

Recommender Systems || Lecture 10

Summer 2022

Dr. Karthik Mohan

UW, Seattle

August 2, 2022

Logistics

- ① Assignment 3 due next Wednesday

Logistics

- ① Assignment 3 due next Wednesday
- ② Project checkpoint due next Monday
- ③ Please pick a time slot for your team to discuss project proposal if you haven't yet!

↓
Some teams yet to discuss!

Logistics

- ① Assignment 3 due next Wednesday
- ② Project checkpoint due next Monday
- ③ Please pick a time slot for your team to discuss project proposal if you haven't yet!
- ④ Anything else?!

Today

① Types of Ranking Problems

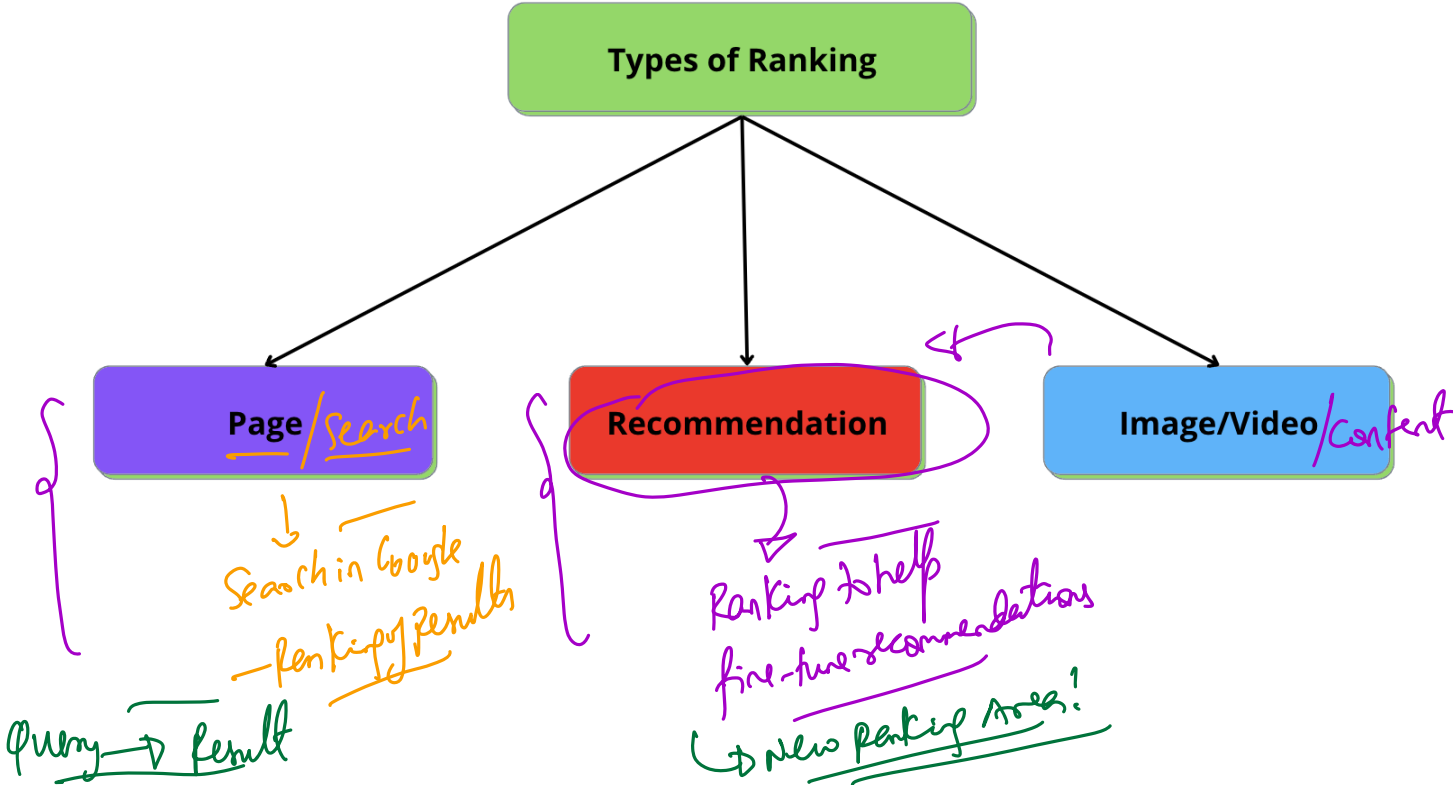
Today

- ① Types of Ranking Problems
- ② Recommendations vs Ranking

Today

- ① Types of Ranking Problems
- ② Recommendations vs Ranking
- ③ Real-world Examples of Recommendations Ranking

Ranking Problems



Page/Search Ranking

Google ml interview questions

Tools

About 131,000,000 results (0.62 seconds)

Ad · https://learn.interviewkickstart.com/ml_interview/preparation

ML Interview Preparation - 17000+ Problems | 1 Pattern

Machine Learning **Interview** Preparation. Learn The Power Patterns & Nail Any Coding Problem. Learn patterns to solve 17000+ tech **interview questions** & nail your next **interview**.

Let's start with some commonly asked machine learning interview questions and answers.

- What Are the Different Types of Machine Learning? ...
- What is Overfitting, and How Can You Avoid It? ...
- What is 'training Set' and 'test Set' in a Machine Learning Model? ...
- How Do You Handle Missing or Corrupted Data in a Dataset?

[More items...](#) • Jul 18, 2022

<https://www.simplilearn.com/tutorials/machine-learn...>

Top 45 Machine Learning Interview Questions Answered for ...

