NVIDIA GTC CONFERENCE

MULTI-OBJECTIVE OPTIMIZATION TO BOOST EXPLORATION IN RECOMMENDER SYSTEMS

SERDAR KADIOGLU

VP of AI, Fidelity Investments

Adjunct Faculty, Brown University

skadio.github.io

TV PAUSE
REC
SOCIAL VIEW
FOOTBALL
DISCOVER
APPS
SEN
RETURN
BACK
DIGITAL/
ANALOG
EXIT
1
ghi
4
pqrs
7
RADIO
TV

Building new applications
with limited or no training data
remains a common challenge in the industry.

Apriori decision in any recommender system:
what is the universe of items $\mathcal{I}$ to consider?

Trade-offs: Item Selection Matters

- Number of items vs. experimentation time
- Number of items vs. diversity and learning objectives
- Item mixture and coverage of outcomes and audience
- Time-to-market and onboarding effort: review, publish, register etc.

Question: What is the right mix of items to start the initial experimentation?

Answer: Principled approach for Item Selection

Today's Talk

- Discrete Optimization
- Natural Language Processing
- Item Selection Problem (ISP)
- Unsupervised Learning

Roadmap

1. Problem Definition
Introduce the ISP problem
Illustrative example

2. High-Level System Design
ISP in the context of
Recommender system pipelines

3. Solution Approach
Multi-objective
Optimization Framework

4. Benefits of the Approach
Numerical results on
recommendation benchmarks

5. Human-in-the Loop Decision Making
Empower business users with
interactive item selection

Item Selection Problem

Illustrative Example

Problem Definition

High-Level System Design

Illustrative Example

Item Metadata Categories and Labels

Category	Label
Title	Spectre – 007
Storyline	A cryptic message sends James Bond on a trail to uncover the existence of ..
Genre	Spy
Language	English
Director	Sam Mendes
Producer	MGM
Stars	Daniel Craig, Christoph Waltz, Monica Bellucci, ..

Item Selection Problem (ISP)

Movie Recommenders

- $\boldsymbol{I}$: All available movies
- $\boldsymbol{S}$: Subset of movies $\boldsymbol{S} \subset \boldsymbol{I}$ for experimentation
- $\boldsymbol{C}$: Categories e.g., language, genre, producer
- $\boldsymbol{L_c}$: Labels within each category e.g., action, comedy for genre
- $\boldsymbol{E(I)}$: Deep latent representation powered by GPU Technology
  text e.g., movie reviews, image e.g., cover art,
  audio e.g., soundtrack, video e.g., trailer
- **Goal:** find a minimum subset and maximize label coverage and diversity

IJCAI 2021, Item Selection meets Active Learning, *Kadioglu et. al.*
CPAIOR 2021, Optimized Item Selection for Recommender Systems, *Kadioglu et. al.*

