IJCAI'21, CPAIOR'21

Optimized Item Selection to Boost Exploration for Recommender Systems

Serdar Kadıoğlu
Adjunct Assoc. Prof, Dept. of Computer Science, Brown University
Group Vice President, AI Center of Excellence, Fidelity Investments
skadio.github.io

Nov. 2022 | Bucknell University


SOCIAL VIEW
FOOTBALL
DISCOVER
APPS
SEN
RETURN
BACK
DIGITAL/
ANALOG
EXIT
1
4
7
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?

Today's Talk

Discrete Optimization

Natural Language Processing

Item Selection Problem (ISP)

Unsupervised Learning

Roadmap

1. Problem Definition
Illustrative example
Introduce the ISP problem

3. Solution Approach
Multi-objective optimization with warm-starts

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

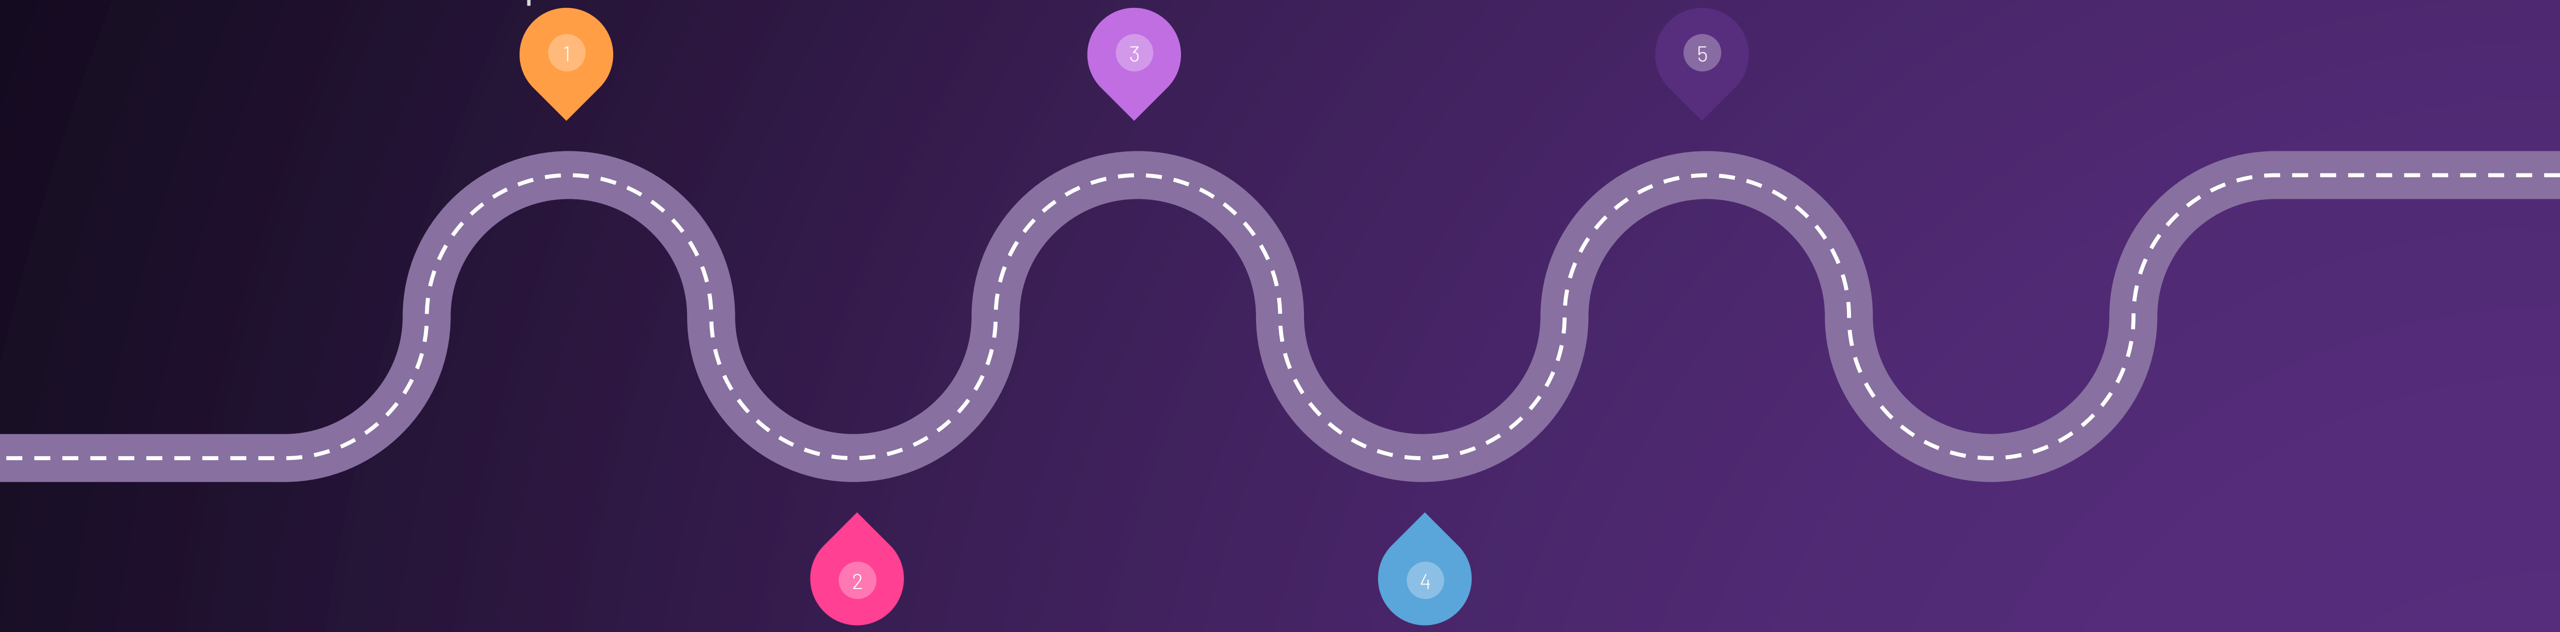

2. High-Level System Design
ISP in the context of Recommender pipelines

4. Benefits of the Approach
Numerical results on recommendation benchmarks

Item Selection Problem

Problem Definition, Illustrative Example

High-Level System Design

Item Selection Problem (ISP)

