Job recommendations leveraging Deep Learning using Analytics Zoo on Apache Spark and BigDL

Job Search

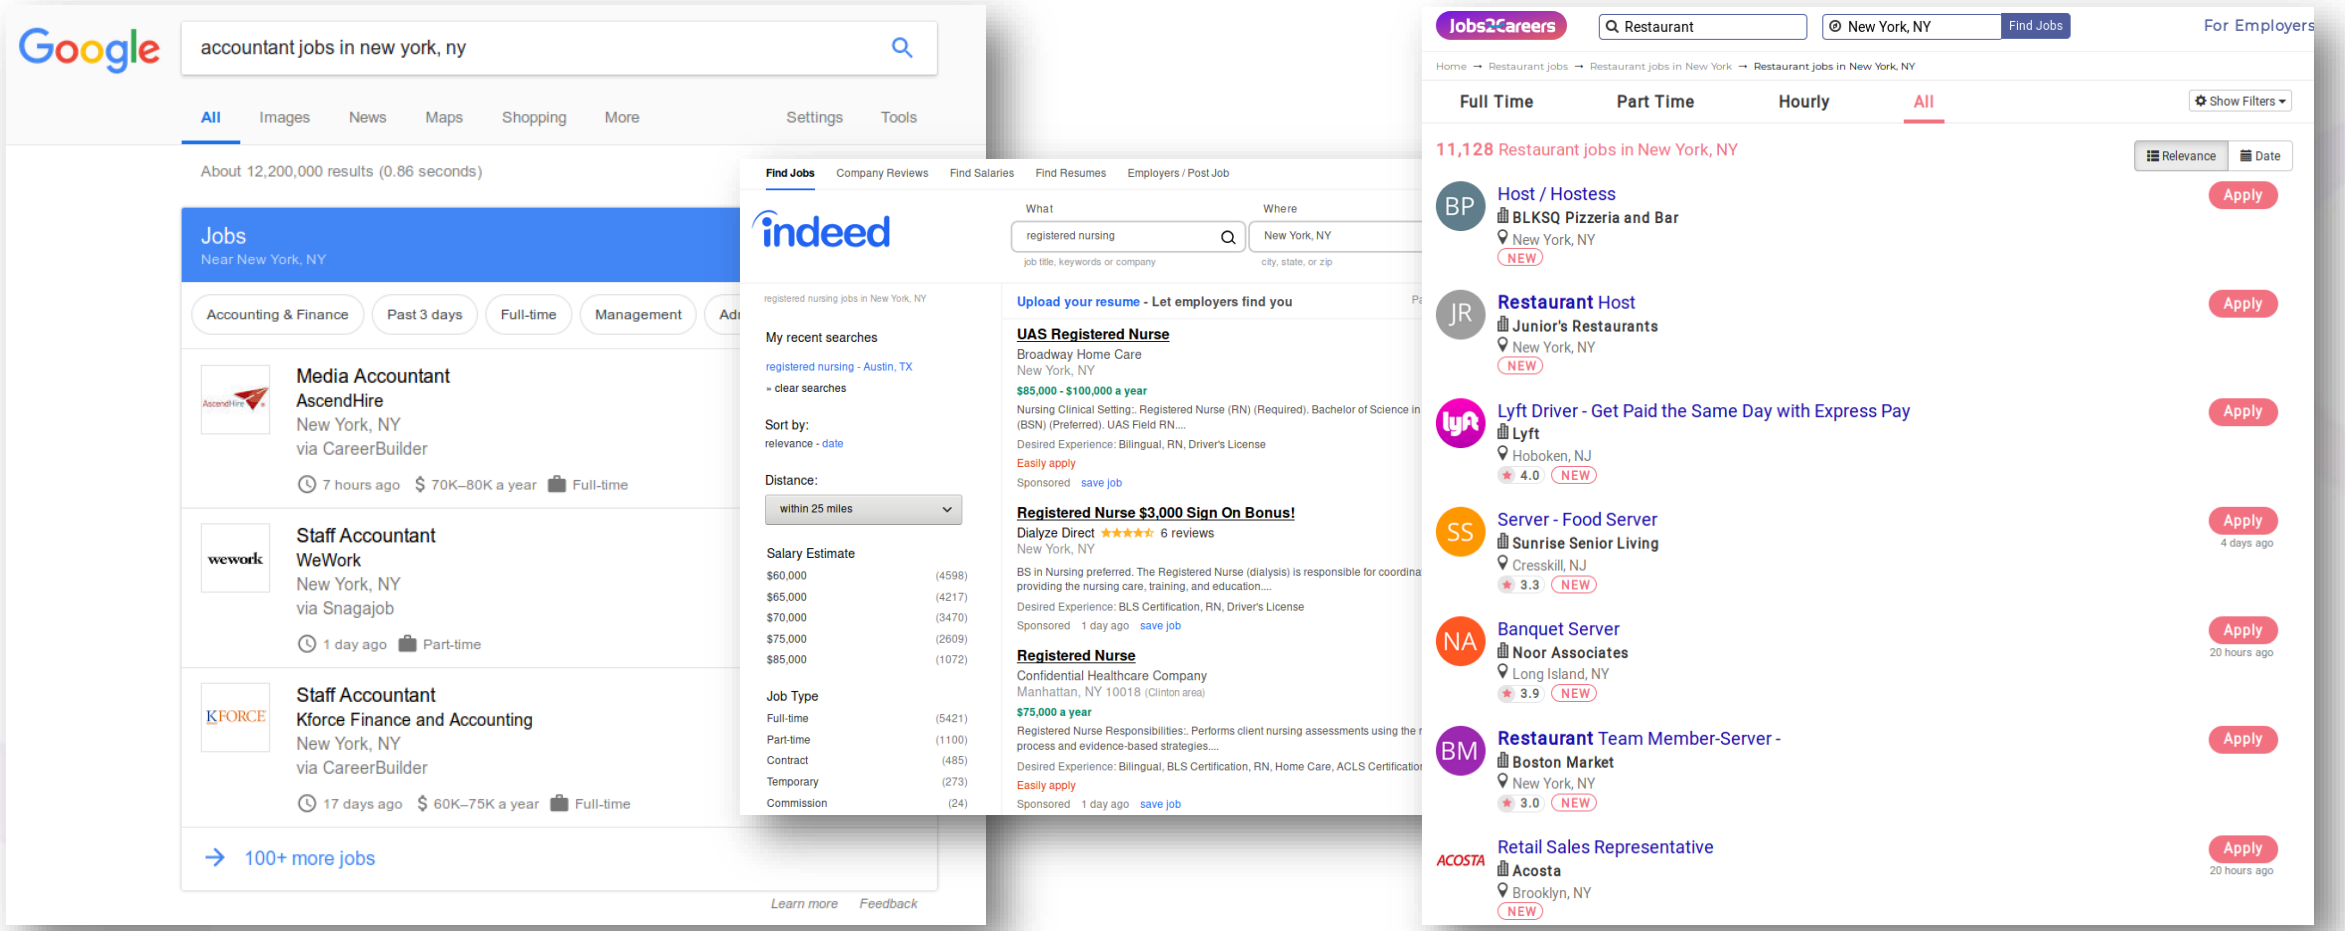
Google
accountant jobs in new york, ny
All Images News Maps Shopping More Settings Tools
About 12,200,000 results (0.86 seconds)
Jobs
Near New York, NY
Accounting & Finance
Past 3 days
Full-time
Management
Media Accountant
AscendHire
New York, NY
via CareerBuilder
7 hours ago
70K–80K a year
Full-time
Staff Accountant
WeWork
New York, NY
via Snagajob
1 day ago
Part-time
Staff Accountant
Kforce Finance and Accounting
New York, NY
via CareerBuilder
17 days ago
60K–75K a year
Full-time
100+ more jobs
Learn more
Feedback
Find Jobs
Company Reviews
Find Salaries
Find Resumes
Employers / Post Job
indeed
What
registered nursing
Where
New York, NY
job title, keywords or company
city, state, or zip
registered nursing jobs in New York, NY
My recent searches
registered nursing - Austin, TX
» clear searches
Sort by:
relevance - date
Distance:
within 25 miles
Salary Estimate
$60,000 (4598)
$65,000 (4217)
$70,000 (3470)
$75,000 (2609)
$85,000 (1072)
Job Type
Full-time (5421)
Part-time (1100)
Contract (485)
Temporary (273)
Commission (24)
Upload your resume - Let employers find you
UAS Registered Nurse
Broadway Home Care
New York, NY
$85,000 - $100,000 a year
Nursing Clinical Setting:. Registered Nurse (RN) (Required). Bachelor of Science in
(BSN) (Preferred). UAS Field RN....
Desired Experience: Bilingual, RN, Driver's License
Easily apply
Sponsored save job
Registered Nurse $3,000 Sign On Bonus!
Dialyze Direct 6 reviews
New York, NY
BS in Nursing preferred. The Registered Nurse (dialysis) is responsible for coordina
providing the nursing care, training, and education....
Desired Experience: BLS Certification, RN, Driver's License
Sponsored 1 day ago save job
Registered Nurse
Confidential Healthcare Company
Manhattan, NY 10018 (Clinton area)
$75,000 a year
Registered Nurse Responsibilities:. Performs client nursing assessments using the
process and evidence-based strategies....
Desired Experience: Bilingual, BLS Certification, RN, Home Care, ACLS Certificatio
Easily apply
Sponsored 1 day ago save job
Jobs2Careers
Restaurant
New York, NY
Find Jobs
For Employers
Home → Restaurant jobs → Restaurant jobs in New York → Restaurant jobs in New York, NY
Full Time
Part Time
Hourly
All
Show Filters
11,128 Restaurant jobs in New York, NY
Relevance
Date
BP
Host / Hostess
BLKSQ Pizzeria and Bar
New York, NY
NEW
Apply
JR
Restaurant Host
Junior's Restaurants
New York, NY
NEW
Apply
Lyft Driver - Get Paid the Same Day with Express Pay
Lyft
Hoboken, NJ
4.0
NEW
Apply
SS
Server - Food Server
Sunrise Senior Living
Cresskill, NJ
3.3
NEW
Apply
4 days ago
NA
Banquet Server
Noor Associates
Long Island, NY
3.9
NEW
Apply
20 hours ago
BM
Restaurant Team Member-Server -
Boston Market
New York, NY
3.0
NEW
Apply
ACOSTA
Retail Sales Representative
Acosta
Brooklyn, NY
NEW
Apply
20 hours ago