Machine Learning Interview Questions For Freshers

- Why was Machine Learning Introduced? ...
- What are Different Types of Machine Learning algorithms? ...
- What is Supervised Learning? ...
- What is Unsupervised Learning? ...
- What is 'Naive' in a Naive Bayes? ...
- What is PCA? ...
- Explain SVM Algorithm in Detail. ...
- What are Support Vectors in SVM?

[More items...](#) • Jun 23, 2022

<https://www.interviewbit.com/machine-learning-interv...>

Machine Learning Interview Questions (2022) - InterviewBit

About featured snippets • Feedback

<https://www.springboard.com/Blog/Data-Science...>

Page/Search Ranking



ml ops

Tools

All Images News Videos Shopping More

About 40,100,000 results (0.55 seconds)

Ad · <https://www.applica.ai/> :
Insights | MLOps - Real-World Application - applica.ai
Discover Applica's Deep Learning-based Automation Solution. Comprehend Even the Most Complex and Variable Business Documents. Get a Demo. Unmatched Precision. For All Complexities. All Document Types. Hyper Learning NLP.

Ad · <https://info.wallaroo.ai/mlops> :
Mlops Platform - Make Machine Learning Easy - wallaroo.ai
Learn How To Optimize Live Models Based On The Latest Data And Insights With No Downtime. Deploy, Run & Optimize Your Data For a Fraction Of The Cost. Learn How With Our eBook. Simple & Secure. Self-Service Toolkit. On-Premises. On Any Cloud.

<https://ml-ops.org> :
ML Ops: Machine Learning Operations
With Machine Learning Model Operationalization Management (MLOps), we want to provide an end-to-end machine learning development process to design, ...
[Crisp-ml\(q\)](#) · [ML Model Governance](#) · [Three Levels of ML Software](#) · [MLOps Principles](#)

People also ask :

- What are ML operations?
- What does an MLOps do?
- Is MLOps a data engineer?
- Is MLOps worth learning?

Feedback

<https://databricks.com> › Glossaries :
MLOps - Databricks
MLOps stands for Machine Learning Operations. MLOps is a core function of Machine Learning engineering, focused on streamlining the process of taking ...
[Take your ML projects from...](#) · [The Big Book of Machine...](#) · [MLOps Virtual Event](#)


<https://en.wikipedia.org> › wiki › MLOps :
MLOps - Wikipedia

MLOps :



MLOps or ML Ops is a set of practices that aims to deploy and maintain machine learning models in production reliably and efficiently. The word is a compound of "machine learning" and the continuous development practice of DevOps in the software field.
[Wikipedia](#)

Mlops Book



Introducing MLOps Practical MLOps:... Engineering MLOps:... Practical MLOps


Feedback



ads
1st
2nd
3rd



Search Ranking jobs on LinkedIn :-)



Machine learning ranking in United States
19 results

Set alert

 **Senior ML Researcher - App Store Search Ranking - AMP AI/ML**
Apple
Cupertino, CA
 You have a preferred skill badge
2 weeks ago

 **Machine Learning Engineer, Ranking, Augmented Reality**
Snap Inc.
Los Angeles, CA
2 benefits
 4 connections work here
5 days ago

 **Machine Learning Engineer, Ranking, Augmented Reality**
Snap Inc.
Palo Alto, CA
2 benefits
 4 connections work here
1 week ago · **23 applicants**

 **Applied Scientist, Ranking & Recommendation Systems**
Amazon
New York, NY
\$144K/yr - \$202K/yr (LinkedIn est.)
 121 connections work here

Senior ML Researcher - App Store Search R...
Apple · Cupertino, CA

[Apply](#) [Save](#) ...

Key Qualifications

Deep understanding of the search and information retrieval fundamentals including indexing, query understanding, retrieval and ranking

Experience working with search relevance and ranking algorithms.

Experience in training large-scale machine-learning models for online search, recommendation, ads ranking, or personalization etc.

Experience in deep learning

Technical leadership in machine learning.

Ability to work independently and in a team to research innovative solutions to challenging business problems.

Experience working with real-world data sets and building scalable models from big data.

Experience with Spark, Hadoop or other distributed frameworks

Effective verbal and written communication skills with non-technical and technical audiences.

Handwritten notes:
- IR (circled) above "information retrieval fundamentals"
- New (circled) with an arrow pointing to "machine-learning models"

Search Ranking jobs on LinkedIn :-)

Machine learning ranking in United States

19 results

Set alert

Snap Inc.

Machine Learning Engineer, Ranking, Augmented Reality

Snap Inc.

Palo Alto, CA

2 benefits



4 connections work here

1 week ago · 23 applicants



Applied Scientist, Ranking & Recommendation Systems

Amazon

New York, NY

\$144K/yr - \$202K/yr (LinkedIn est.)



121 connections work here

2 weeks ago



Staff Machine Learning Engineer - Home Ranking

Spotify

New York, NY



1 connection works here

2 weeks ago



Sr. Applied Scientist, Ranking & Recommendation Systems, Percolate Experience Optimization

Amazon

Seattle, WA

\$140K/yr - \$160K/yr (LinkedIn est.)



121 connections work here

Applied Machine Learning Manager, Search...

Etsy · Brooklyn, NY

Apply

Save

...

What's the role?