Problem Definition and Illustrative Example

Given a set of items I, the goal of the item selection problem (ISP) is to find the minimum subset $S \subseteq I$ that covers a set of labels L_c within each category $c \in C$ while maximizing the diversity of the selection S in the latent embedding space of items $E(I)$

Illustrative Example: Movie Recommendations

- $\boldsymbol{I}$: All available movie titles to be recommended
- $\boldsymbol{S}$: Subset of movie titles to be included in experimentation
- $\boldsymbol{C}$: Categories of interest (e.g., language, genre, producer)
- $\boldsymbol{L_c}$: Labels within each category (e.g., action, comedy for genre)
- $\boldsymbol{E}(\boldsymbol{I})$: Latent representation based on textual data (e.g., synopses, movie reviews), image data (e.g., cover art), audio data (e.g., soundtracks), video data (e.g., trailer)

High-Level System Design

Recommender System Components

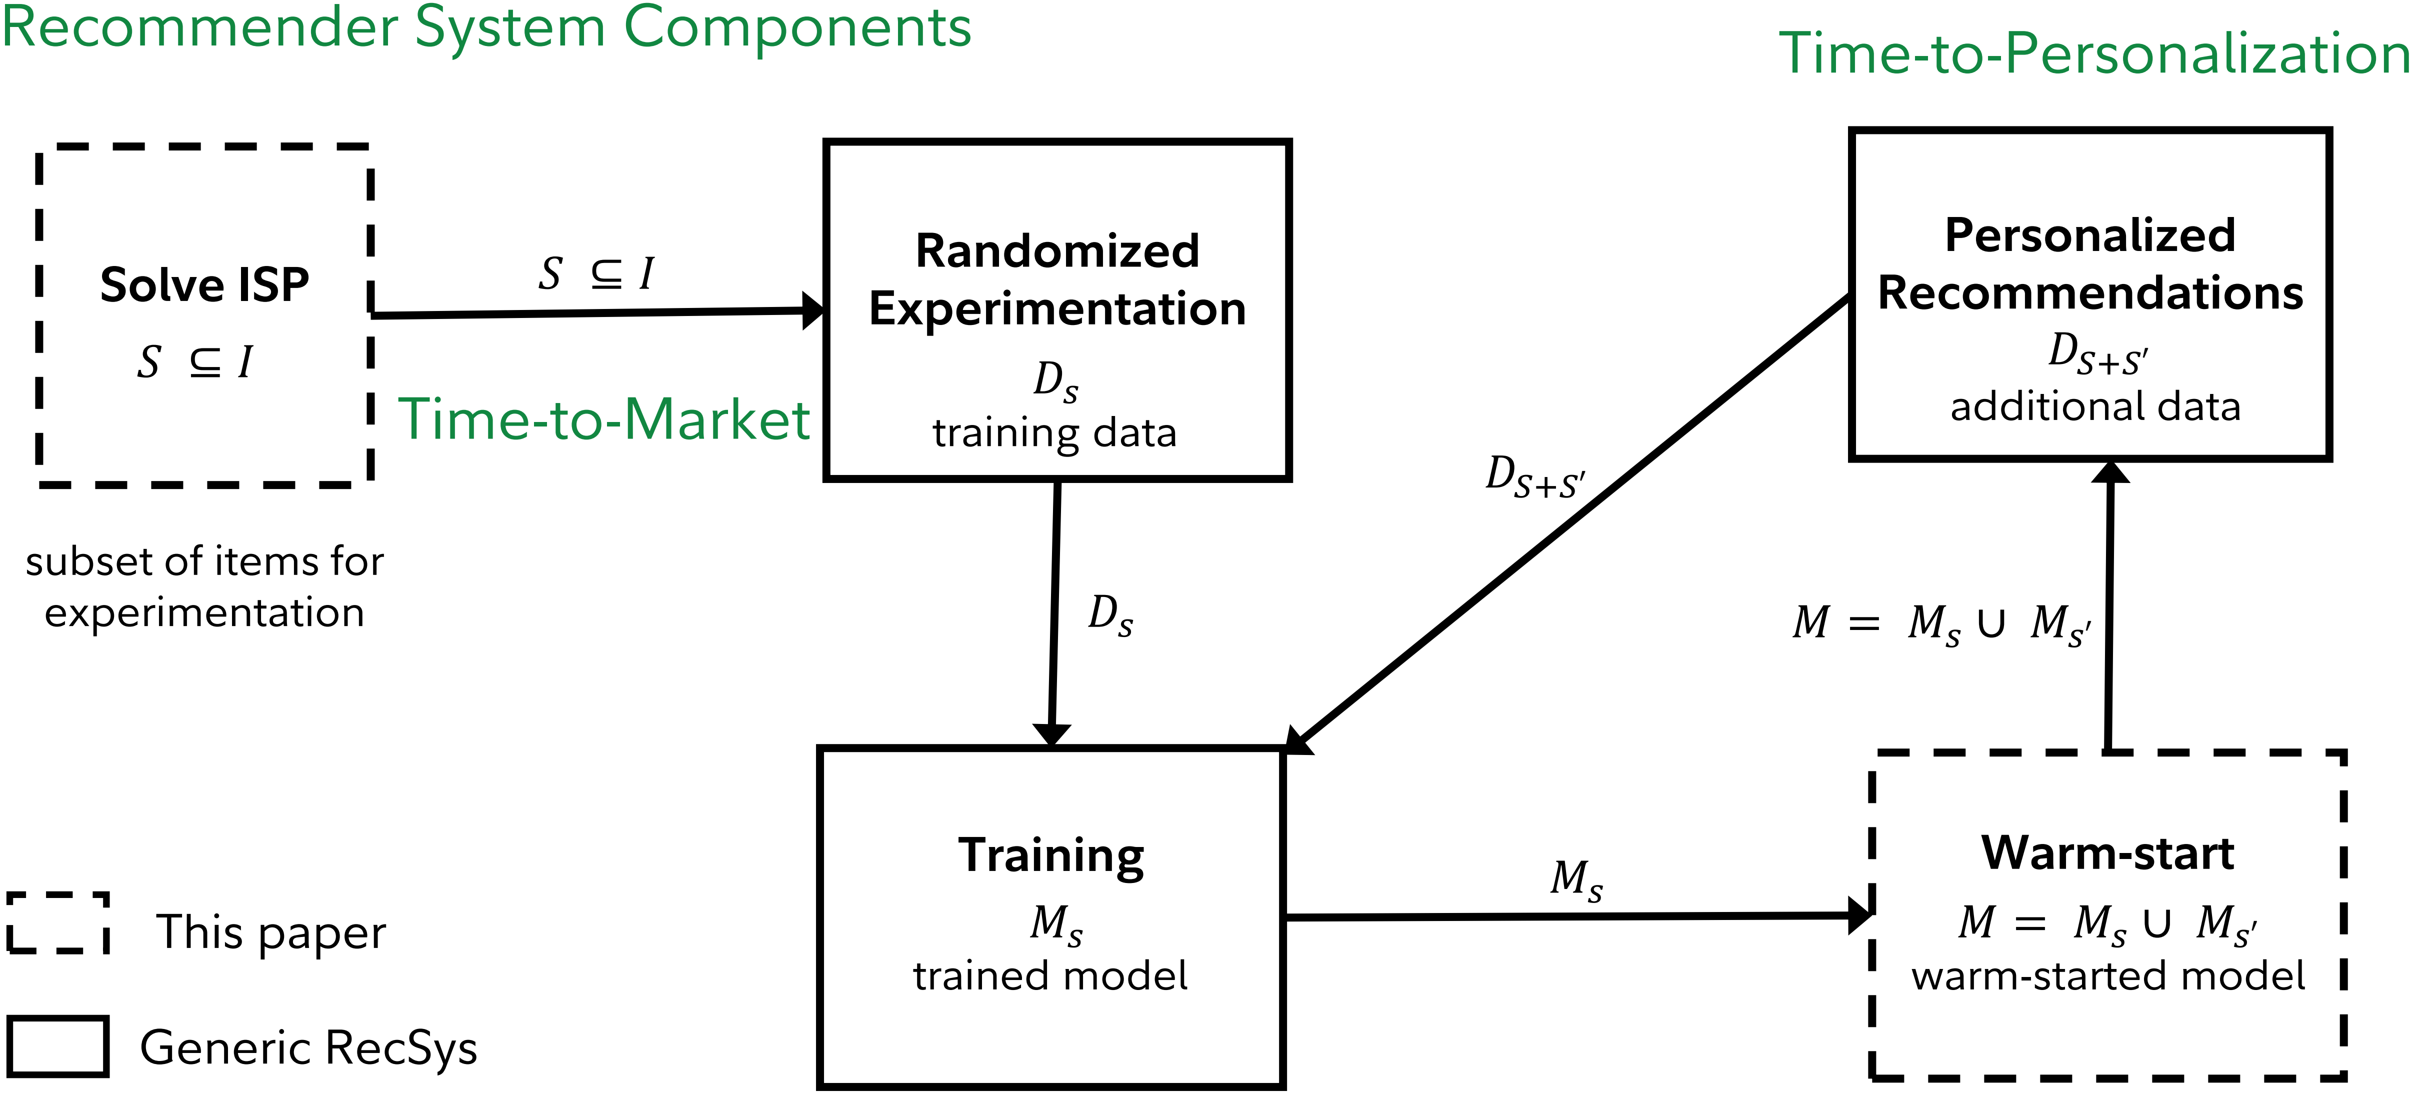

Mab2Rec

https://github.com/fidelity/mab2rec

Implications of Item Selection

Time-to-Market vs. Time-to-Personalization

Hypothetical Scenario

- 3M visits/week
- 1% CTR
- Uniform impressions

Implications

- Clicks per item
- Model training
- Feature space

Average Clicks per Item

Number of items	Weeks: 2	4	6	8
15	4,000	**8,000**	**12,000**	**16,000**
20	3,000	**6,000**	**9,000**	**12,000**
25	2,400	4,800	**7,200**	**9,600**
30	2,000	4,000	**6,000**	**8,000**
35	1,714	3,429	**5,143**	**6,857**
40	1,500	3,000	4,500	**6,000**
45	1,333	2,667	4,000	**5,333**
50	1,200	2,400	3,600	4,800

Multiple Trade-offs

Conflicting criteria for item selection

- Number of items vs. Experimentation Time
- Item Diversity vs. Learning Objectives
- Item Mix vs. Coverage Outcomes
- Item Onboarding (creation, review, publication, maintenance)

