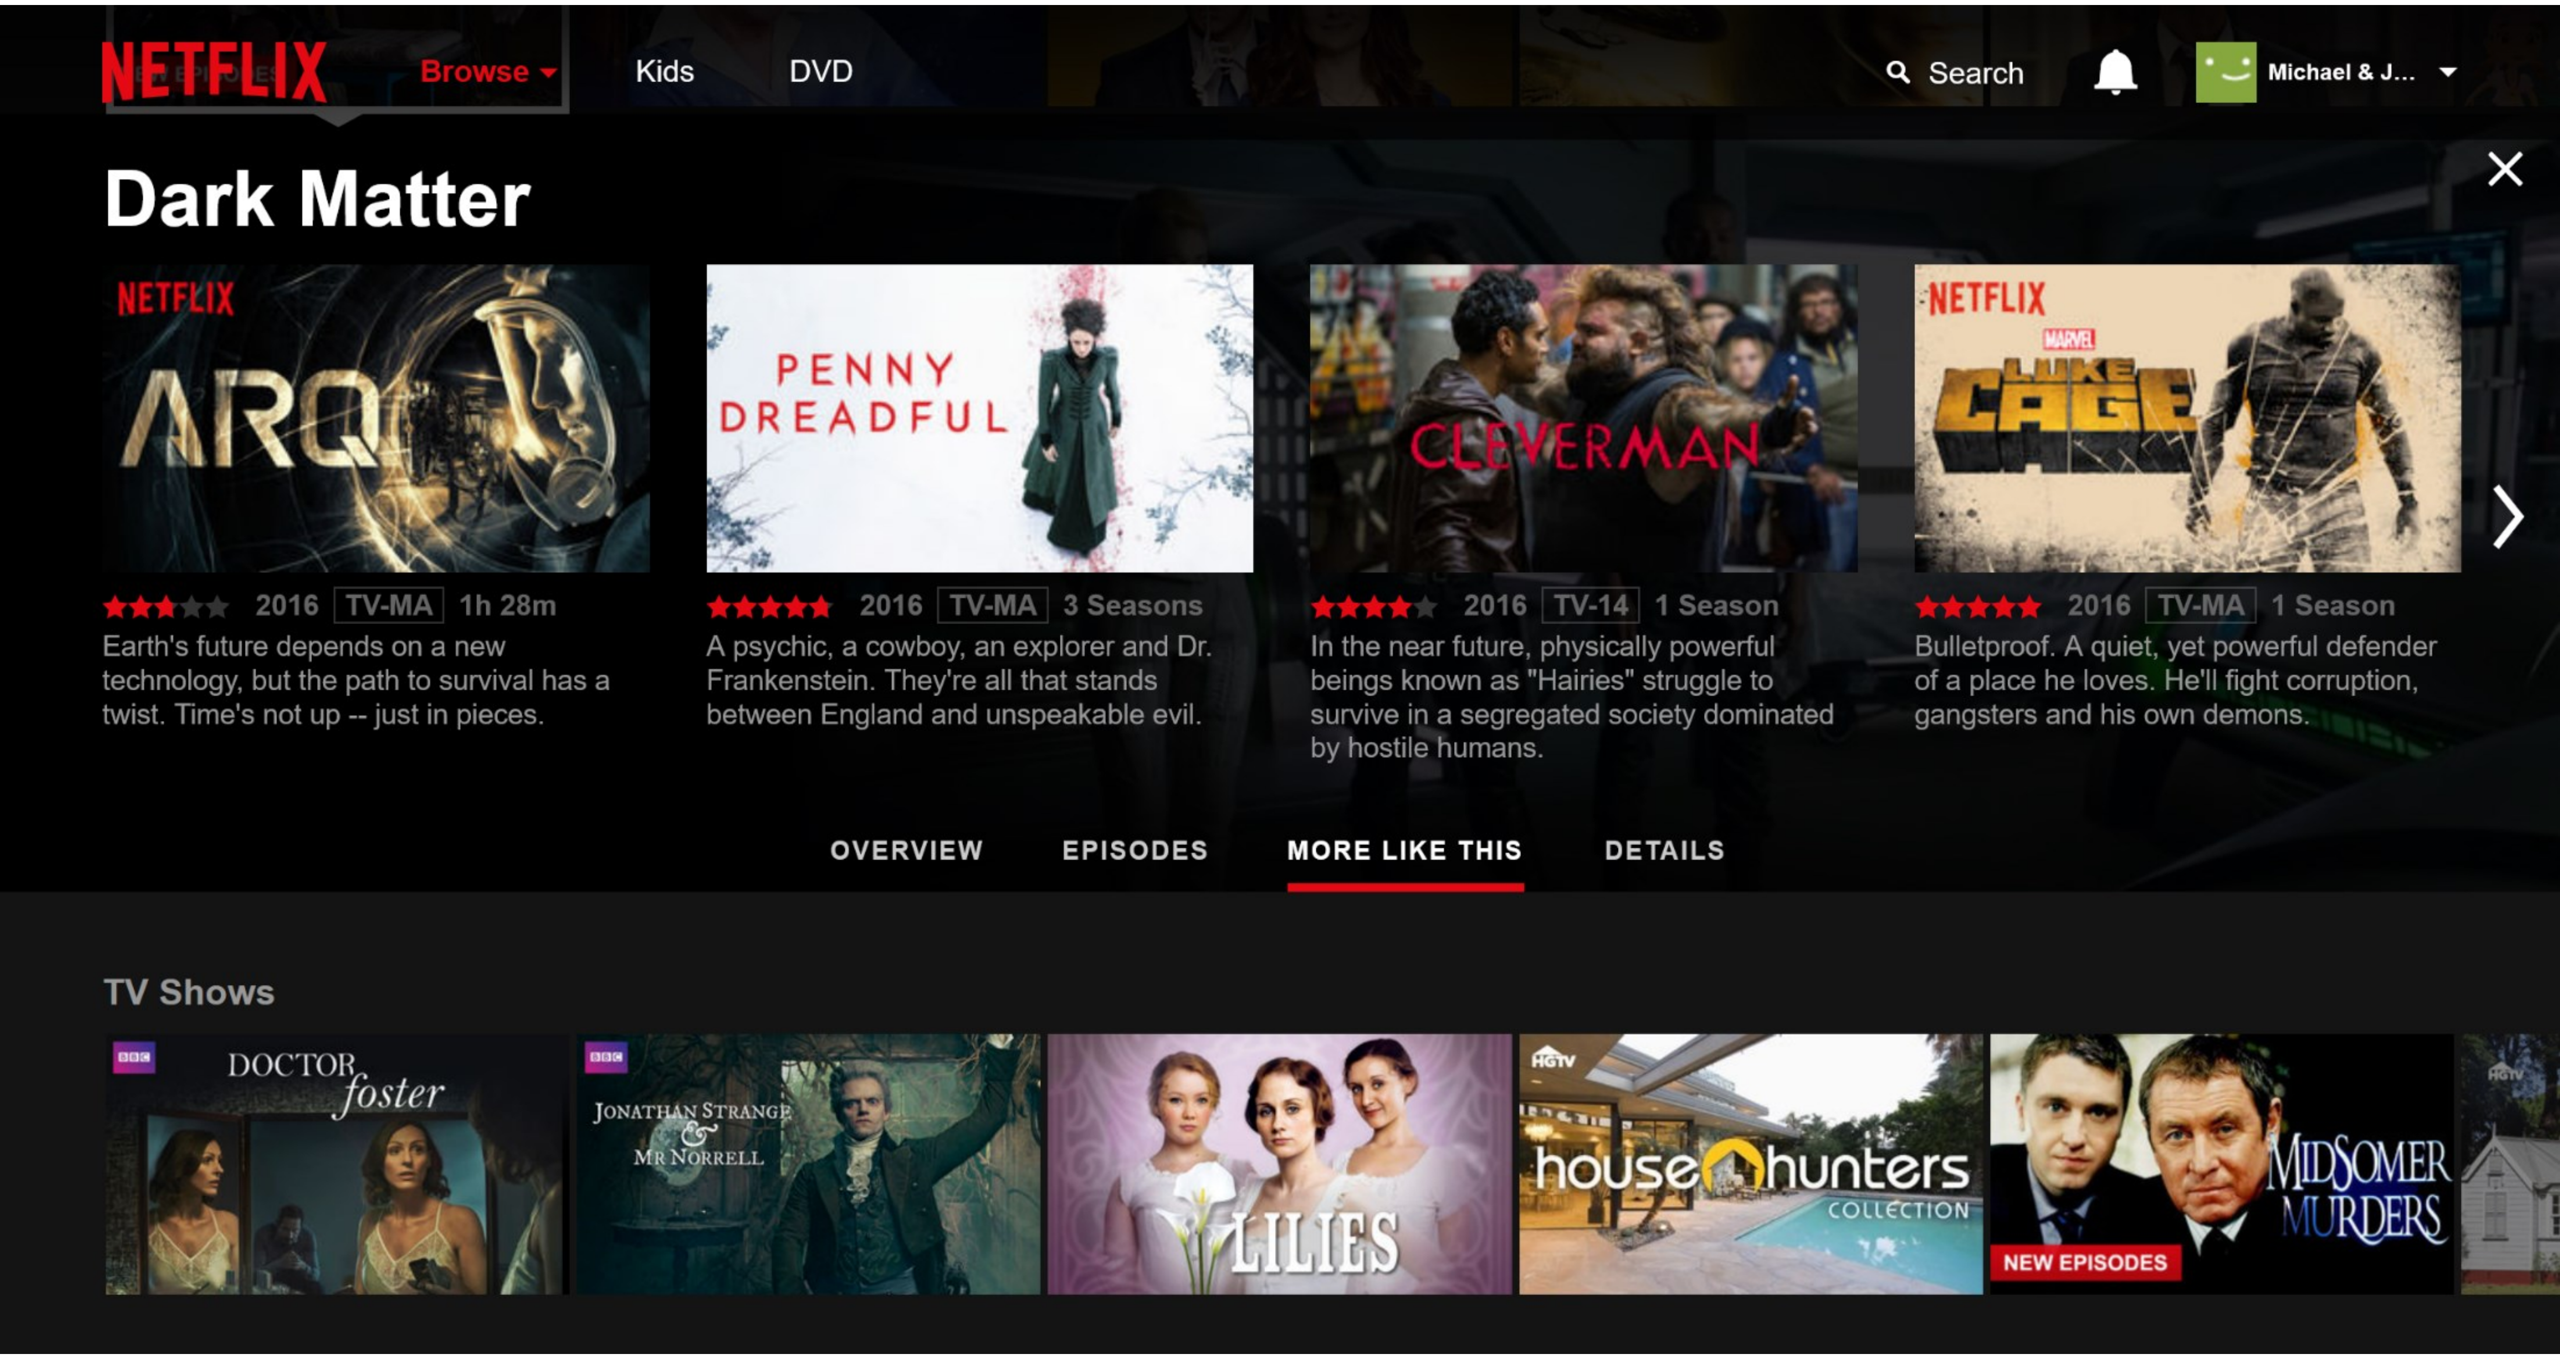
NETFLIX
Browse
Kids
DVD
Search
Michael & J...
Dark Matter
NETFLIX
ARQ
2016 TV-MA 1h 28m
Earth's future depends on a new technology, but the path to survival has a twist. Time's not up -- just in pieces.
PENNY DREADFUL
2016 TV-MA 3 Seasons
A psychic, a cowboy, an explorer and Dr. Frankenstein. They're all that stands between England and unspeakable evil.
CLEVERMAN
2016 TV-14 1 Season
In the near future, physically powerful beings known as "Hairies" struggle to survive in a segregated society dominated by hostile humans.
NETFLIX
MARVEL
LUKE CAGE
2016 TV-MA 1 Season
Bulletproof. A quiet, yet powerful defender of a place he loves. He'll fight corruption, gangsters and his own demons.
OVERVIEW
EPISODES
MORE LIKE THIS
DETAILS
TV Shows
DOCTOR foster
JONATHAN STRANGE & MR NORRELL
LILIES
house hunters
COLLECTION
MIDSOMER MURDERS
NEW EPISODES