Etsy is looking for an engineering manager who can lead a team of Applied Scientists and ML Engineers in Search to drive the machine learning strategy for Etsy's Search Ranking team, with a focus on Learning-to-Rank personalization and deep learning approaches. You will lead a team with the mission to build customized machine learning applications to drive innovation and optimize business values for the Etsy marketplace in production. To learn more about Machine Learning and Engineering at Etsy, check out <https://dsml.etsy.com> and <https://codeascraft.com/>

For this role, we are considering candidates based in the United States who are either remote, flex, or office-based. Etsy offers different work modes to meet the variety of needs and preferences of our team. Learn more about our flexible work options and vaccination policy here.

What's this team like at Etsy?

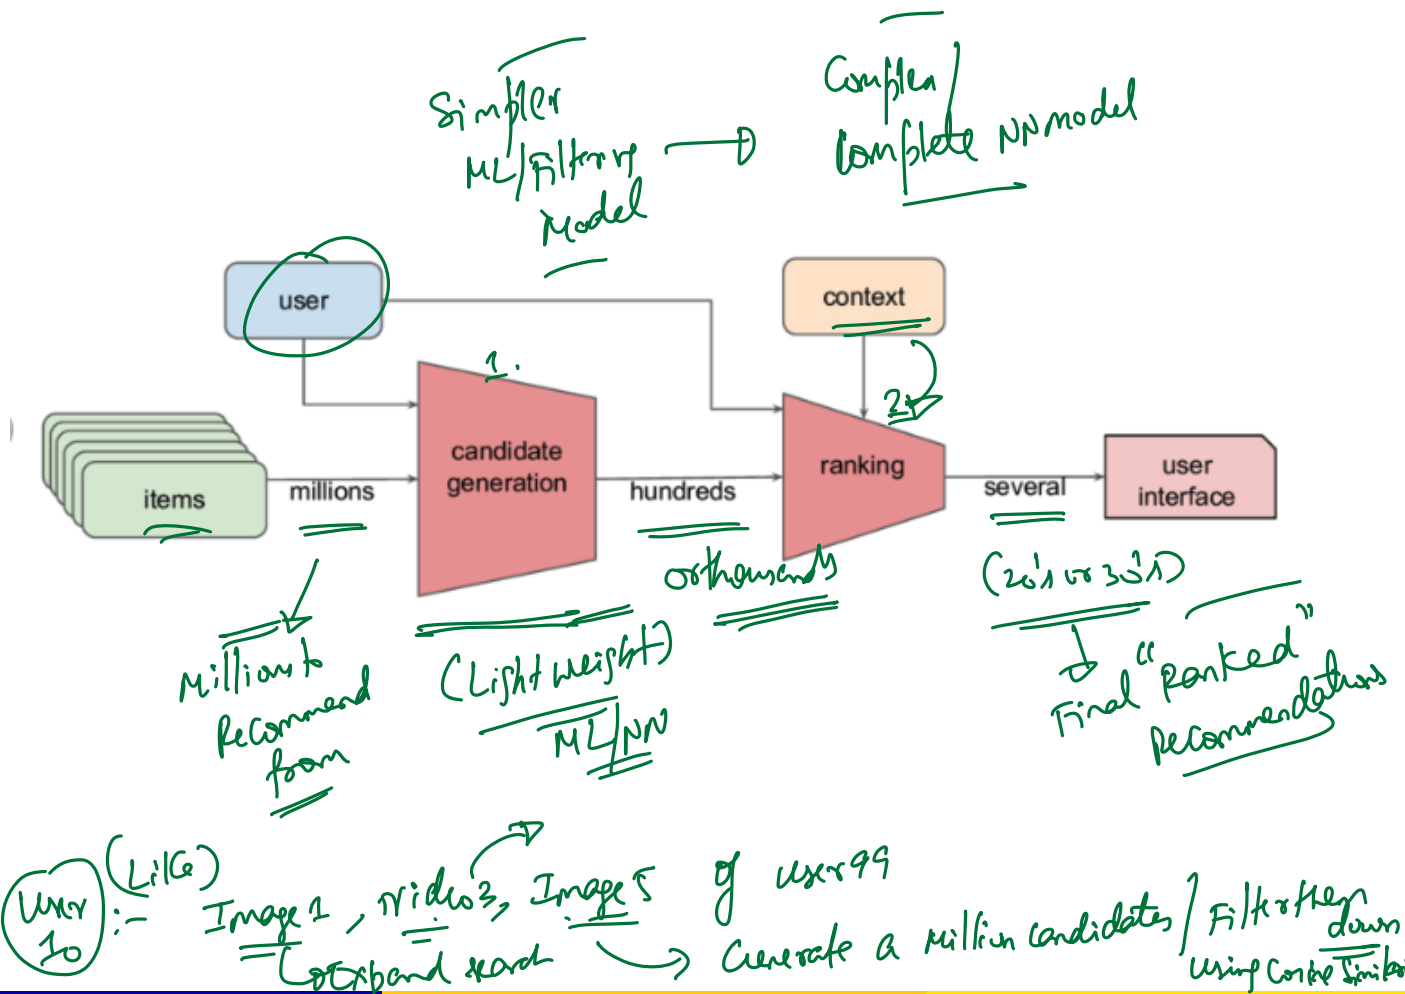
- Provide technical guidance and mentorship in search ranking, relevance and personalization algorithms.
- Communicate to executives and product/engineering leaders; drive the strategy of ML innovation.
- Lead technical discussions, be a domain expert and have strategic influence.
- Define the roadmap and supervise its execution.
- Cultivate a strong applied science and engineering culture, responsible for delivering business results.

What does the day-to-day look like?