Solving the ISP

Cover formulations, Multi-objective optimization framework

Warm-start procedure

Solving the ISP

Set Covering Refresher

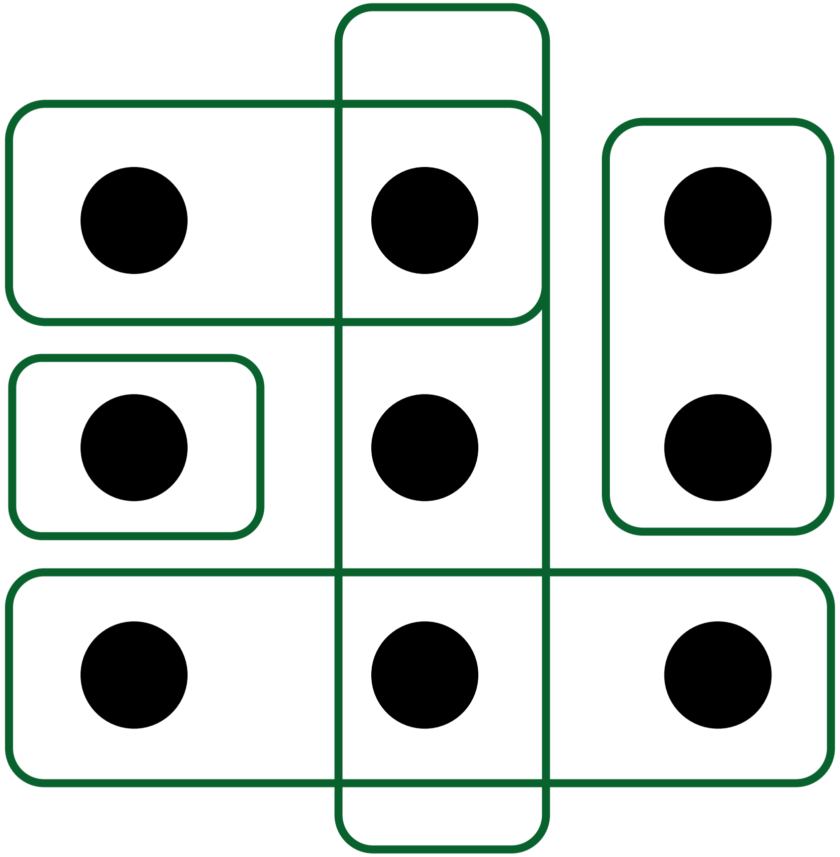

[1] Beasley, J.E.: An algorithm for set covering problem. European Journal of Operational Research 31(1), 85–93 (1987)

Solving the ISP

Multi-Objective Optimization

1	**Minimize Subset Size**	Use standard set covering [1] formulation to select subset of items that cover all predefined labels
2	**Maximize Diversity**	Reformulate unicost selection to yield minimum subset of items that are most spread in embedding space $E(I)$ while covering all labels
3	**Bounded Subset Size**	Constrain number of selected items in #2 while maximizing the number of labels covered

[1] S. Kadioglu, B. Kleynhans, X. Wang, Active learning meets optimized item selection (IJCAI'21)
[2] B. Kleynhans, X. Wang, S. Kadioglu, Optimized item selection to boost exploration for recommender systems (CPAIOR'21)

Solving the ISP

#1 Minimizing the Subset Size

Standard covering formulation to select a subset of items that cover all predefined labels

$\boldsymbol{P_{unicost}}$

$$
\begin{aligned}
& min \sum_{i}^{I} c_i x_i \\
& \sum_{i \in I} M_{l,i} x_i \geq 1 \qquad \forall l \in L_c, \forall c \in C \\
& x_i \in \{0,1\}, c_i = 1 \qquad \forall i \in I
\end{aligned}
$$

Assume $unicost\_selection \subseteq I$ is the solution to $P_{unicost}$ where
$\mathrm{k} = |unicost\_selection|$ is the number of selected items

Solving the ISP

#2 Maximizing Diversity

Given k from the solution of $P_{unicost}$, cluster the embedding space of items $\mathrm{E(I)}$ into k clusters and let K denote the cluster centers

Reformulate $P_{unicost}$ by changing its cost structure such that the inclusion of item i incurs cost, c_i based on the distance to its closest cluster

$\boldsymbol{P_{diverse}}$

$$\boxed{c_i = \min \quad distance(i, k) \quad k \in K \quad \forall i \in I}$$

Solving the ISP

#3 Maximize Bounded Subset Size

Given a constant t such that t ≤ $|Pdiverse|$ select up to t items from $diverse\_selection$ such that coverage is maximized

$$\boldsymbol{P_{max\_cover@t}}$$

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

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

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

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

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

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

Solving the ISP

Bringing it Together

1

// First Level: Minimize the subset size
Formulate $P_{unicost}(I, M)$
$unicost\_selection \leftarrow$ **solve**$(P_{unicost})$

2

// Second Level: Maximize diversity
$k \leftarrow |unicost\_selection|$
$K \leftarrow$ **cluster**$(E(I), num\_clusters = k)$
Initialize $cost \leftarrow$ zeros$(|I|)$
for all item $\in I$ **do**
  $cost_{item} \leftarrow$ **min**(**distance**$(item, centroids \in K)$)
end for
Formulate $P_{diverse}(I, M, cost, unicost\_selection)$
$diverse\_selection \leftarrow$ **solve**$(P_{diverse})$

3

// Third Level: Maximize bounded coverage
$t \leftarrow |diversity\_selection|$
Formulate $P_{max\_cover@t}(diverse\_selection, M, t)$
$S = max\_coverage \leftarrow$ **solve**$(P_{max\_cover@t})$

Warm-start Procedure

Exploiting the Exploration

Exploration with ISP yields the training data D_S which is used to build model M_S

Warm-start items $s' \in S': I \backslash S$ to build $M_{S'}$ sharing knowledge from M_S

Calculate item-to-item similarities using item embedding $E(I)$

Find a distance threshold, w, based on similarities of all items in S

Given pairwise distances, we find the closest item $s \in S$ for each untrained item s' such that $distance(s, s') \leq w$

Leverage the training data D_S or trained parameters of model M_S to warm-start s' [1, 2]

[1] Caruana, R., Niculescu-Mizil, A., Crew, G., Ksikes, A.: Ensemble selection from libraries of models. ICML 2004
[2] Caruana, R., Munson, A., Niculescu-Mizil, A.: Getting the most out of ensemble selection. ICDM 2006

Numerical Results

Effectiveness of the ISP Solution & Warm-start

Experiments

Research Questions

1. What is the **minimum number of items** required to cover all labels?
2. How much **speed-up is enabled** when ISP is used to collect response data?
3. How effective is the **warm-start procedure**?
4. How **sensitive is the ISP** to the choice of latent embedding space of items?

