

RevMax: Revenue-Maximizing Recommendation System Competition

EAAI-26 Model AI Assignment Track

Fang Sun, Paul Zhang, Pranav Subbaraman, Yizhou Sun

UCLA Computer Science Dept.

fts@cs.ucla.edu

Motivation: Beyond Accuracy

Traditional Recommender System assignments often focus on:

- Static datasets (MovieLens, Netflix Prize).
- Accuracy metrics (RMSE, Precision).
- Offline evaluation only.

RevMax shifts the focus to Real-World Dynamics:

- **Business Metric:** Revenue maximization (Price \times Purchase Probability).
- **Dynamic Environment:** Multi-iteration simulation where models *must* adapt.
- **Exploration vs. Exploitation:** Balancing immediate profit with learning user preferences.

The RevMax Platform

Built on the **Sim4Rec** simulation framework.

1. **Simulation Loop:**

- Recommender proposes Top-K items.
- Simulated Users respond (purchase/ignore).
- System calculates Revenue.
- Recommender receives feedback and updates model.

2. **Key Challenge:**

- Optimizing for **Cumulative Discounted Revenue** over time.
- Handling "Cold Start" for new users/items introduced during simulation.

Checkpoint 1: Content-Based Filtering

"Users like items with similar properties."

- **Focus:**
 - **Feature Engineering:** Extracting value from static user/item attributes (price, category, demographics).
 - **Cold Start Problem:** Handling new items or users with zero interaction history—a critical real-world challenge.
- **Technologies:** Scikit-learn, Random Forests, Gradient Boosting (XGBoost/LightGBM).
- **Learning Goal:** Mastering tabular data processing and traditional supervised learning pipelines.

Checkpoint 2: Sequence-Based Modeling

"What did the user interact with *just now* matters most."

- **Focus:**
 - **Temporal Dynamics:** Capturing evolving user interests and short-term intent.
 - **Sequential Dependencies:** Modeling the order of actions (e.g., *view phone* → *view case* → *purchase*).
- **Technologies:** PyTorch, RNNs/LSTMs, Transformers (SASRec/BERT4Rec).
- **Learning Goal:** Implementing deep learning architectures for sequential data.

Checkpoint 3: Graph-Based Approaches

"Users connected to similar items share latent preferences."

- **Focus:**
 - **Bipartite Graphs:** Modeling the explicit User-Item interaction network.
 - **Link Prediction:** Framing recommendation as predicting missing edges in the graph.
 - **High-Order Connectivity:** capturing collaborative signals beyond direct neighbors.
- **Technologies:** Graph Neural Networks (GCN, GraphSAGE), PyTorch Geometric.
- **Learning Goal:** Understanding geometric deep learning and structural data representations.

Technical Stack & Student Experience

Designed for **CS145: Introduction to Data Mining** at UCLA.

- **Language:** Python 3.10+
- **Infrastructure:**
 - **PySpark:** Exposure to distributed computing concepts.
 - **PyTorch:** Deep learning implementation.
 - **uv:** Modern, fast Python package management.
- **Workflow:**
 - Students implement `fit()` and `predict()` methods.
 - Provided `Analyzer` tool visualizes learning curves and revenue accumulation.

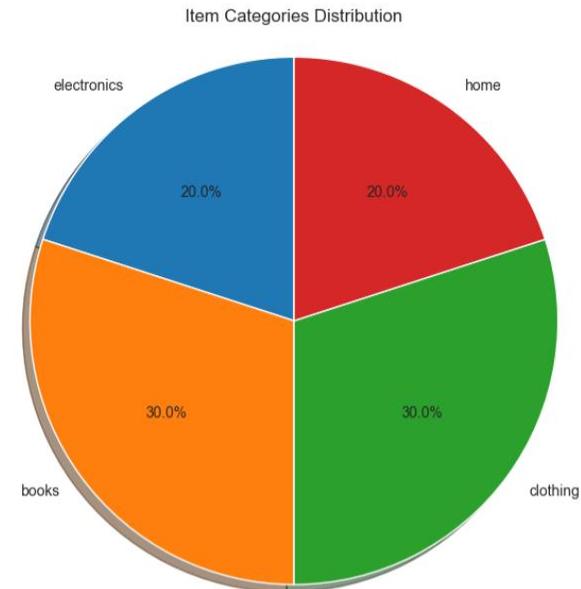
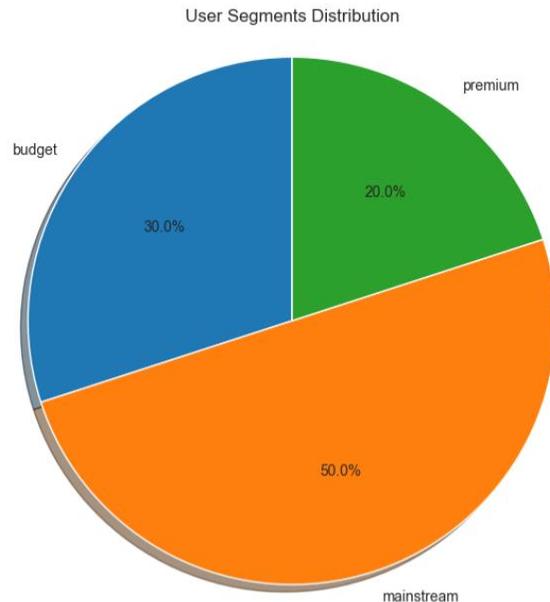
Data Analysis: Understanding the Simulation