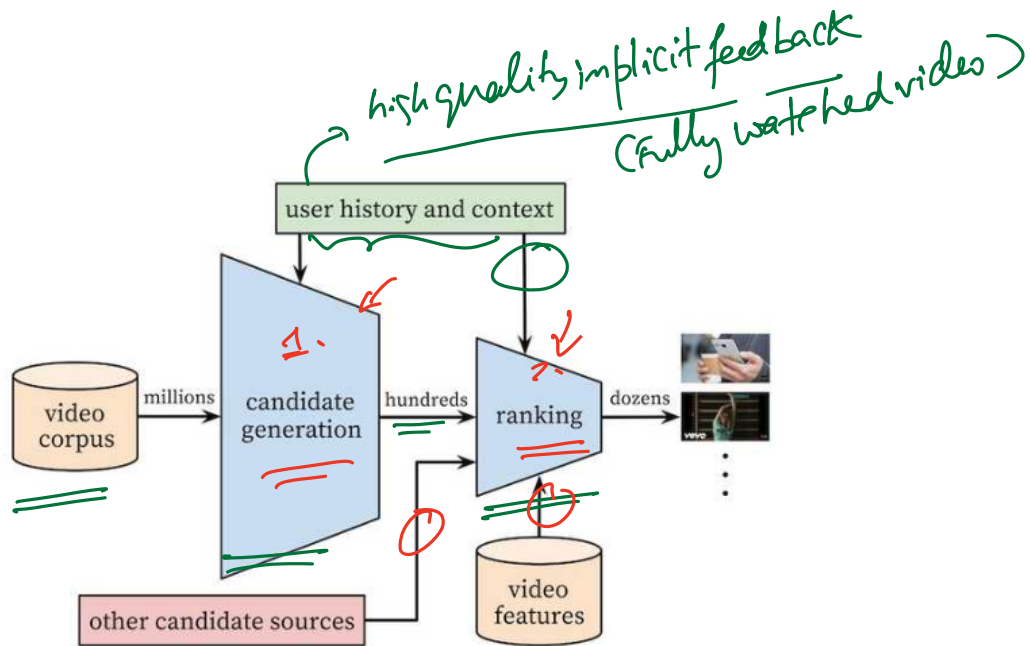
- Increase conversion rate and gross merchandise sales through personalization and ranking algorithms with quality signals, and improve personalized buyer experience and buyer confidence such that they would revisit or re-purchase again.
- Develop embeddings to collect salient signals of our search queries, users, and listings, including aspects like content, budget, preferences and buyer

Recommendations Ranking - Generic Template

2/



Recommendations Ranking (YouTube)



Real-time recommendations on YouTube ([source](#))

YouTube Recommendations

(reference) → 2016

Recommendations Ranking (YouTube)

Challenges to building a YouTube scale Recommender:

- 1 Scale
- 2 Freshness
- 3 Noise

§ YouTube Recommendations

Recommendations Ranking (YouTube)


Challenges to building a YouTube scale Recommender:

- 1 Scale
- 2 Freshness
- 3 Noise → *Noise is the implicit feedback*
- 4 1B parameters and 100 MM examples (this was in 2016). Anyone's guess on what the updated model looks like! *Divides (user-video)*

YouTube Recommendations

Candidate Generation (YouTube)

- 1 ML Modeling approach: Multi-class classification



YouTube Recommendations

Candidate Generation (YouTube)

- 1 **ML Modeling approach:** Multi-class classification
- 2 **Model:** NN model that uses ~~Like~~ a non-linear ~~NN~~ model based on user watch history data / Content data on the videos MF

YouTube Recommendations

Candidate Generation (YouTube)

- ① **ML Modeling approach:** Multi-class classification
- ② **Model:** NN model that uses Like a non-linear NN model based on user watch history data
- ③ **Positive example:** A fully video watch by a user counts as a positive example (target)

YouTube Recommendations

Candidate Generation (YouTube)

- ① **ML Modeling approach:** Multi-class classification
- ② **Model:** NN model that uses Like a non-linear NN model based on user watch history data
- ③ **Positive example:** A fully video watch by a user counts as a positive example (target)
- ④ **Negative examples:** Are randomly sampled (Importance) ^{sampling}

YouTube Recommendations

Candidate Generation (YouTube)