Speakers

Wenjing Zhan
Data Scientist, Talroo

Guoqiong Song
Software Engineer, Intel

Jacob Eisinger
Director of Data, Talroo

Agenda

- Talroo Overview
- Job Search Challenges
- Resume Search Opportunity
- Analytics Zoo and BigDL Overview
- Resume Search Analytics Zoo Solution
- Lessons Learned

Agenda

- Talroo Overview
- Job Search Challenges
- Resume Search Opportunity
- Analytics Zoo and BigDL Overview
- Resume Search Analytics Zoo Solution
- Lessons Learned

Where Talent and Recruiting Intersect

- Data-driven job ad network
- Self-funded growth
- Collaborative and social culture
- Austin, TX headquarters

Inc. 5000 HONOR ROLL FIVE-TIME HONOREE

Reach

- Email
- SMS/text
- Search engine
- Associations
- Blogs
- Push notifications
- Niche job sites
- General job boards

Talroo reaches active candidates where they are searching.

Efficiency

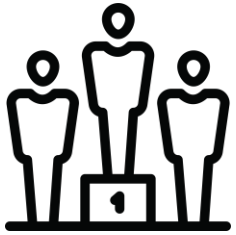

Unique candidates

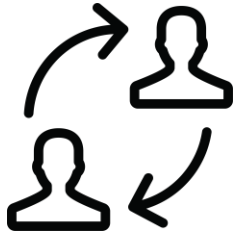

Better matching

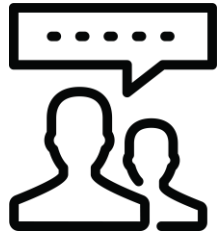

Client service

- Talroo turns up what works, and turns down what doesn't, ultimately reducing cost-per-hire.

Search

We power search.

Talroo Monthly Scale

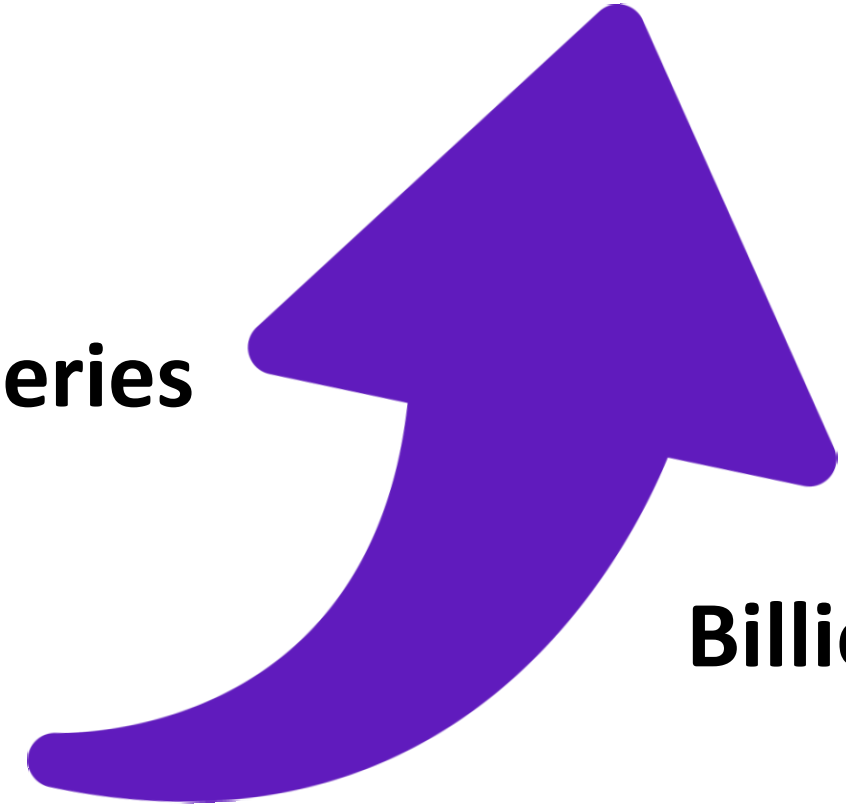

Tens of Millions of Interactions

Billions of Queries

Billions of Jobs Events

Tens of Billions Impressions

Agenda

- Talroo Overview
- Job Search Challenges
- Resume Search Opportunity
- Analytics Zoo and BigDL Overview
- Resume Search Analytics Zoo Solution
- Lessons Learned