High-Level System Design

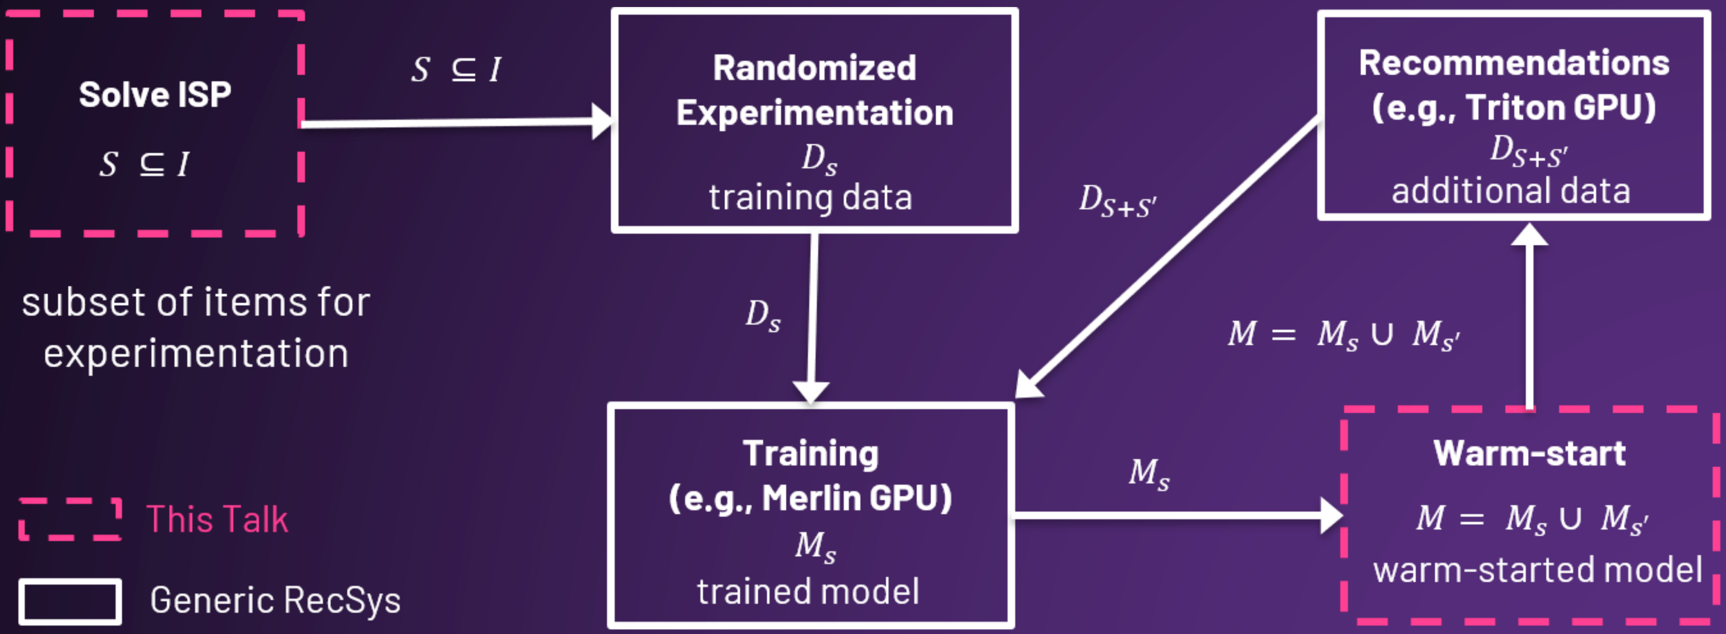

Solution Approach

Multi-objective optimization framework with cover formulation

Multi-Objective Optimization Framework

MINIMIZE
ITEM SIZE

MAXIMIZE
DIVERSITY

MAXIMIZE
LABEL COVERAGE

ITEM EMBEDDING
SPACE

TextWiser Library, AAAI 2021, *Kilitcioglu et. al.*

Multi-Objective Optimization Framework

1	**Minimizing the Subset Size**	Use standard set covering formulation to select subset of items that cover all predefined labels
2	**Maximizing Diversity**	Reformulate the loss function to consider items that are most spread out in embedding space $E(I)$ while still covering all labels
3	**Bounded Experiment Time**	Constrain number of selected items from #2 while maximizing the number of labels covered

The High Level Idea

$$\min \sum_{i}^{I} c_i x_i$$

$$\sum_{i \in I} M_{l,i} x_i \geq 1 \qquad \forall l \in L_c, \forall c \in C$$

$$x_i \in \{0,1\}, c_i = 1 \qquad \forall i \in I$$

$$\max \sum_{l \in L_c, c \in C} is\_label\_covered_l$$

$$\sum_{i \in I} x_i \leq t$$

$$M_{l,i} x_i \leq is\_label\_covered_l \quad \forall l \in L_c,\ \forall c \in C\ \forall i \in I$$

$$\sum_{i \in I} M_{l,i} x_i \geq is\_label\_covered_l \quad \forall l \in L_c,\ \forall c \in C$$