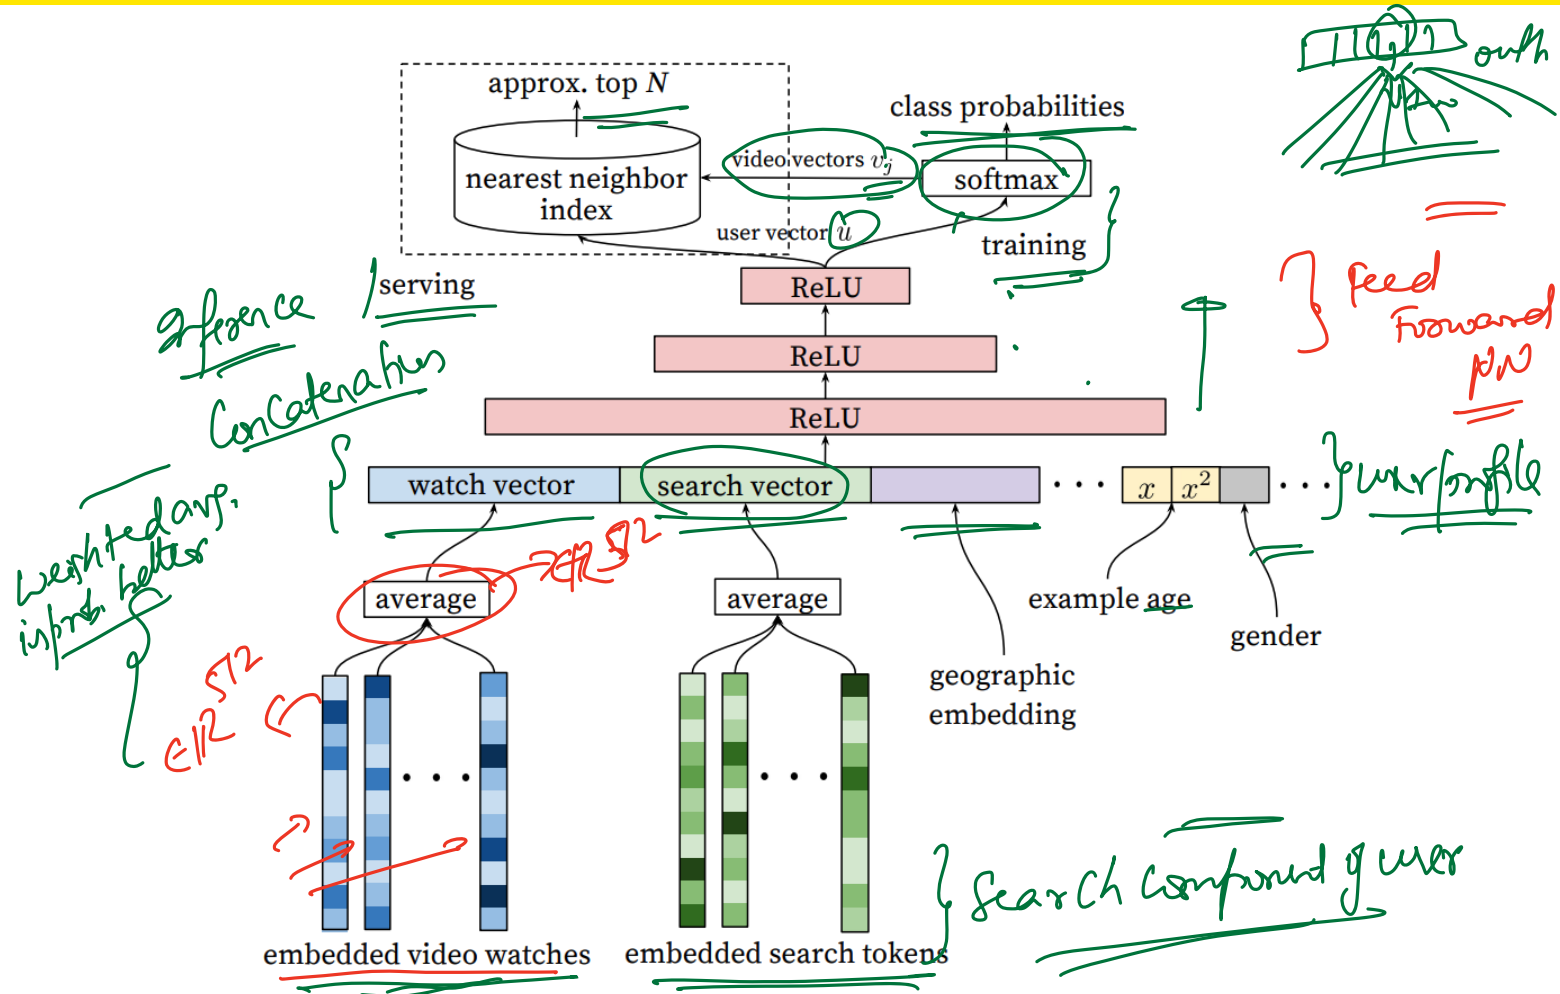
- 1 **ML Modeling approach:** Multi-class classification
- 2 **Model:** NN model that uses Like a non-linear NN model based on user watch history data
- 3 **Positive example:** A fully video watch by a user counts as a positive example (target)
- 4 **Negative examples:** Are randomly sampled
- 5 **Inference time:** Latency requirement in tens of milliseconds - So use approximate kNN based on user and video embeddings

YouTube Recommendations

reference (cos θ)



Candidate Generation (YouTube)



ICE #1

Candidate Generation

The Candidate Generation model at YouTube is an example of:

- 1 Content Based Filtering
- 2 Collaborative Filtering
- 3 Hybrid Model
- 4 None of the above

Poll

ICE #2


Candidate Generation

The loss function for candidate generation stage of the YouTube model would most likely be:

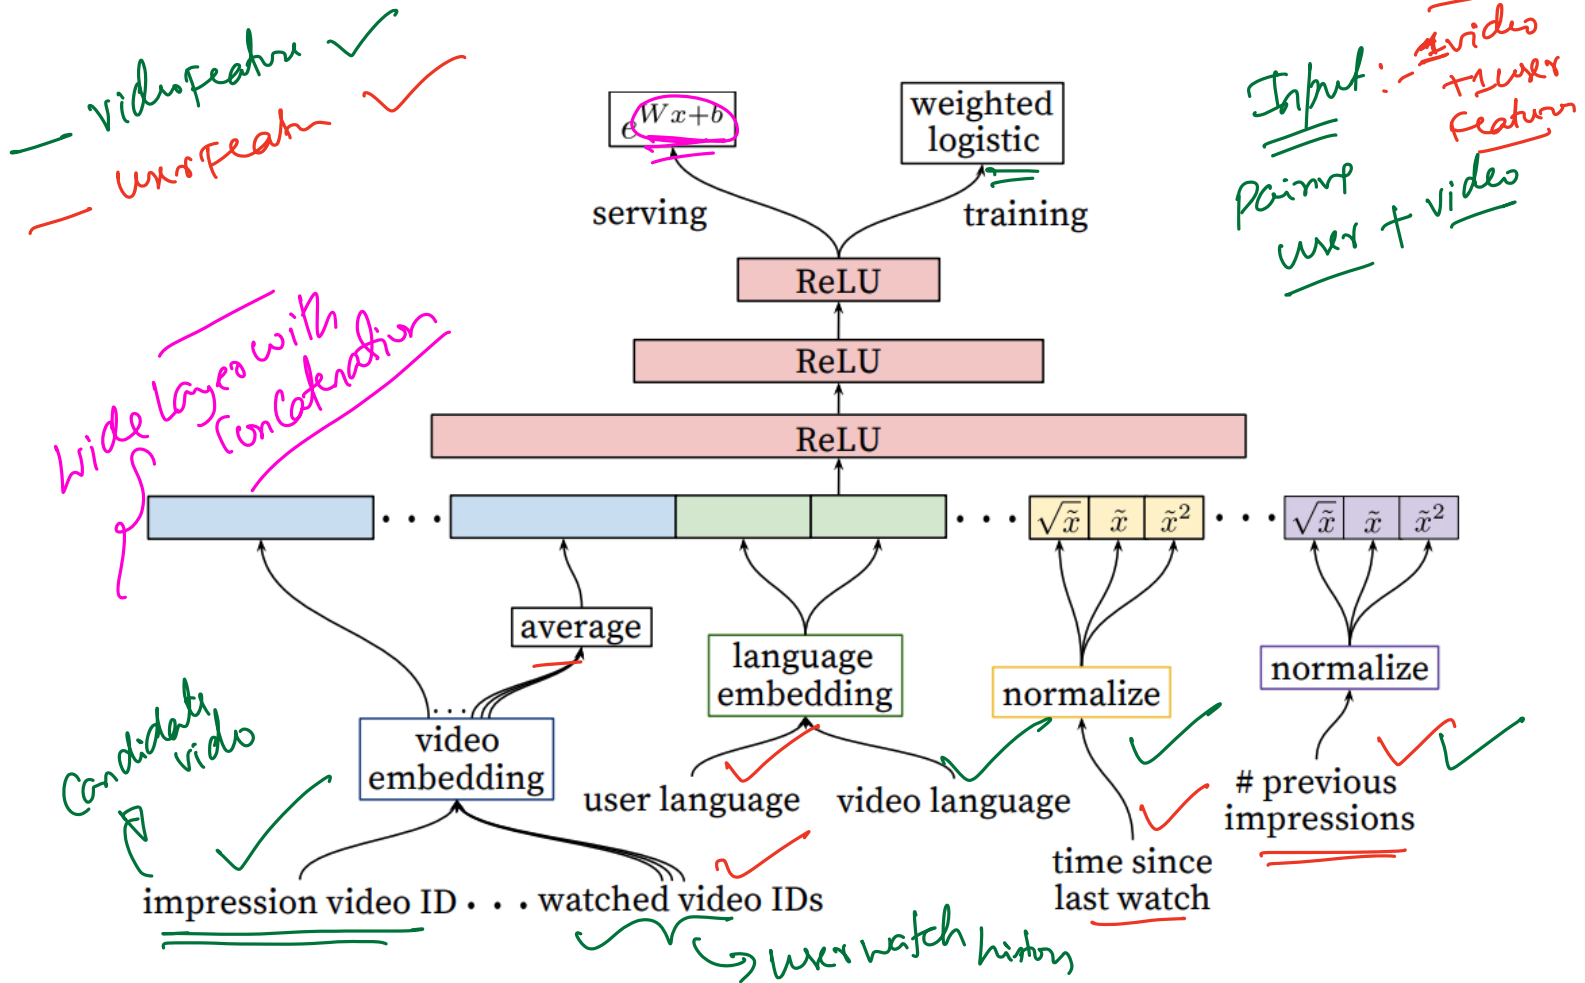
- 1 Cross-entropy loss
- 2 Logistic loss
- 3 Ranking loss
- 4 Quadratic loss

Poll

Candidate Generation (YouTube)



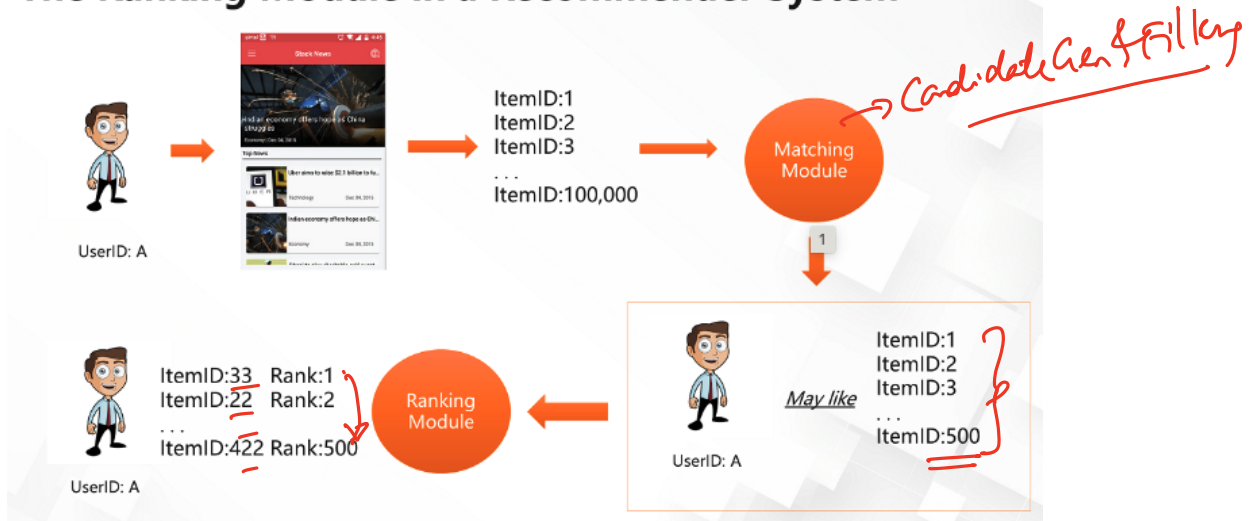
Ranking Model (YouTube)