Traditional Information Retrieval Sufferings

Solution challenges:
stemming, synonyms,
ontologies, sensitivity

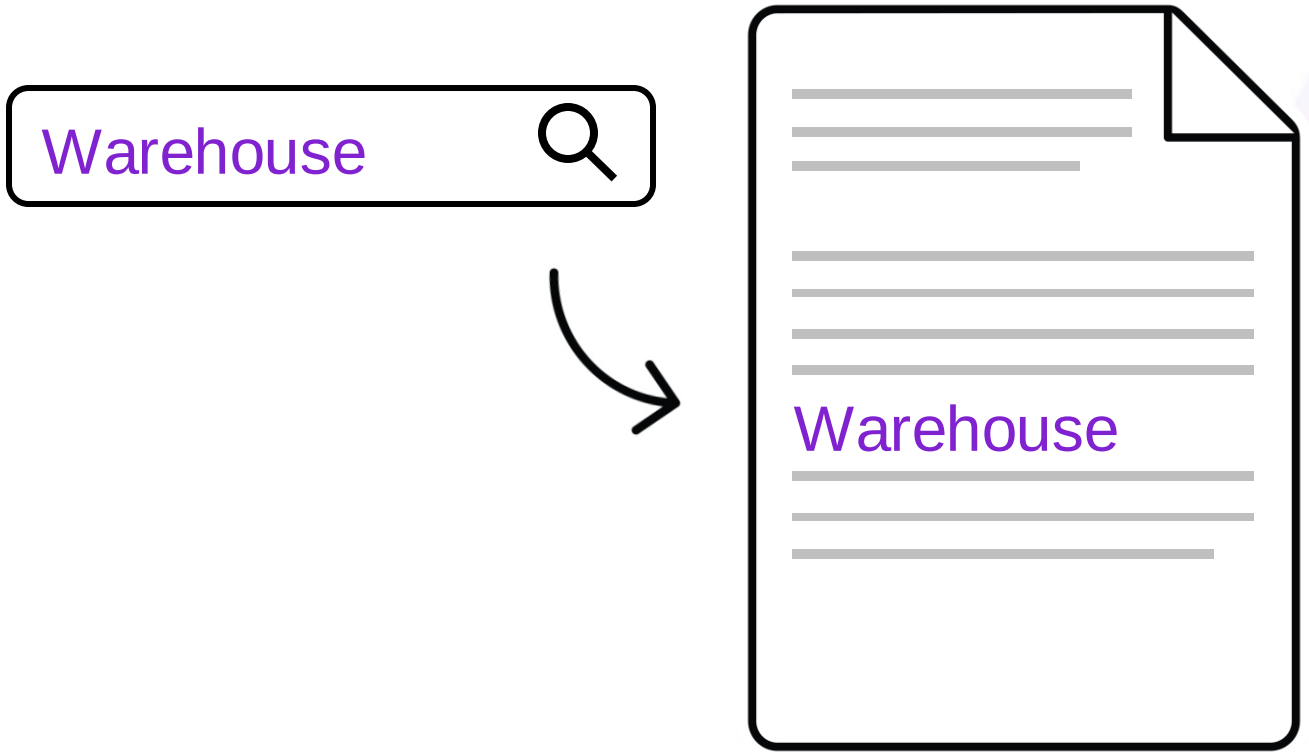

Stemming

Accountant =!= Accounting

Stemming Solution

Accountant =!= Accounting

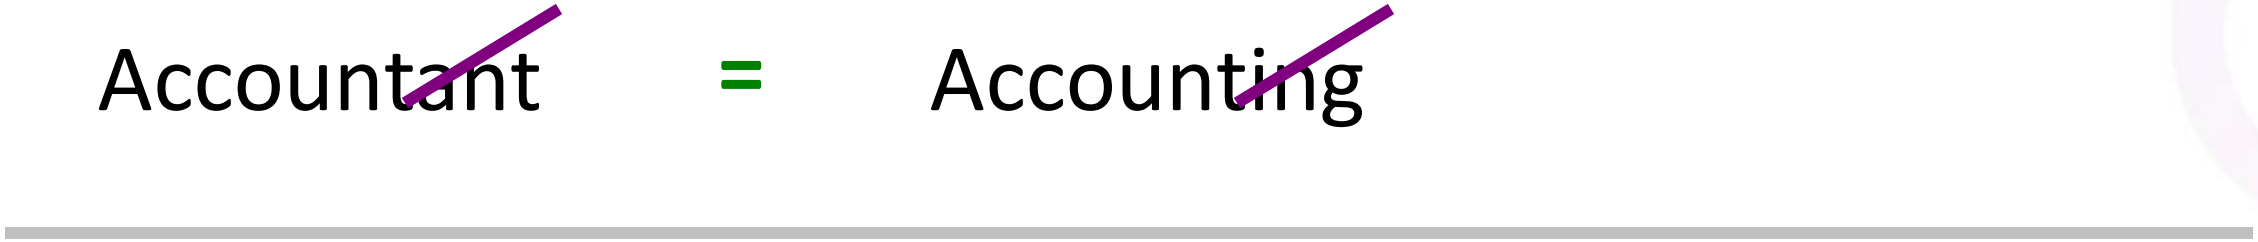

Stemming Sufferings

Accountant =!= Accounting

Accountant = Accounting

Accountant = Accounting = Account Representative

Synonyms

Registered Nurse =!= RN

Synonyms Solution

Registered Nurse	=!=	RN
Registered Nurse	=	RN → registered nurse

Synonym Sufferings

Registered Nurse	=!=	RN
Registered Nurse	=	RN → registered nurse

. . .

Ontologies

Dishwasher =!= Back of House

Ontologies Solution

Dishwasher =!= Back of House

Dishwasher → Restaurant = Back of House → Restaurant

Ontologies Sufferings

Talroo™

Dishwasher =!= Back of House

Dishwasher → Restaurant = Back of House → Restaurant

. . .