Experiments

Data & MIP Solver

Two well-known datasets:

- Goodreads Book Reviews with 11,123 books (items)
- MovieLens (ml-25m) Movie Recommendations with 62,423 movies (items)

Randomly selected subsets with 1,000 and 10,000 items

Python-MIP with COIN-OR CBC Solver

Dataset	# Items	Categories	# Labels
Goodreads [1]	1,000 10,000	{Genre, Publisher, Genre × Publisher}	574 1,322
MovieLens [2]	1,000 10,000	{Genre, Producer, Language, Genre × Language}	473 1,011

[1] Wan, M., McAuley, J.J.: Item recommendation on monotonic behavior chains. In: Pera, S., Ekstrand, M.D., Amatriain, X., O'Donovan, J. (eds.)
[2] Harper, F., Konstan, J.: The movielens datasets: History and context.

Experiments

Comparisons

Random+

Uniform random selection as a simple baseline. The Random method uses the subset size k from the solution for $P_{unicost}$

Greedy

The classical greedy [1] heuristic for set covering that adds items iteratively, whereby at each step, the item with the best $\frac{cost}{coverage}$ ratio is selected

Kmeans+

Clusters the latent space $E(I)$ into k centers and then selects items closest to the centroids, independent of item labels

[1] Vazirani, V.V.: Approximation algorithms. Springer (2001)

Experiments

Q1: What is the minimum number of items required to cover all labels?

Dataset	# Items	# Labels	|P_unicost|	% Reduction	Greedy # Items	Random Coverage	Greedy Coverage	KMeans Coverage
Goodreads	1,000	574	374	63%	374	57%	100%	58%
	10,000	1,322	1,080	89%	1,080	46%	100%	45%
MovieLens	1,000	473	243	76%	249	46%	100%	43%
	10,000	1,011	523	95%	703	29%	100%	31%

Summary

- $P_{unicost}$ yields substantial reduction in number of items and still cover all labels
- Coverage from $Random$ and $KMeans$ methods are markedly lower for same number of items
- $Greedy$ did not always yield the optimal solution

Experiments

Q1: What is the minimum number of items required to cover all labels?

Dataset	# Items	# Labels	|P_unicost|	% Reduction	Greedy # Items	Random Coverage	Greedy Coverage	KMeans Coverage
Goodreads	1,000	574	374	63%	374	57%	100%	58%
	10,000	1,322	1,080	89%	1,080	46%	100%	45%
MovieLens	1,000	473	243	76%	249	46%	100%	43%
	10,000	1,011	523	95%	703	29%	100%	31%

Summary

- $P_{unicost}$ yields substantial reduction in number of items and still cover all labels
- Coverage from $Random$ and $KMeans$ methods are markedly lower for same number of items
- $Greedy$ did not always yield the optimal solution

Experiments

Q2: How much speed-up is enabled in exploration phase when using ISP?