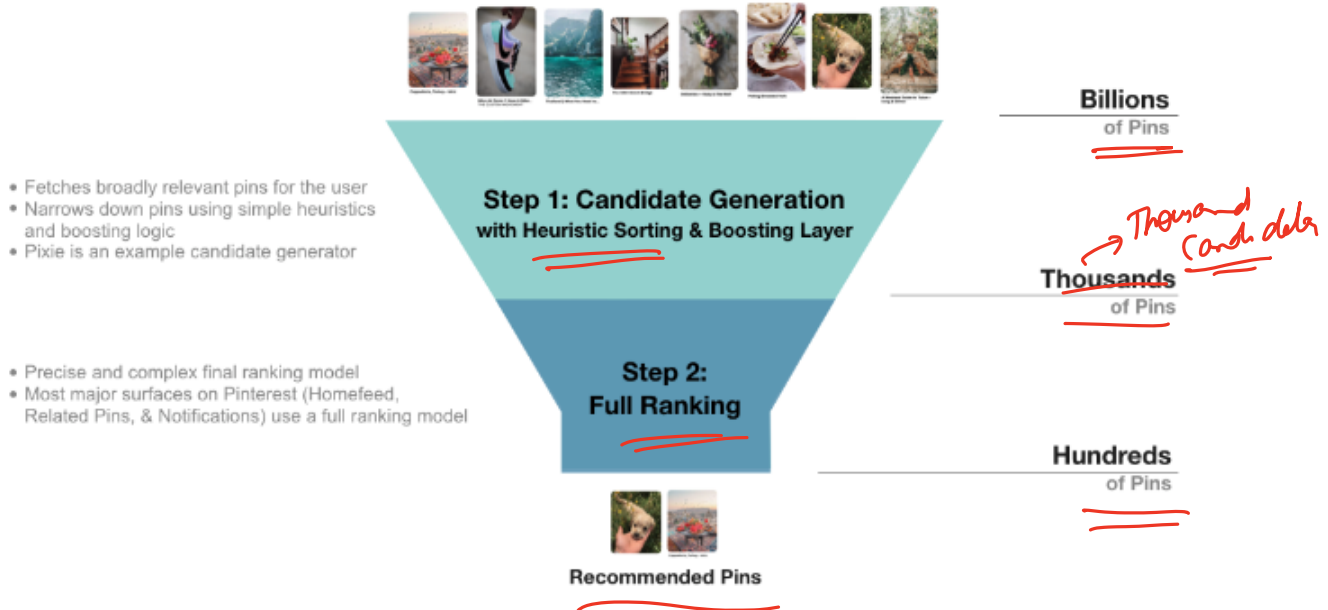
Recommendations Ranking (AliBaba)

The Ranking Module in a Recommender System

阿里云




Recommendations Ranking (Pinterest)





Recommendations Ranking on LinkedIn :-)

Recommendations ranking in United States
505 results Set alert





Applied Scientist, Ranking & Recommendation Systems

Amazon
New York, NY
\$144K/yr - \$202K/yr (LinkedIn est.)
 121 connections work here
2 weeks ago





Applied Scientist II, Consumer Payments

Amazon
Seattle, WA
\$140K/yr - \$160K/yr (LinkedIn est.)
 121 connections work here
2 weeks ago




Senior ML Researcher - App Store Search Ranking - AMP AI/ML

Apple
Cupertino, CA
 Premium tip: you'd be a top applicant
2 weeks ago





Senior Applied Scientist, Amazon Advertising


Amazon
Seattle, WA
 121 connections work here
1 week ago · **3 applicants**


Applied Scientist, Ranking & Recommendation Systems

Amazon · New York, NY · 2 weeks ago · **5 applicants**

 \$144,000/yr - \$202,000/yr (LinkedIn est.) · Full-time

 10,001+ employees · Internet Publishing

 121 connections · 1,732 company alumni · 5,302 school alumni

 Your profile shows most of the skills for this job

[Apply](#) [Save](#)