Students begin by exploring the generated synthetic data.

- **User Segments:** Distributions of user types (e.g., price-sensitive vs. quality-focused).
- **Item Categories:** Product hierarchies and price distributions.
- **Interaction Patterns:** Heatmaps of user-item engagement.

Generated Data: User & Item Distributions

Left: User Segments / Right: Item Categories



- Visualizing these distributions helps students identify bias and class imbalance.
- Encourages segment-specific strategies.

Generated Data: Price & Interactions

Left: Price Distribution / Right: Interaction Heatmap



- **Price:** Essential for optimizing Revenue (not just CTR).
- **Heatmap:** Reveals which user segments prefer which categories.

Baseline Models To Start With

The platform includes standard baselines for comparison:

1. **Random:** Stochastic selection (lower bound).
2. **Popularity:** Recommends most purchased items.
3. **SVM/Linear:** Simple supervised learning.

Evaluation & The Leaderboard

Competition drives engagement.

- **Primary Metric:** Discounted Revenue (Price weighted by rank).
- **Secondary Metrics:** NDCG, Precision@K, Hit Rate.
- **Robustness:**
 - **Hidden Test Environments:** Submissions evaluated on unseen random seeds.
 - Prevents overfitting to the training simulation parameters.

Live Leaderboard

- Real-time feedback for students.
- Gamified experience encourages iterative improvement.
- Supports multiple baselines (Random, Popularity) for reference.

Submit Your Recommender

Team Name *

Choose a unique team name for the leaderboard

Email Address *

We'll use this to contact you if needed

Submission File *

Choose File No file chosen

Upload your `submission.py` file (max 16MB)

[Submit & Evaluate](#)

Submission Guidelines

- Your file must be named exactly `submission.py`
- Implement your recommender as `MyRecommender` class
- Follow the template structure provided in the course materials
- Rankings are based on `test_discounted_revenue` metric
- Evaluation may take several minutes to complete
- Check the leaderboard for your results

Top Performers [View Full Leaderboard](#)

Rank	Team	Score
1	Eddy	5619.8770
2	66	2831.2164
3	DRAV	2831.2164
4	66	2831.2164
5	66	2831.2164

Competition Stats

50	50	8
Total Submissions	Successful	Unique Teams

How It Works

- 1. Upload**
Submit your `submission.py` file with your custom MyRecommender implementation.
- 2. Evaluate**
Our server runs your recommender against the standardized test dataset using train-test split.
- 3. Rank**
Results are ranked by `test_discounted_revenue` - the key performance metric.
- 4. Compete**
See your position on the leaderboard and compete with your classmates!

Classroom Integration

Flexible & Modular:

- **Duration:** Can be run as a semester-long project or in short 2-week modules.
- **Difficulty:** Checkpoints can be assigned individually.
 - *Undergrad:* Focus on Content-based & basic Sequences.
 - *Grad:* Full Deep Learning & Graph implementations.

Provided Resources:

- Detailed Directives for each checkpoint.
- Baseline implementations (Starter code).
- Visualization tools for error analysis.
- Auto-grading scripts.

Learning Outcomes

By the end of RevMax, students can:

1. **Translate Business Goals to Code:** Optimizing for *revenue* instead of just *clicks*.
2. **Handle Dynamic Data:** Implementing online learning and frequent model updates.
3. **Implement Modern Architectures:** From Logistic Regression to Transformers and GNNs.
4. **Manage Trade-offs:** Latency vs. Accuracy, Exploration vs. Exploitation.

Thank You!

RevMax is available for the EAAI Community.

- **Website:** <https://modelai.gettysburg.edu/2026/revmax/>
- **Code:** <https://github.com/FrancoTSolis/RevMax>
- **Contact:** fts@cs.ucla.edu

Ready to maximize revenue in your classroom?