Related Tasks

Support Decision

provide data to help 👽 decide whether to buy 🚀

- recommender algorithm outputs may be useful
- can recommend *information*

Well-studied in decision support literature; less in recsys

Overview

1. What does the recommender do?
2. **What data do we use?**
3. How do we compute recommendations?
4. How do we evaluate effectiveness?
5. What next?

Explicit Feedback

Users rate items

- 1-5 stars
- Thumbs up/down
- ‘Like’

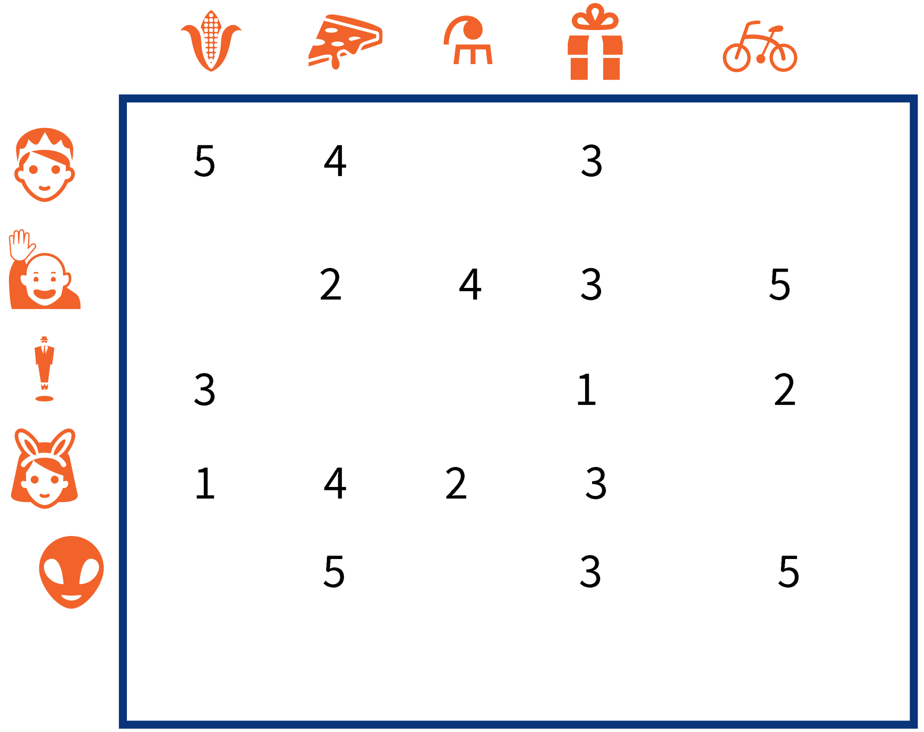

Rating Matrix

Implicit Feedback

Users take action

- Buy
- Click
- Watch
- Listen

Often used to synthesize a rating matrix

User and Item Data

- Demographics
- Text
- Metadata
- Tags
- External data (e.g. reviews)
- Whatever features are available!

Overview

1. What does the recommender do?
2. What data do we use?
3. **How do we compute recommendations?**
4. How do we evaluate effectiveness?
5. What next?

Broad Categories

Personalized or Non-personalized?

Content-Based or Collaborative?

Designed or Learned?

Non-Personalized Recommendation

- Most Popular
- Best Rated
- Human-Curated and Published
    - Movie reviews
    - Zagat guides

Light Personalization

- Demographics
  - People in your zip code like…
- Purchase-Based
  - People who bought the thing you're looking at bought…
  - Often done with **association rules**
    - Conditional probability or lift

Content-Based Recommendation

- Draws methods from IR and machine learning
- Uses item data
  - Metadata
  - Text content
  - Item features (e.g. acoustic analysis, mise-en-scène features)
- Traditionally ad-hoc
  - TF-IDF models over text representations
- Machine learning possible
  - Deep learning has promise here!

Collaborative Filtering

- Ignore the items!
  - Really. Ignore them.
- Learn everything from *user-item interactions*
  - Association rules +++
- Independent of item type/characteristics!
  - Good for *interchangeable* domains
  - Not so good for functional dependencies

Nearest-Neighbor CF

User-based

Find users 🕴 and 👯 who like the same things as 👽

Recommend that 👽 buy them!

$$s(i;u) = \frac{\sum_{v \in N(u;i)} (r_{vi} - \mu_v) w_{uv}}{\sum_{v \in N(u;i)} |w_{uv}|} + \mu_u$$

Nearest-Neighbor CF

Item-based

People who buy and/or 🚲 often buy 🌽

So recommend that 🧑 buy 🌽 too

$$s(i;u) = \frac{\sum_{j \in N(i;u)} r_{uj} w_{ij}}{\sum_{j \in N(i;u)} \left| w_{ij} \right|}$$

Matrix Factorization

$$R \approx PQ^{\mathrm{T}}$$

Perfectly clear!