Specificity Suffering

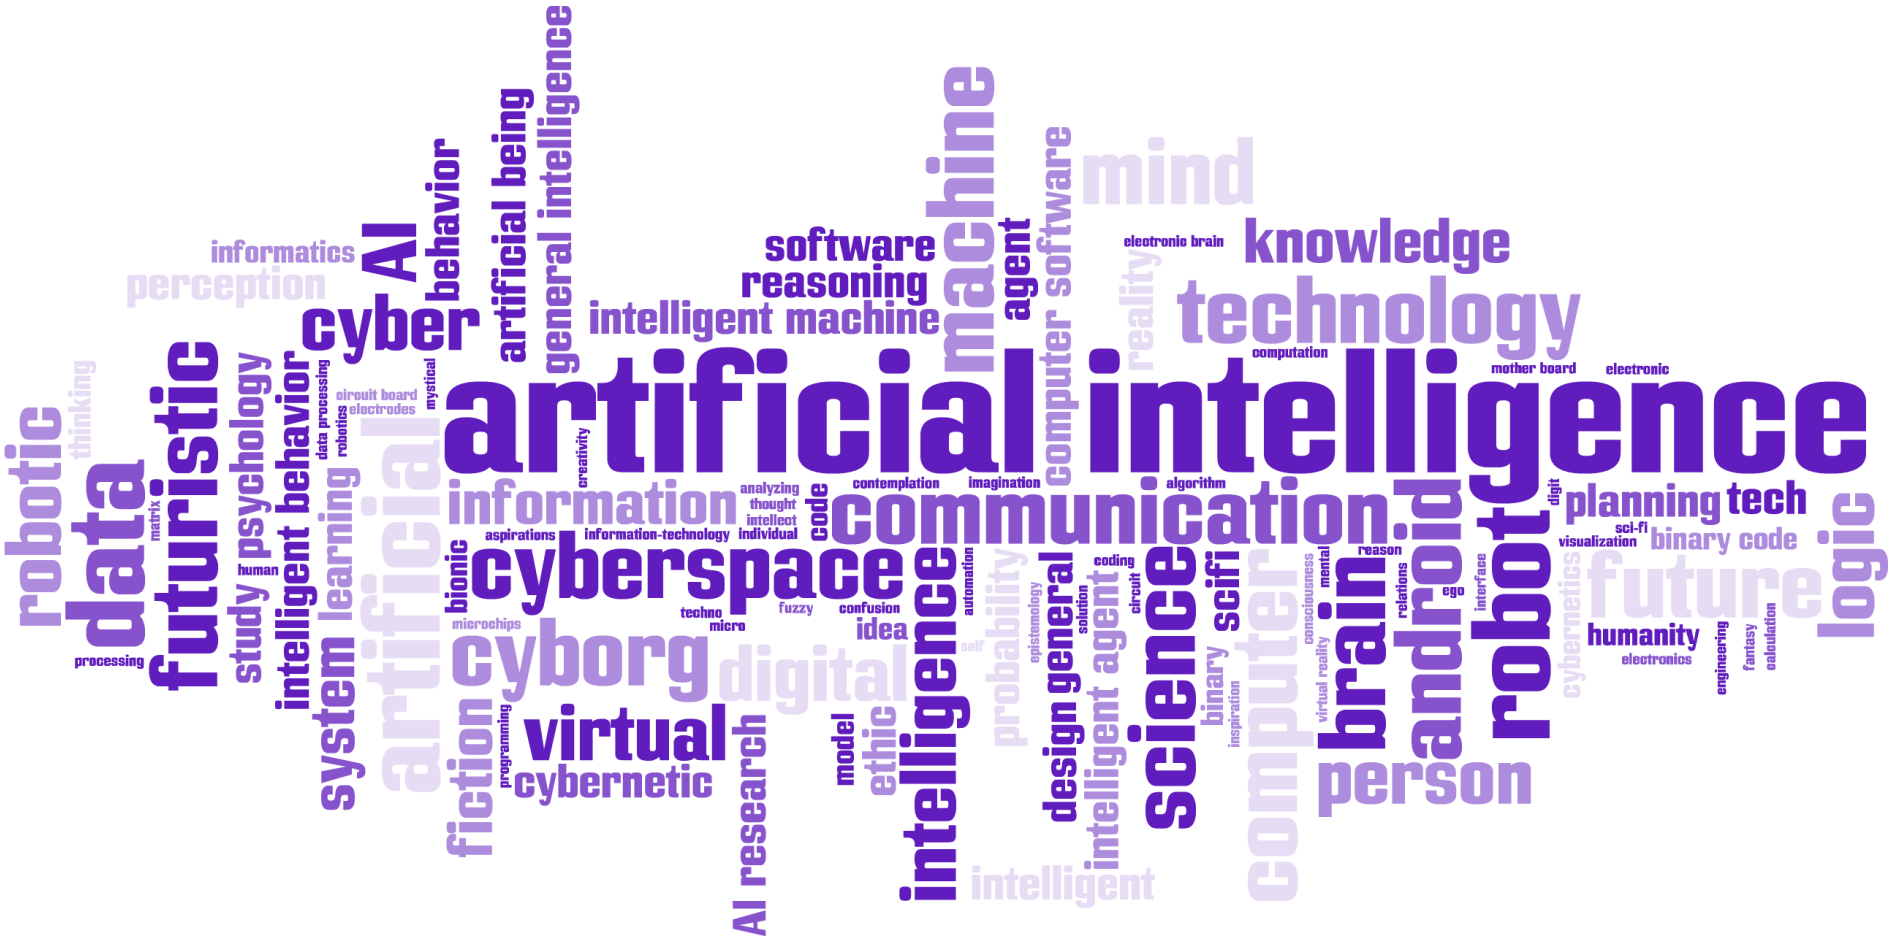

Personalize Results Value

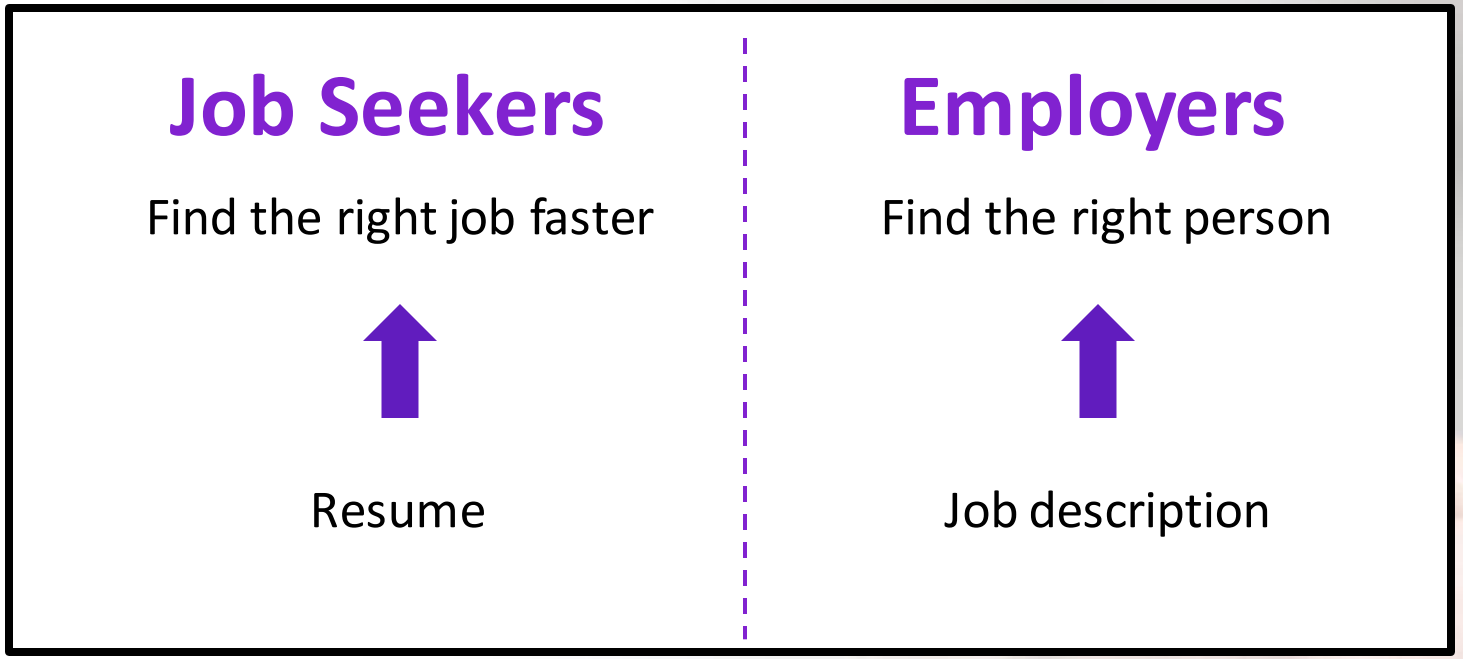

Agenda

- Talroo Overview
- Job Search Challenges
- Resume Search Opportunity
- Analytics Zoo and BigDL Overview
- Resume Search Analytics Zoo Solution
- Lessons Learned

Framing The Problem

- **What information** is logged ?
- **What do we do** with the logged information ?
- **What did we find** from logged information ?
- **How** to further tackle the findings ?