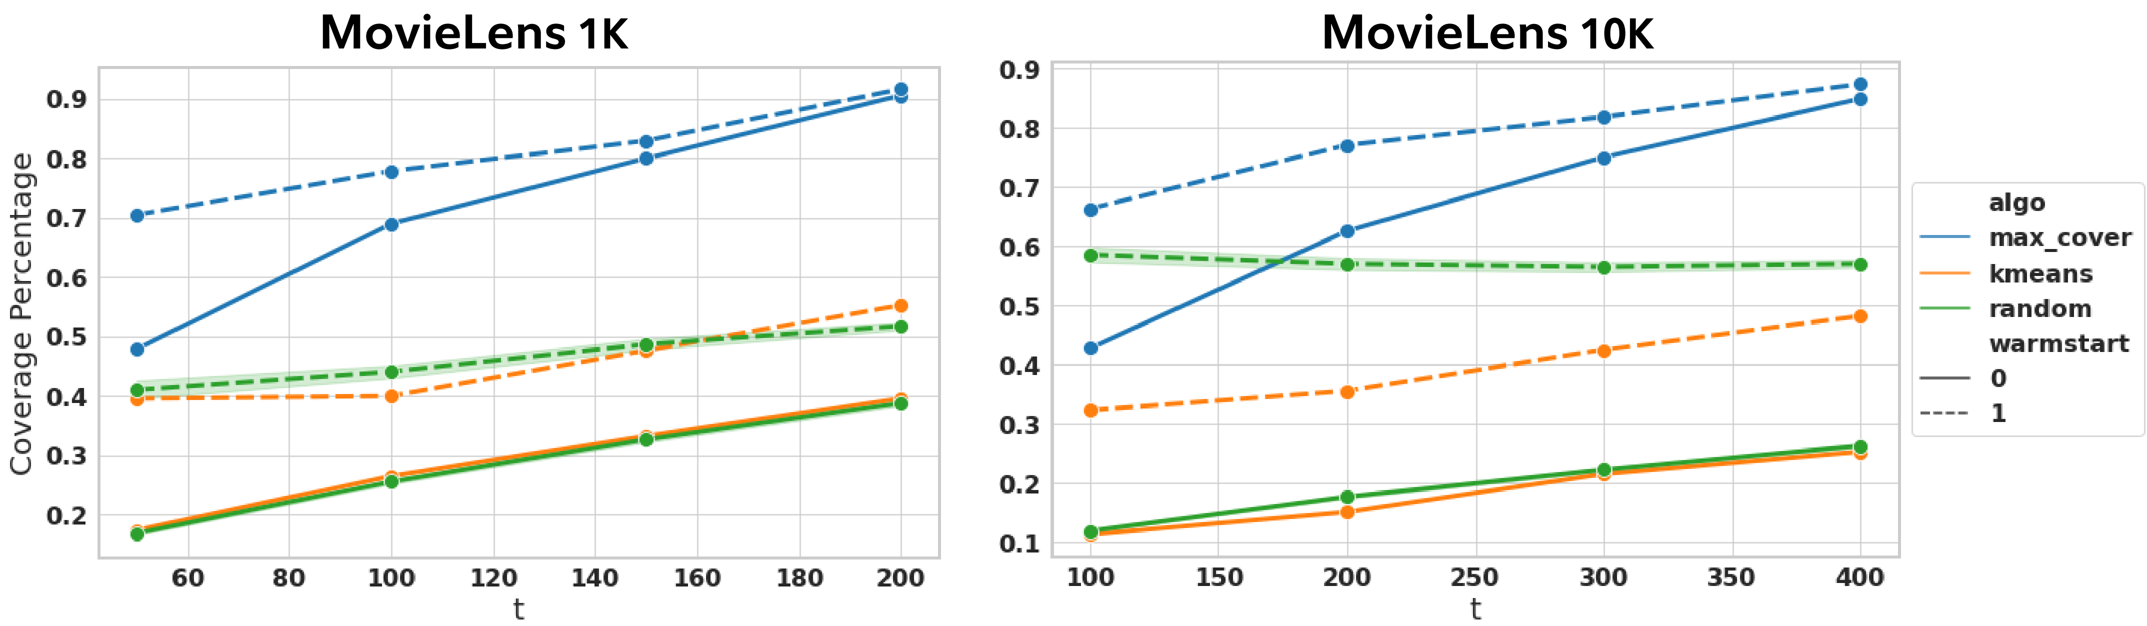

Summary

- For a given coverage level, $P_{max\_cover@t}$ requires much less items, hence learning time
- For a given threshold, coverage percentage higher for $P_{max\_cover@t}$ compared to other methods
- After warm-start (dashed lines), coverage for $P_{max\_cover@t}$ continues to rank highest

Experiments

Q2: How much speed-up is enabled in exploration phase when using ISP?