Matrix Factorization

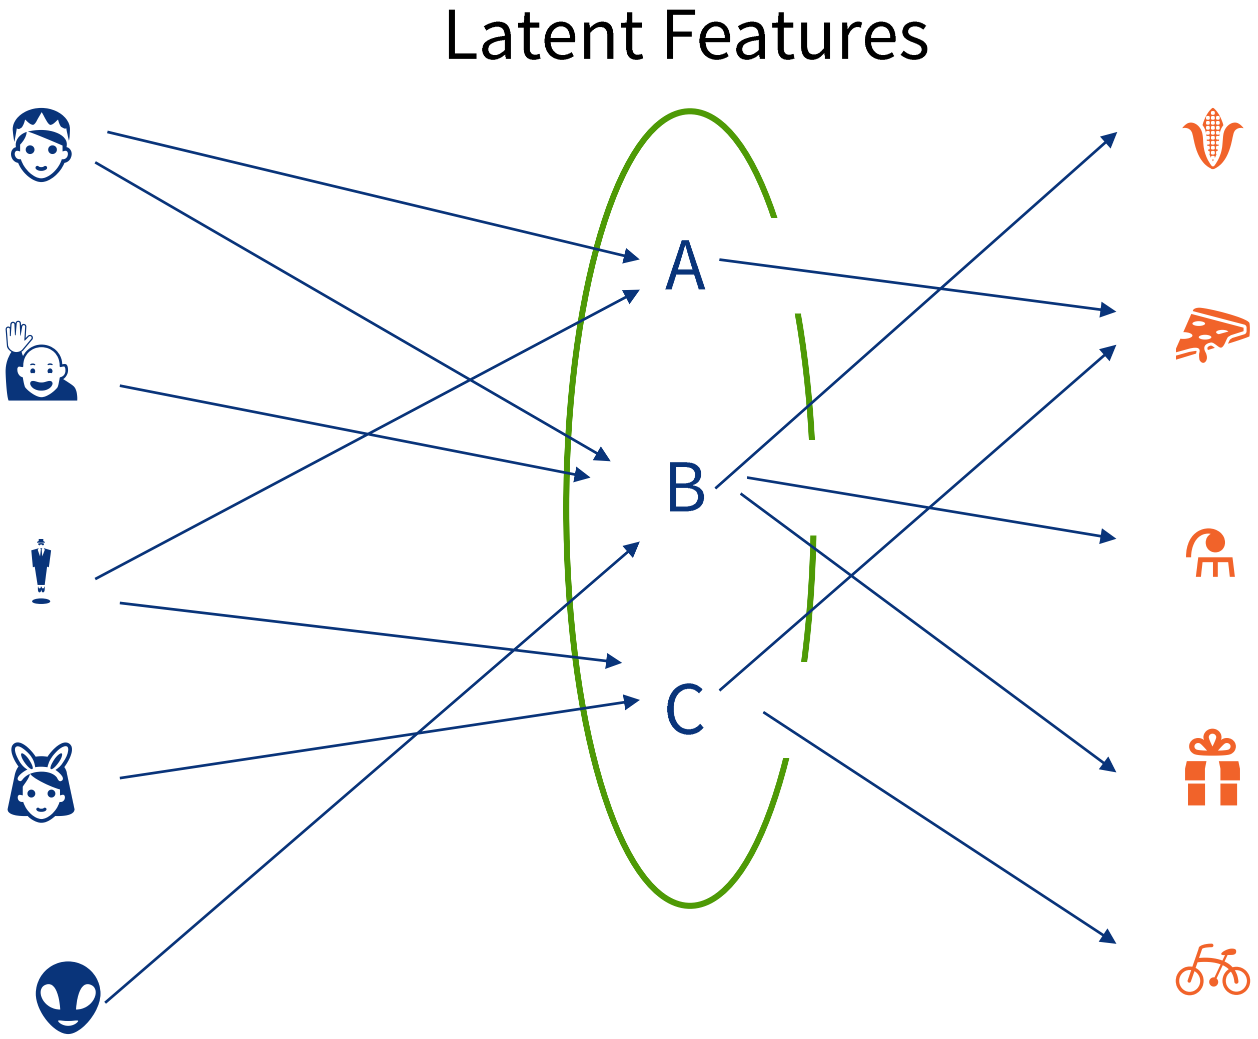

Matrix Factorization

$$R \approx PQ^{\mathrm{T}}$$

Derived from **Singular Value Decomposition**

$$s(i;u) = \sum_f p_{uf} q_{uf}$$

Typically a machine learning method (minimize SSE)