$$x_i \in \{0,1\} \quad \forall i \in I$$

$$is\_label\_covered_l \in \{0,1\} \quad \forall l \in L_c,\ \forall c \in C$$

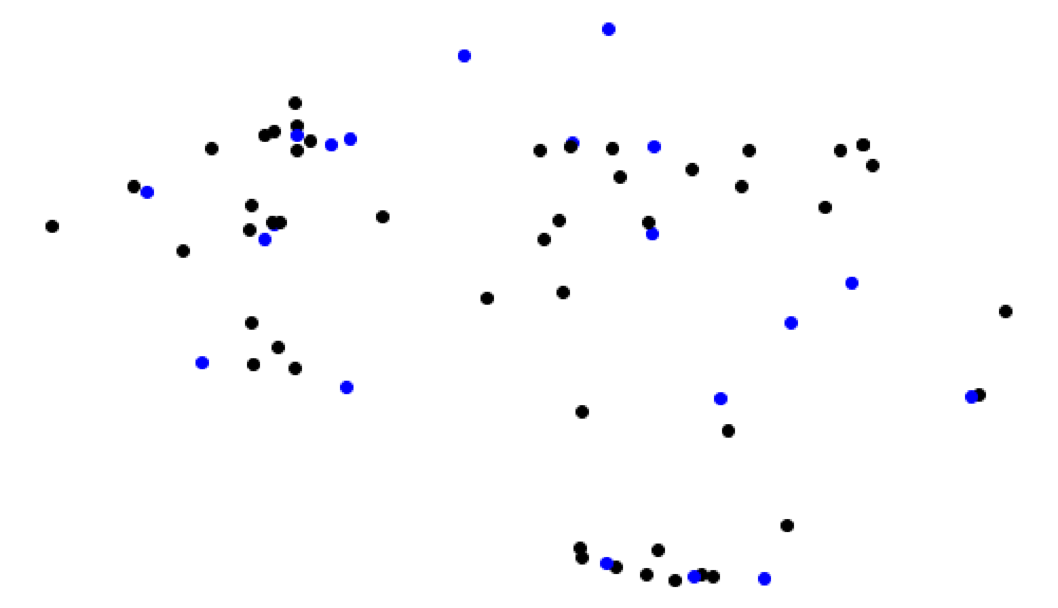

IJCAI 2021, Item Selection meets Active Learning, *Kadioglu et. al.*

CPAIOR 2021, Optimized Item Selection for Recommender Systems, *Kadioglu et. al.*

Benefits of the Approach

[Q1] How much speed-up is possible?

[Q2] How sensitive is the embedding space?

Movie and Book Recommenders

	Items	Categories	Unique Labels
MovieLens Movie Recommenders[1]	1,000 10,000	Genre, Producer, Language, Genre x Language	473 1,011
Goodreads Book Recommenders[2]	1,000 10,000	Genre, Publisher, Genre X Publisher	574 1,322

[1] Harper, F., Konstan, J.: The movielens datasets: History and context.
[2] Wan, M., McAuley, J.J.: Item recommendation on monotonic behavior chains.

[Q1] How much speed-up?

Significantly less number of items ***while covering all labels***
against Randomized, Greedy and Clustering baselines

~75% Reduction (4X faster)

Small dataset with 1000 items

~90% Reduction (10X faster)

Large dataset with 1000 items

[Q2] Sensitivity of Embeddings

Similar unit coverage for different embedding methods

Deep/GPU methods lead to rich latent representations

TextWiser Embedding[1]	1K	10K
TFIDF[2]	1.2	0.5
Word2Vec[3]	1.4	0.7
GloVe[4]	1.4	0.6
Byte-Pair[5]	1.3	0.6

[1] Kilitcioglu, D., Kadioglu, S. Representing the Unification of Text Featurization using a Context-Free Grammar. AAAI 2021

[2] Jones, K.S.: A statistical interpretation of term specificity and its application in retrieval.

[3] Grave, E., Bojanowski, P., Gupta, P., Joulin, A., Mikolov, T.: Learning word vectors for 157 languages. ACL 2018

[4] Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. ACL 2014

[5] Sennrich, R., Haddow, B., Birch, A.: Neural machine translation of rare words with subword units. ACL 2016