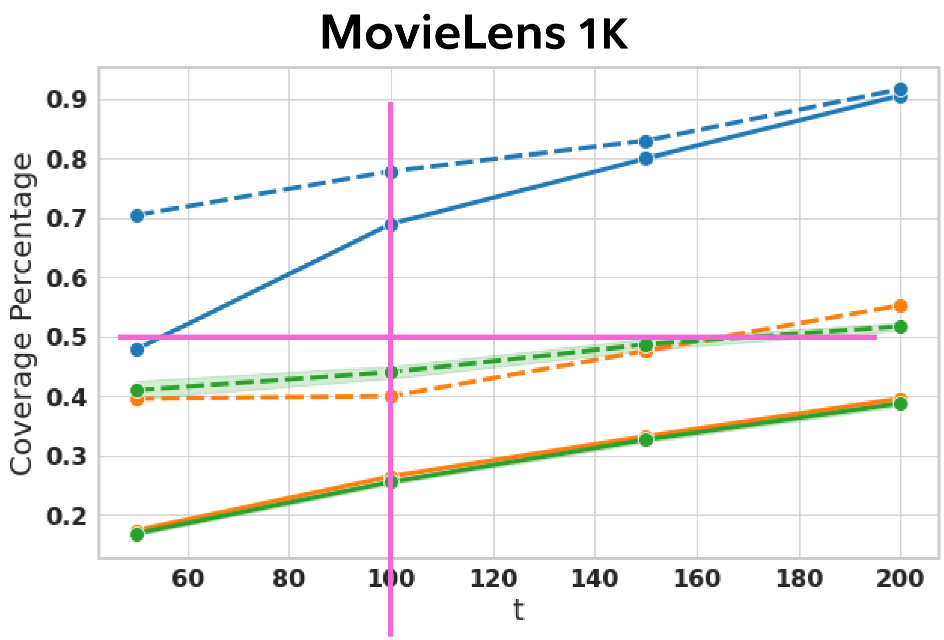

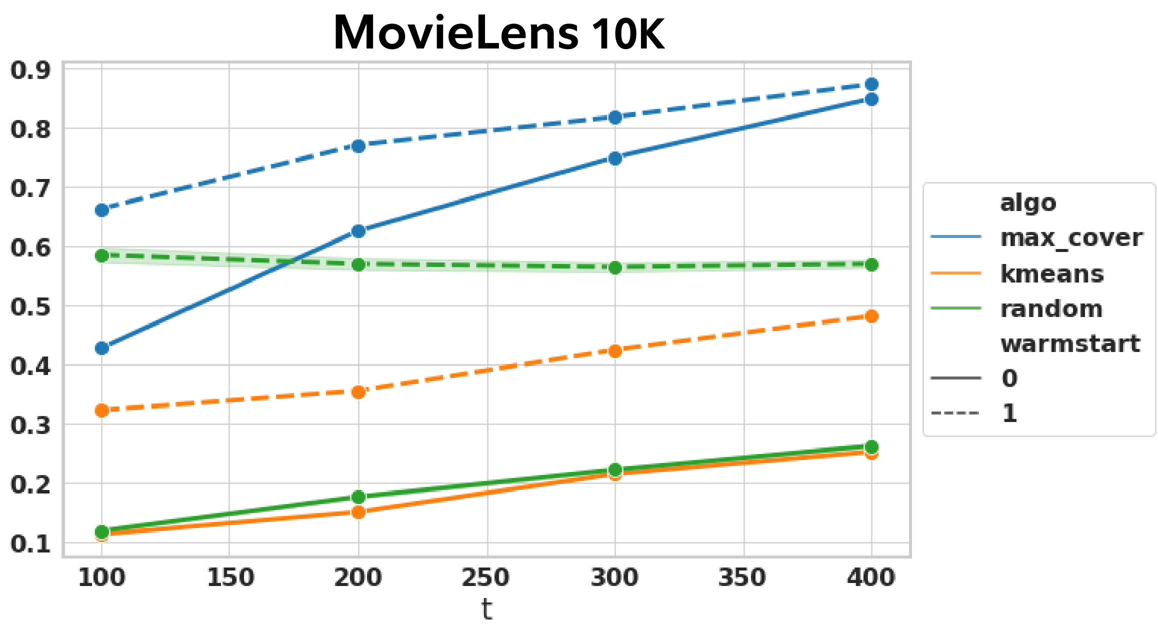

Summary

- For a given coverage level, $P_{max\_cover@t}$ requires much less items, hence learning time
- For a given threshold, coverage percentage higher for $P_{max\_cover@t}$ compared to other methods
- After warm-start (dashed lines), coverage for $P_{max\_cover@t}$ continues to rank highest

Experiments

Q3: How effective is the warm-start procedure as a function of distance quantile?

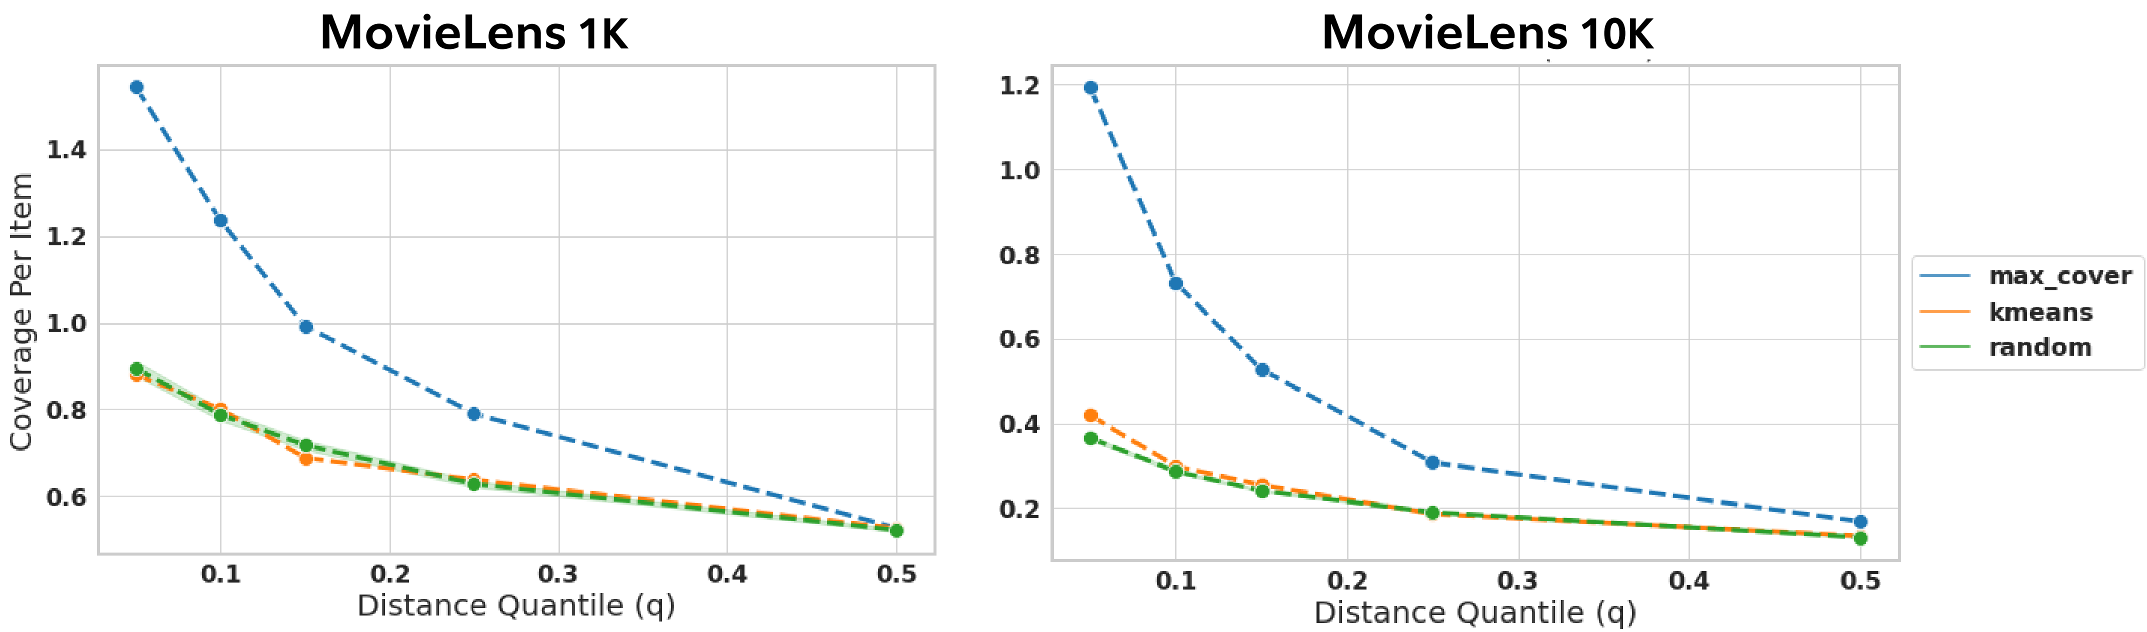

Summary

- As the distance quantile, q, is increased coverage per item decreases for all methods
- Consistent with the coverage analysis, $P_{max\_cover@t}$ is the most effective approach in terms of the number of labels covered per item, significantly better than $Random$ and $KMeans$, especially for the top (semi-) decile, i.e., $q \leq 0.1$

Experiments

Q4: How sensitive is the ISP to the choice of latent embedding space of items?