What Information Do We Log

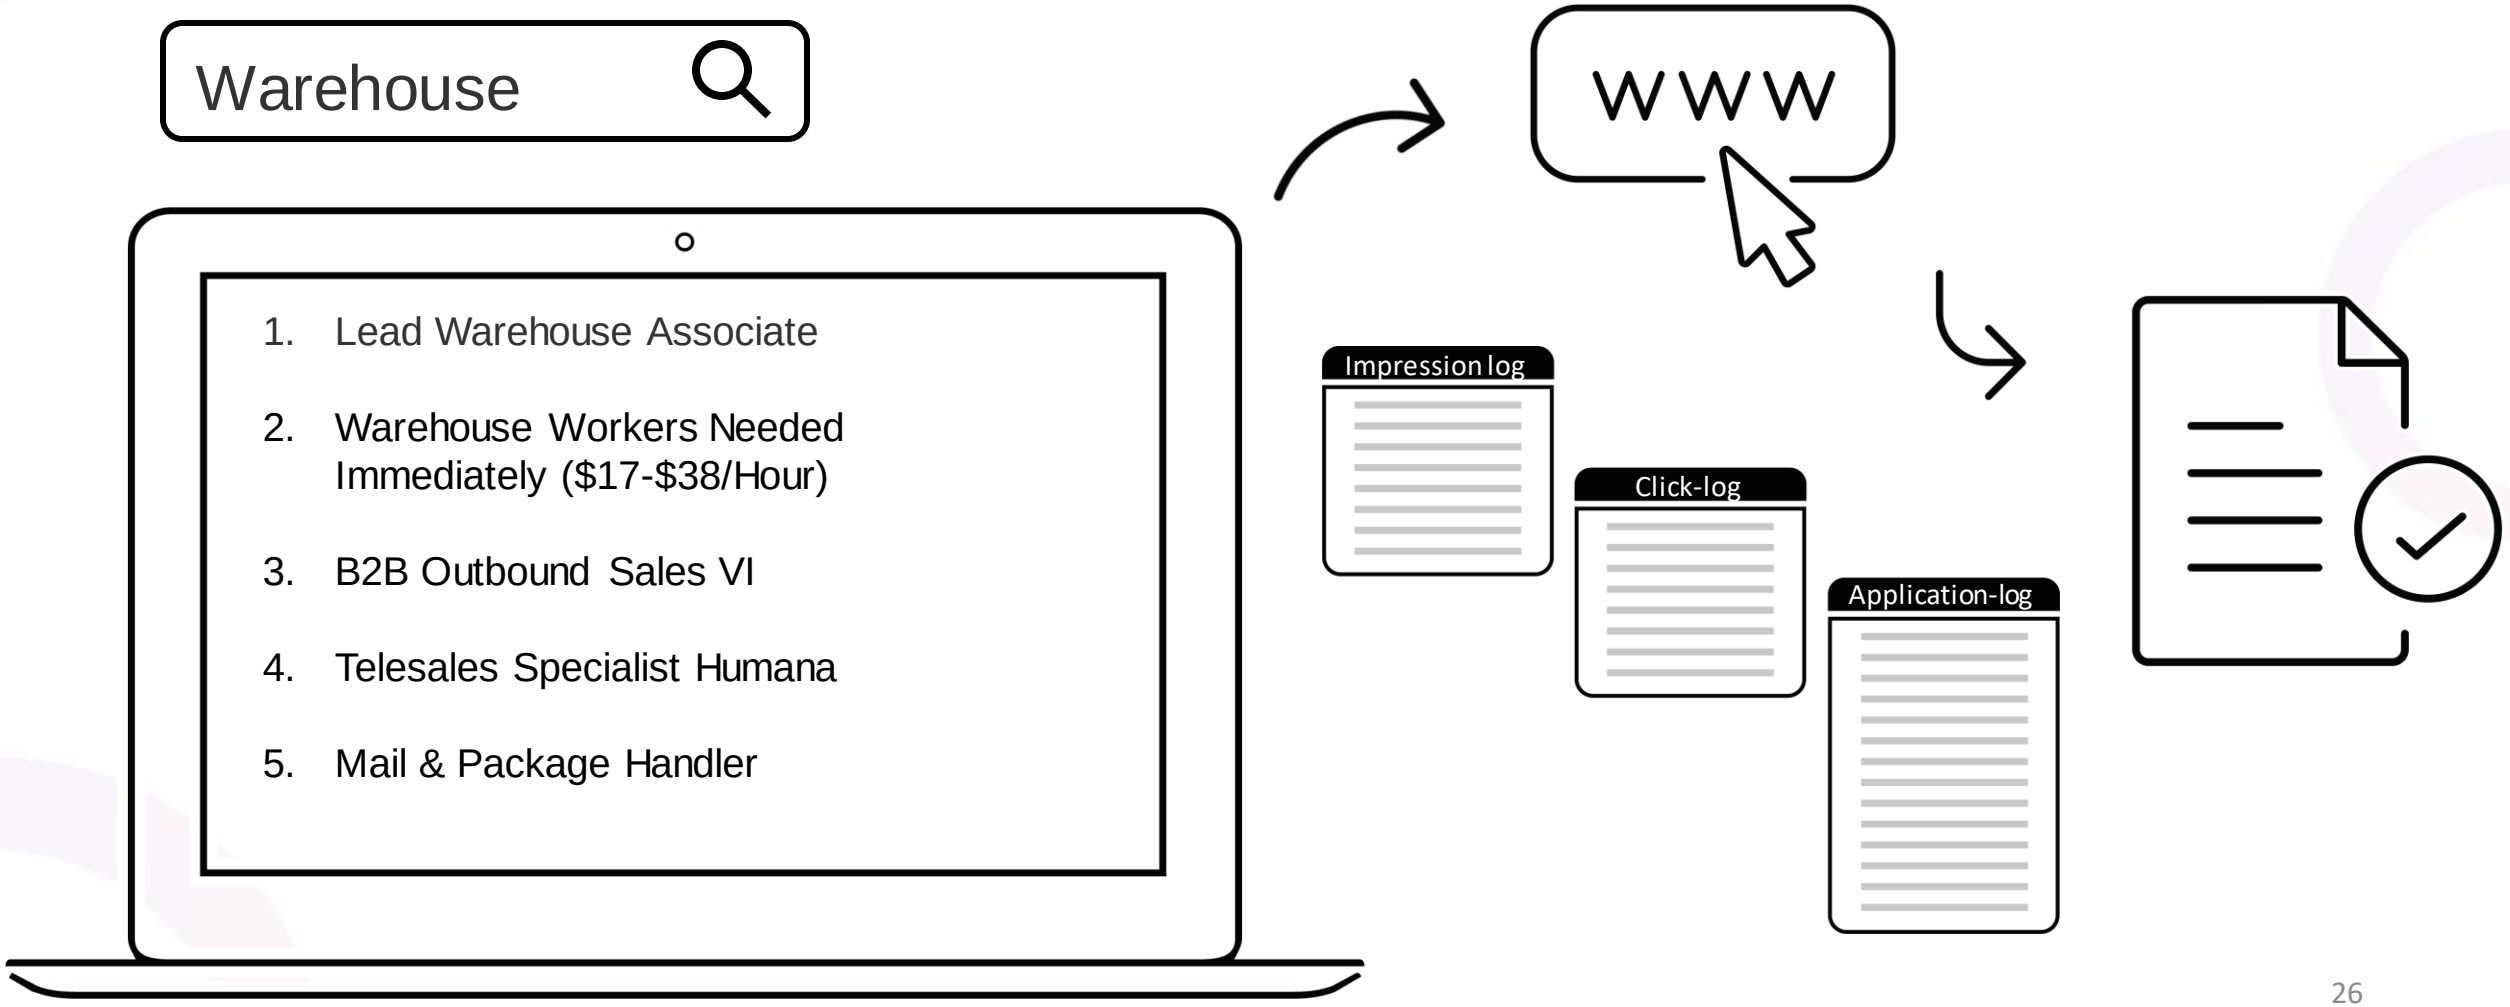

Replaying the Job Application Process

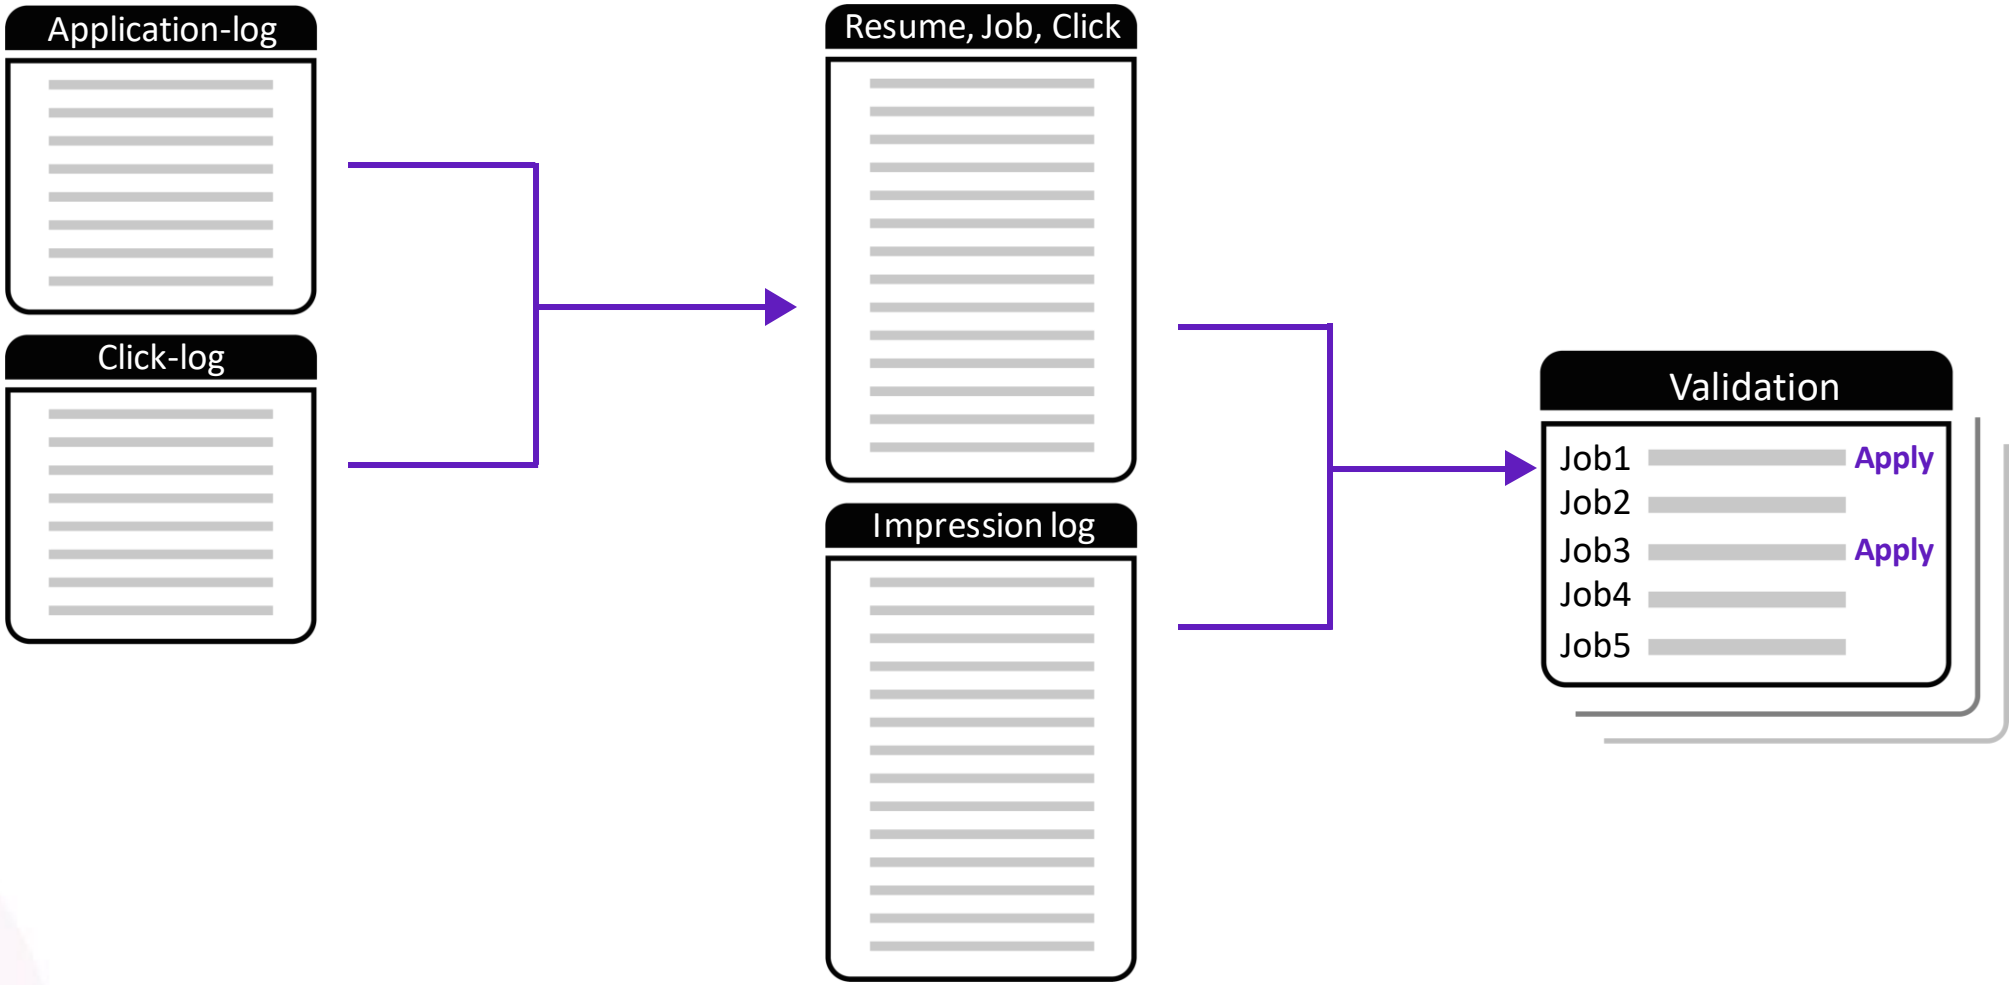

Real World Example

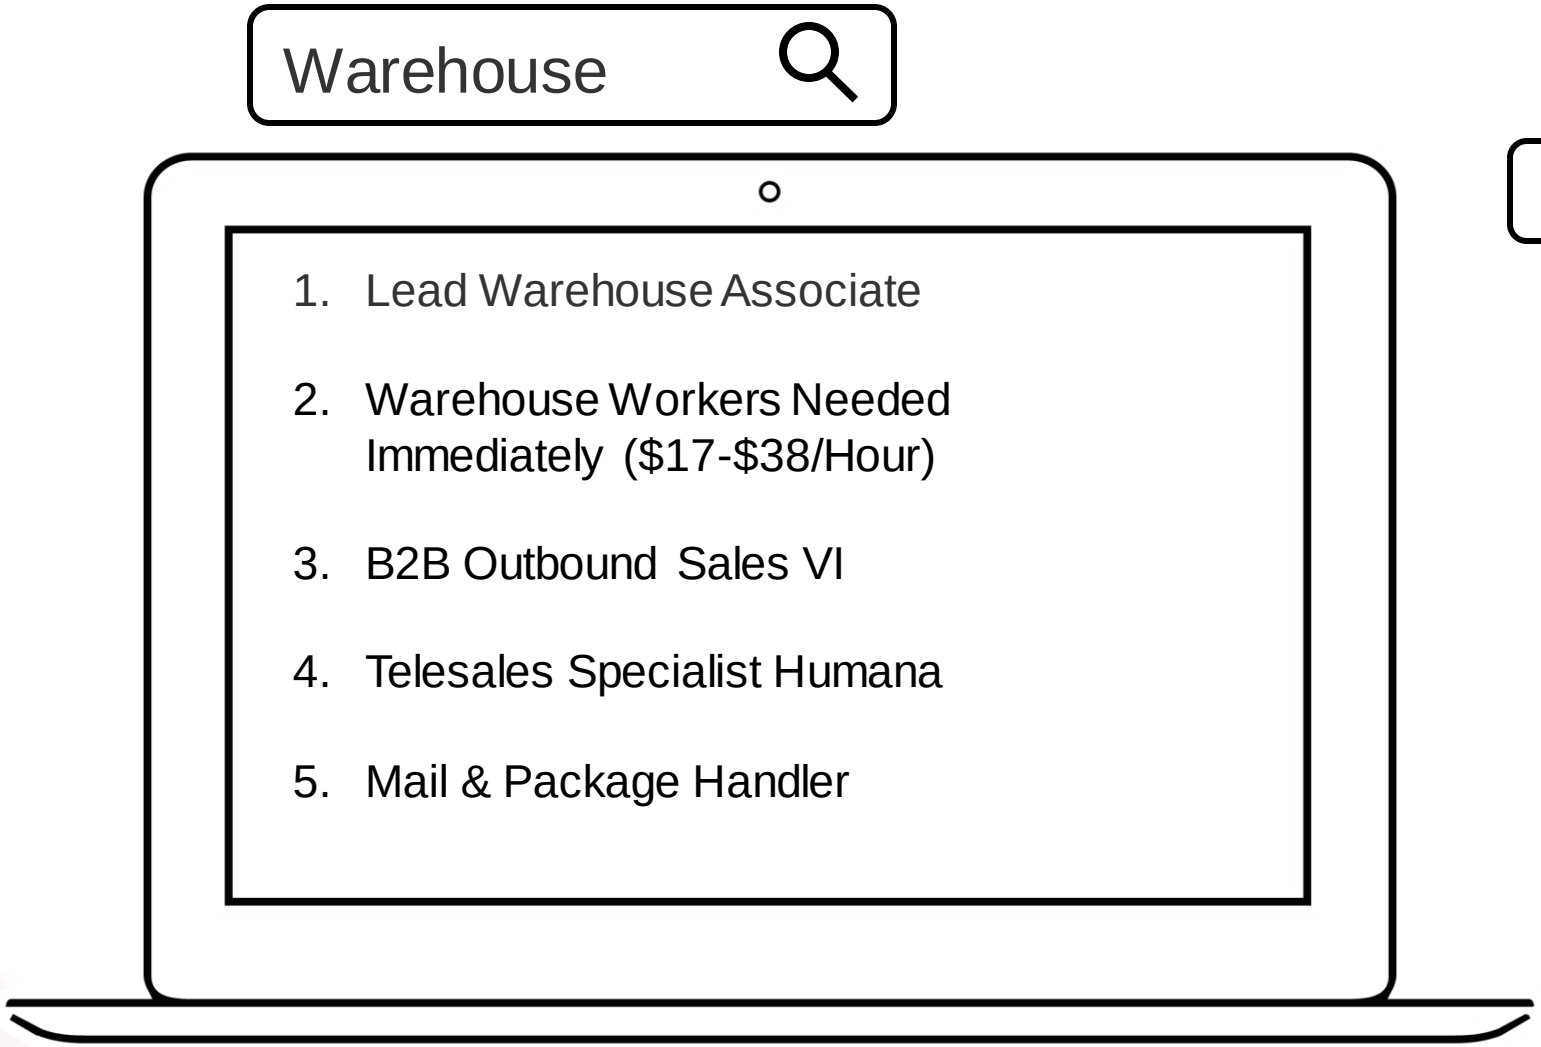

expect

1

2

3

4

5

Real World Example

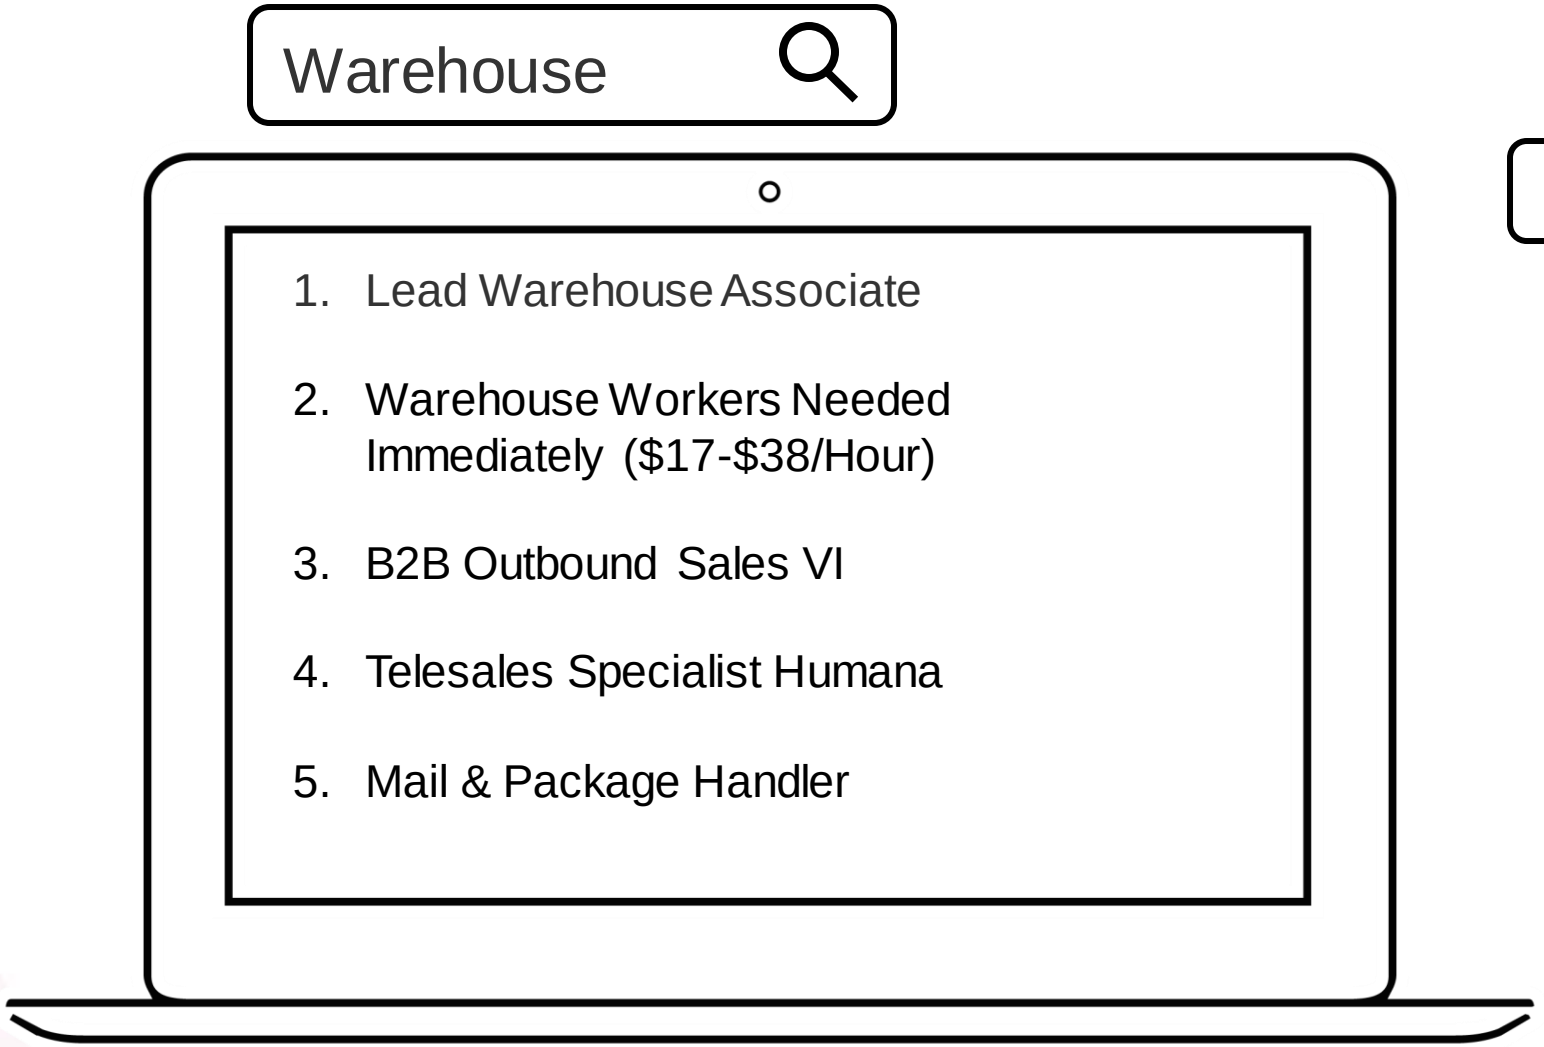

expect	reality
1	1
2	2
3	3
4	4
5	5

Real World Example

Bilingual International Sales Executive

09/2007 - 09/2017
Store Manager