Extended Matrix Factorization

SVD++

Incorporates who-rated-what along with rating values

PMF

Reinterprets matrix factorization probabilistically

GPMF

Adds content features to PMF

All trained as machine learning approaches

Hybrid Recommenders

- Combine one or more algorithms
- Strengths of each!
- Common: **linear blends**
  - Often: blend weights depend on user/item features
- Most production recommenders are hybrids
- Typically learned machine-learning style

Learning to Rank

Learning to Score

Predict individual ratings/purchases

Minimize error

Not the real problem!

Learning to Rank

Produce good rankings (good item at top, most good items)

Hard, but more directly solves real problem

Overview

1. What does the recommender do?
2. What data do we use?
3. How do we compute recommendations?
4. **How do we evaluate effectiveness?**
5. What next?

Evaluation Strategies

Offline evaluation

Can we predict existing data?

Online evaluation (A/B testing)

What do users do?

Lab-style user studies

What do users think?

Offline Evaluation

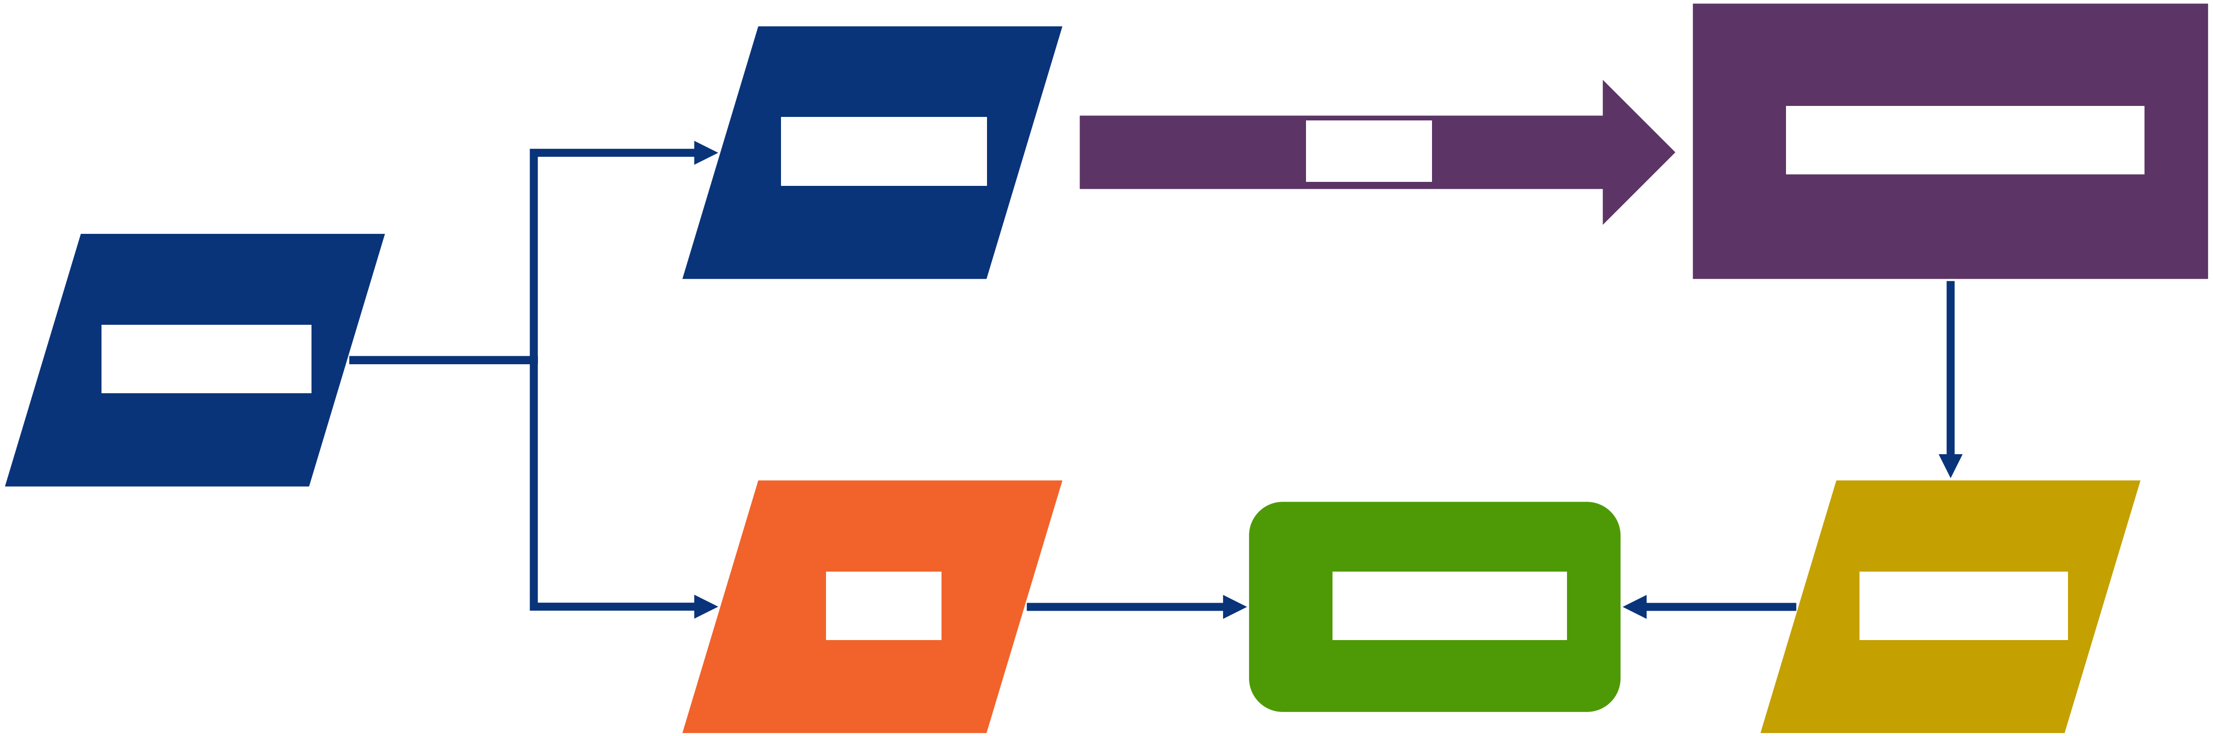

Offline Metrics

Prediction accuracy

RMSE, MAE

Top-N accuracy

Precision, Mean Reciprocal Rank, Mean Average Precision

Heavily dependent on setup

Other metrics