Human-in-the-Loop Decision Making

Interactive item selection

Dynamic settings

Active learning

Optimized Content Selection

2D embeddings of contents

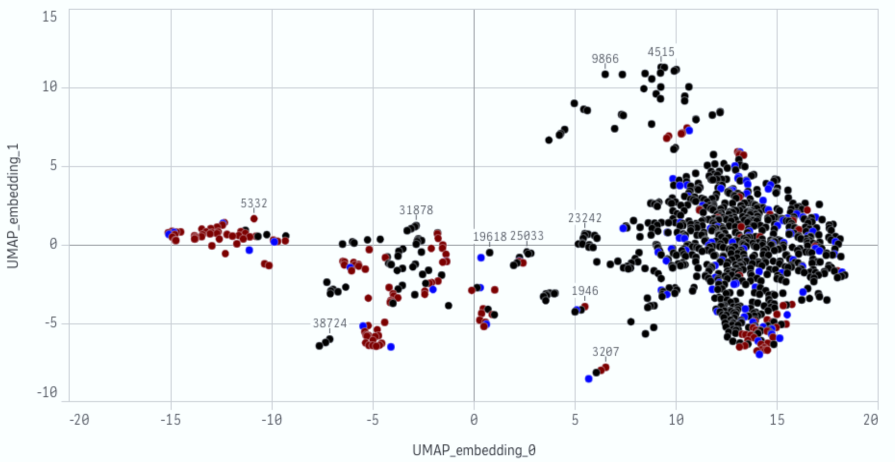

Content data

conte...	Genre	Title	Description	se
19173	poetry	The Divine Comedy Vol. 2: Purgatory	_The Divine Comedy_ is perhaps the greatest Christian classic ever written, and probably the greatest adventure story ever told. Dante wrote	e
28716	fiction	Blood Meridian	'Blood Meridian' presents an epic novel of the violent American West. The story is loosely based on accounts of murder along the border	e
45205	fiction	No Name	'Mr Vanstone's daughters are Nobody's Children'. Magdalen Vanstone and her sister Norah learn	e
10404	fiction	Doctor Who: Cat's Cradle-Witch Mark	'Spare no sympathy for those creatures. They were witches, they deserved to die.' A coach crashes on the M40. All the passengers	e
11499	fiction	The Promise	"A superb mirror of a place, a time, and a group of people who capture our immediate interest and hold it tightly." --The Philadelphia Inquirer	e
45066	fiction	The Mill on the Floss	"Backgrounds" includes fifteen letters from the 1859-69 period centering on the novel's content and composition; "Brother and Sister"	e
35958	children	Sagwa the Chinese Siamese Cat	"Before you go out into the world," Ming Miao told her five kittens, "you must know the true story of your ancestors...."	e
38700	fiction	Rabbit Hole	"David Lindsay-Abaire has crafted a drama that's not just a departure but a revelation--an intensely emotional examination of grief, laced	e

Number of Ratings

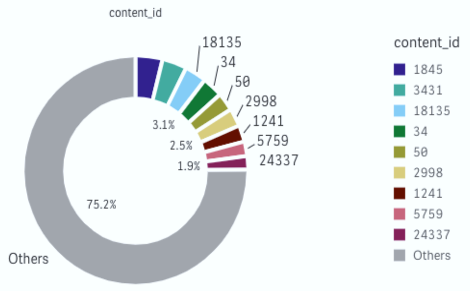

Number of Ratings v.s. Year Published

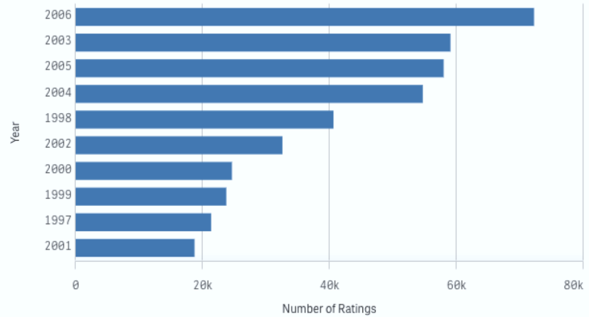

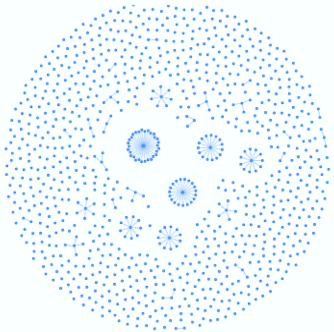

Special thanks to our collaborators!

- [IJCAI’21] Active learning meets optimized item selection
- [CPAIOR’21] Optimized item selection to boost exploration for recommender systems
- [AAAI’21] Representing the unification of text featurization using a context-free grammar
- [AAAI’22] Seq2Pat: Sequence-to-Pattern generation
- [AAAI’22] Dichotomic pattern mining for prediction from clickstream datasets
- [ICMLA’21] Surrogate ground truth to enhance binary fairness in uplift modelling
- [IJAIT’21] Parallelizable contextual multi-armed bandits
- [JDSA’21] Modeling uncertainty to improve personalized recommendations via Bayesian DL
- [ICTAI’19] Bayesian DL-based exploration-exploitation for personalized recommendations
-

- Recommenders Mab2Rec: https://github.com/fidelity/mab2rec
- Multi-armed Bandits MABWiser: https://github.com/fidelity/mabwiser
- Text/NLP TextWiser: https://github.com/fidelity/textwiser
- Pattern Mining Seq2Pat: https://github.com/fidelity/seq2pat
- Feature Selection Selective: https://github.com/fidelity/selective
- AI Fairness & Bias Jurity: https://github.com/fidelity/jurity

skadio.github.io

Du Cheng Doruk Kilitcioglu

Bernard Kleynhans Filip Michalsky Xin Wang