03/2001 - 08/2007
International Sales Executive

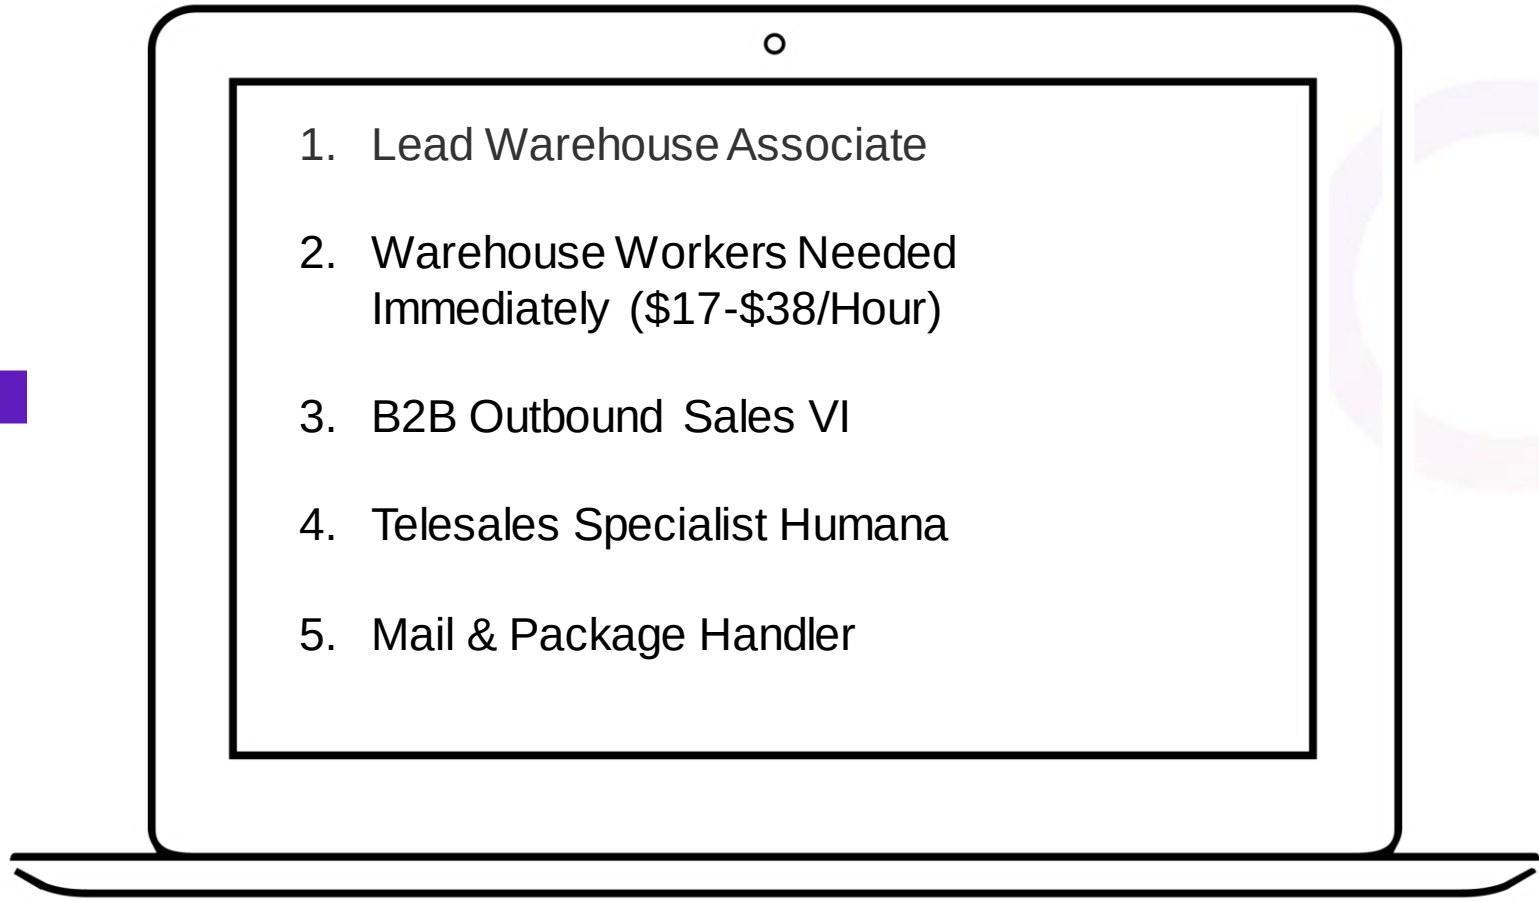

Real World Example

Bilingual International Sales Executive

09/2007 - 09/2017
Store Manager

03/2001 - 08/2007
International Sales Executive

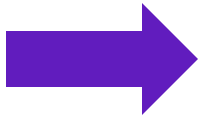

1. Lead Warehouse Associate
2. Warehouse Workers Needed Immediately ($17-$38/Hour)
3. B2B Outbound Sales VI
4. Telesales Specialist Humana