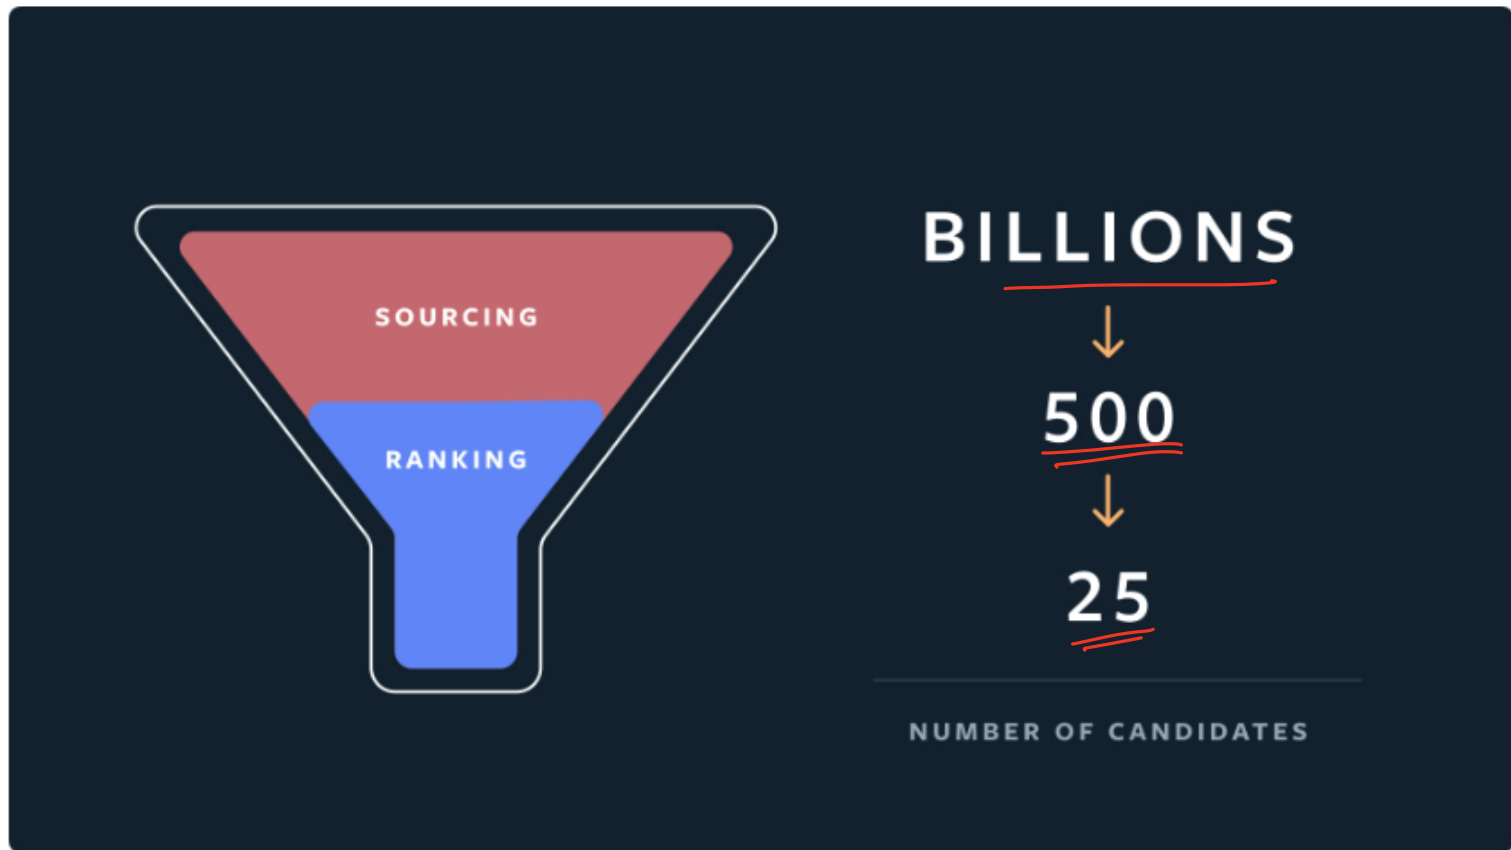
Description

Are you passionate about AI for recommendation systems? Do you want to influence the content that customers see at Amazon.com? Our recommendation services team designs and implements scalable machine learning solutions to personalize and optimize customer experience across Amazon retail pages. We are currently expanding in New York, and are looking for an applied scientist to join us in this exciting journey.

As An Applied Scientist, You Will

- Push the boundaries of real-world ranking, recommendation, and optimization systems
- Support science, engineering and product development on a scale only seen at Amazon.
- Champion and define best practices to maximize learnings while mentoring more junior scientists and engineers.
- Obsess over customer needs and satisfaction

Image/Video Ranking (Instagram)



Instagram Recommendations

Image/Video Ranking (Instagram Candidate Generation)

One type of Explore source



Instagram Recommendations → ref

Image/Video Ranking (Amazon Video)

Behavior-based Popularity Ranking on Amazon Video

With the growth in the number of video streaming services, providers have to strive hard to make relevant content available and keep customers engaged. A good experience would help customers discover new and popular videos to stream with ease. Customer streaming behavior tends to be a strong indicator of whether they found a video engaging. Aggregate customer behavior serves as a useful predictor of popularity. We discuss the use of past streaming behavior to learn patterns and predict a video's popularity using tree ensembles.

CCS Concepts: • **Information systems** → **Recommender systems**; **Content ranking**.

Additional Key Words and Phrases: video streaming; popularity; recommendations

ACM Reference Format:

. 2020. Behavior-based Popularity Ranking on Amazon Video. In *Fourteenth ACM Conference on Recommender Systems (RecSys '20)*, September 22–26, 2020, Virtual Event, Brazil. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3383313.3411555>

1 INTRODUCTION

On streaming platforms customers use two main approaches to find videos—search and discovery. Customers can enter a query to search for a video or browse the catalog to discover content. Streaming platforms use customer preferences,

References for lecture

Useful references both for this course and for an interview!

- ① YouTube Recommendation System
- ② Instagram Recommendation System
- ③ AliBaba Recommendations
- ④ Pinterest Recommendation System
- ⑤ Amazon Video Recommendations



Next Class



- 1 Ranking Loss functions
- 2 Learning to Rank (classic ML problem in Search)
- 3 Search Ranking Challenges
- 4 Ranking metrics (different from recommendation metrics)