Goodreads: Coverage per item for $P_{max\_cover@100}$

Embedding [1]	1K	10K
TFIDF [2]	1.2	0.4
Word2Vec [3]	1.4	0.7
GloVe [4]	1.4	0.6
Byte-Pair [5]	1.3	0.6

Summary

- Similar unit coverage for different embedding methods
- More complex methods appear to provide better unit coverage compared to TFIDF

[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 201

Optimized Content Selection

2D embeddings of contents

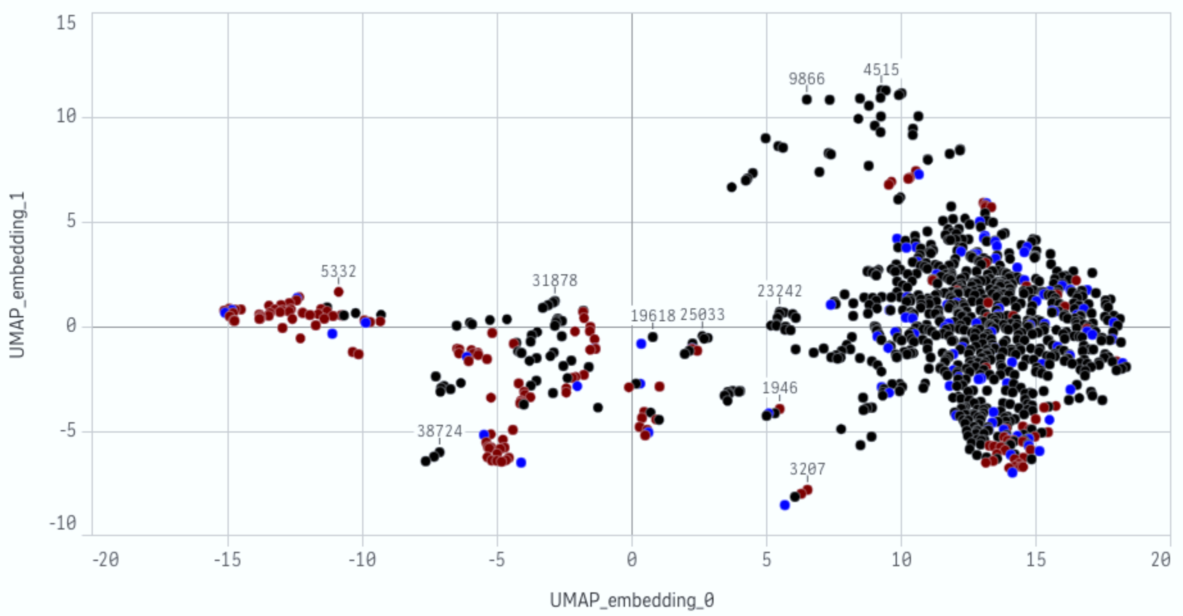

Content data

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

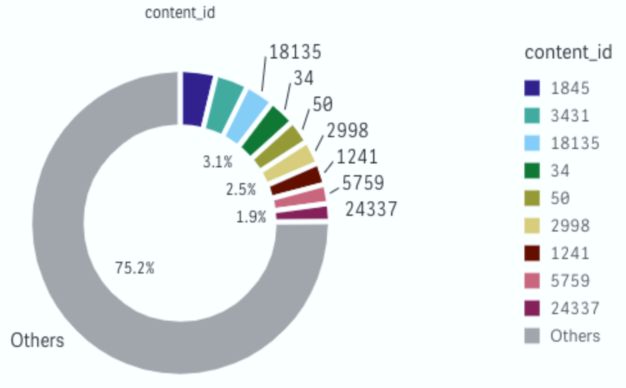

Number of Ratings v.s. Year Published

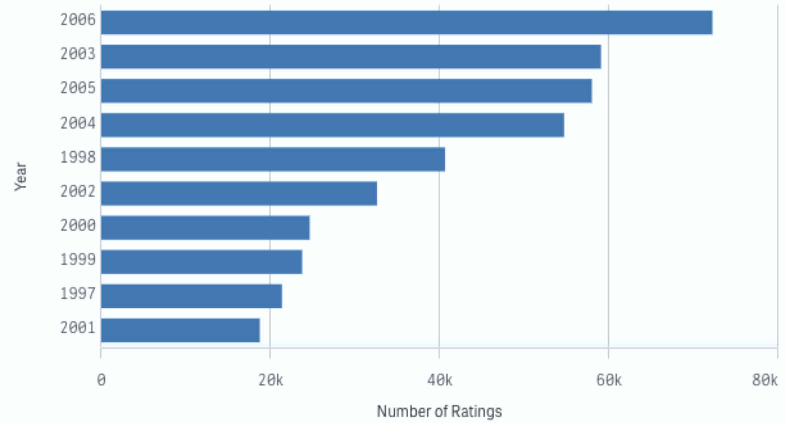

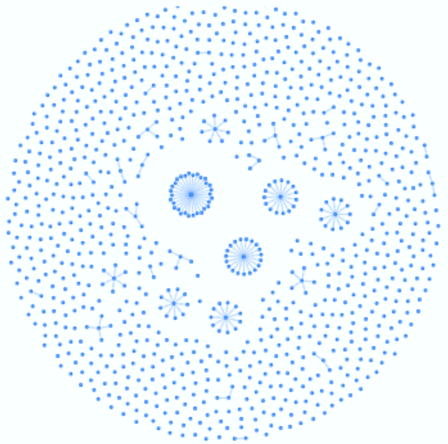

Special thanks to my collaborators!

- [TMLR'22] Non-deterministic behavior of Thompson sampling
- [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
- 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