5. Mail & Package Handler

Investigation Of The Findings

-- First Attempt

Bilingual International Sales Executive

09/2007 - 09/2017
Store Manager

03/2001 - 08/2007
International Sales Executive

3. B2B Outbound Sales VI

Industry

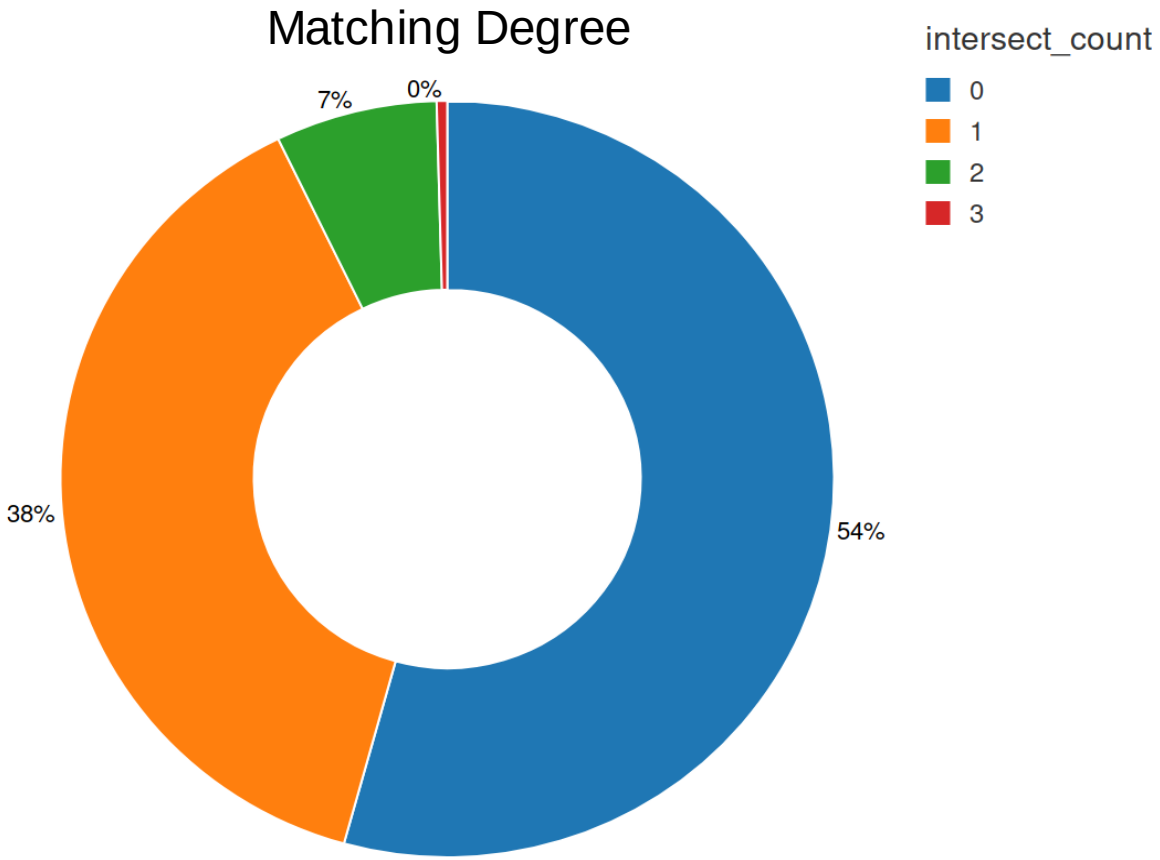

Investigation Of The Findings

-- Reframe The Solution

- Handle different terminology
- Identify and respect multiple industries
- Recognize other topics