Diversity (intra-list similarity), novelty (popularity), etc.

Offline Problems

- Weak correlation with online performance
- Tests *predicting existing data*, not *finding new things*
- Data from recommenders is ‘tainted’

Still: necessary for pre-validation, tuning, etc.

  - Machine learning requires metrics!

A/B Testing

- Split users to site
- Give different treatments
- Measure result
  - Sales
  - Video plays
  - Bounce rate
- Apply statistics, rinse, repeat

User Studies

- Ask users what they think!
- Design rigorous surveys
  - Difficult
  - Many questions, tiring
  - Hard to get participants
- Benefit: subjective perception data

Overview

1. What does the recommender do?
2. What data do we use?
3. How do we compute recommendations?
4. How do we evaluate effectiveness?
5. **What next?**

Research Directions

- Ongoing work on algorithms
    - Often a small part of the problem
- Evaluation is hard
- What are drivers of user satisfaction, adoption, etc.?
- What is social impact?

Further Study

Build a tool

- LensKit (http://lenskit.org), several others in R, Python, etc.

Read a book

- *Practical Recommender Systems* (forthcoming from Manning)
- *Programming Collective Intelligence* (O'Reilly)

Take a class

- Boise State CS 597 this spring
- *Recommender Systems* on Coursera

Thank you!