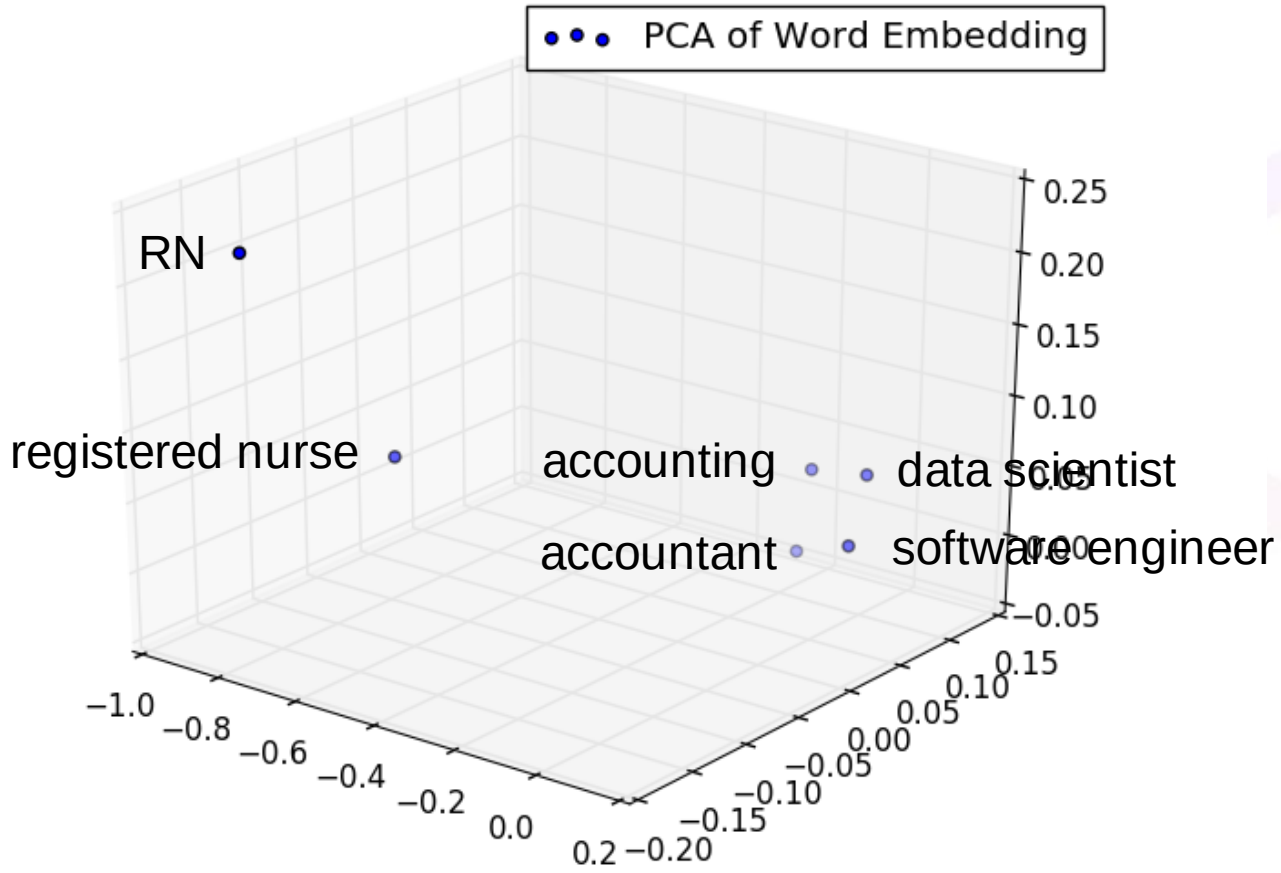

Ideal Resume Search Requirements

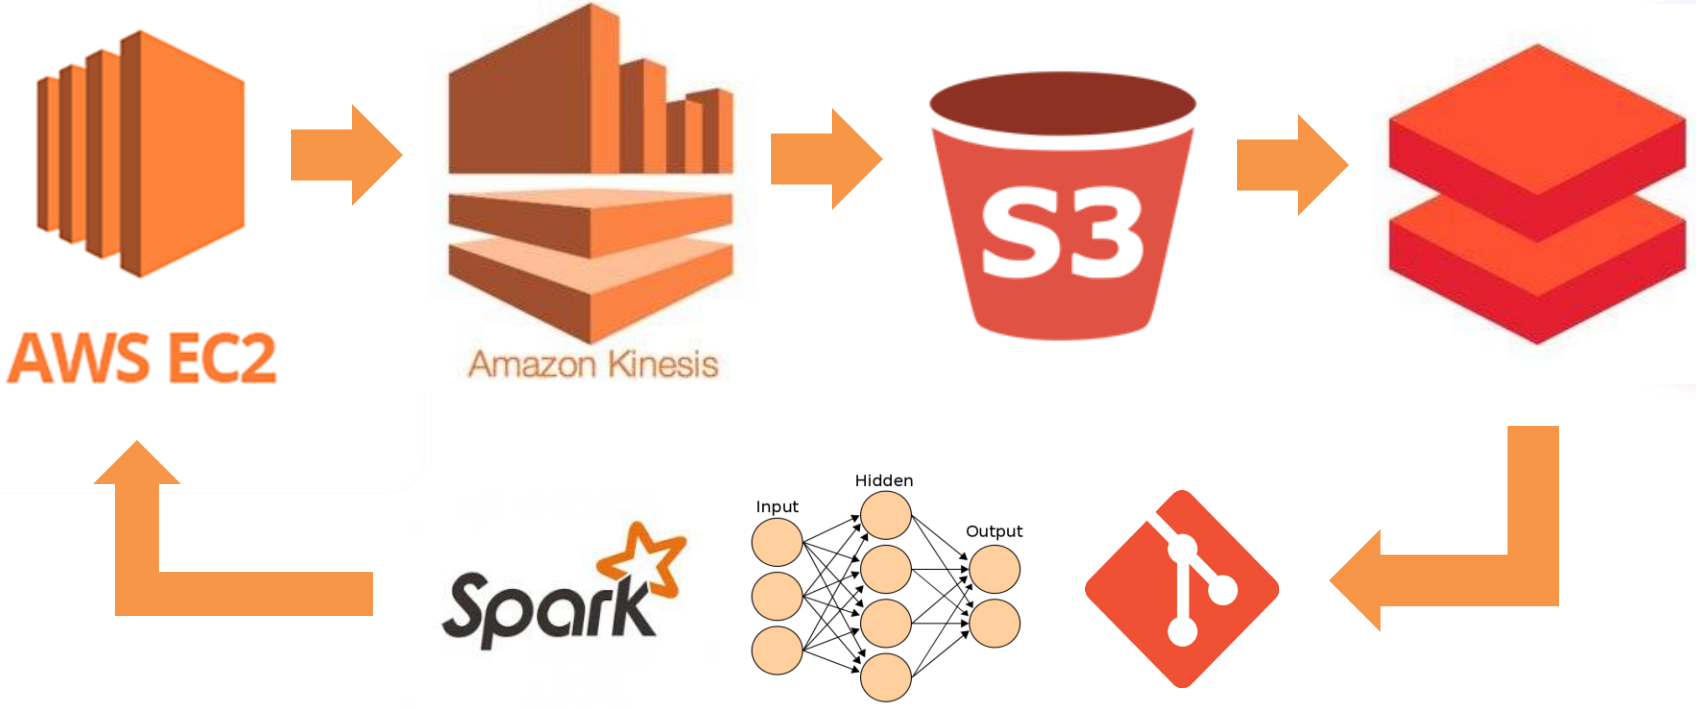

Scoring Before

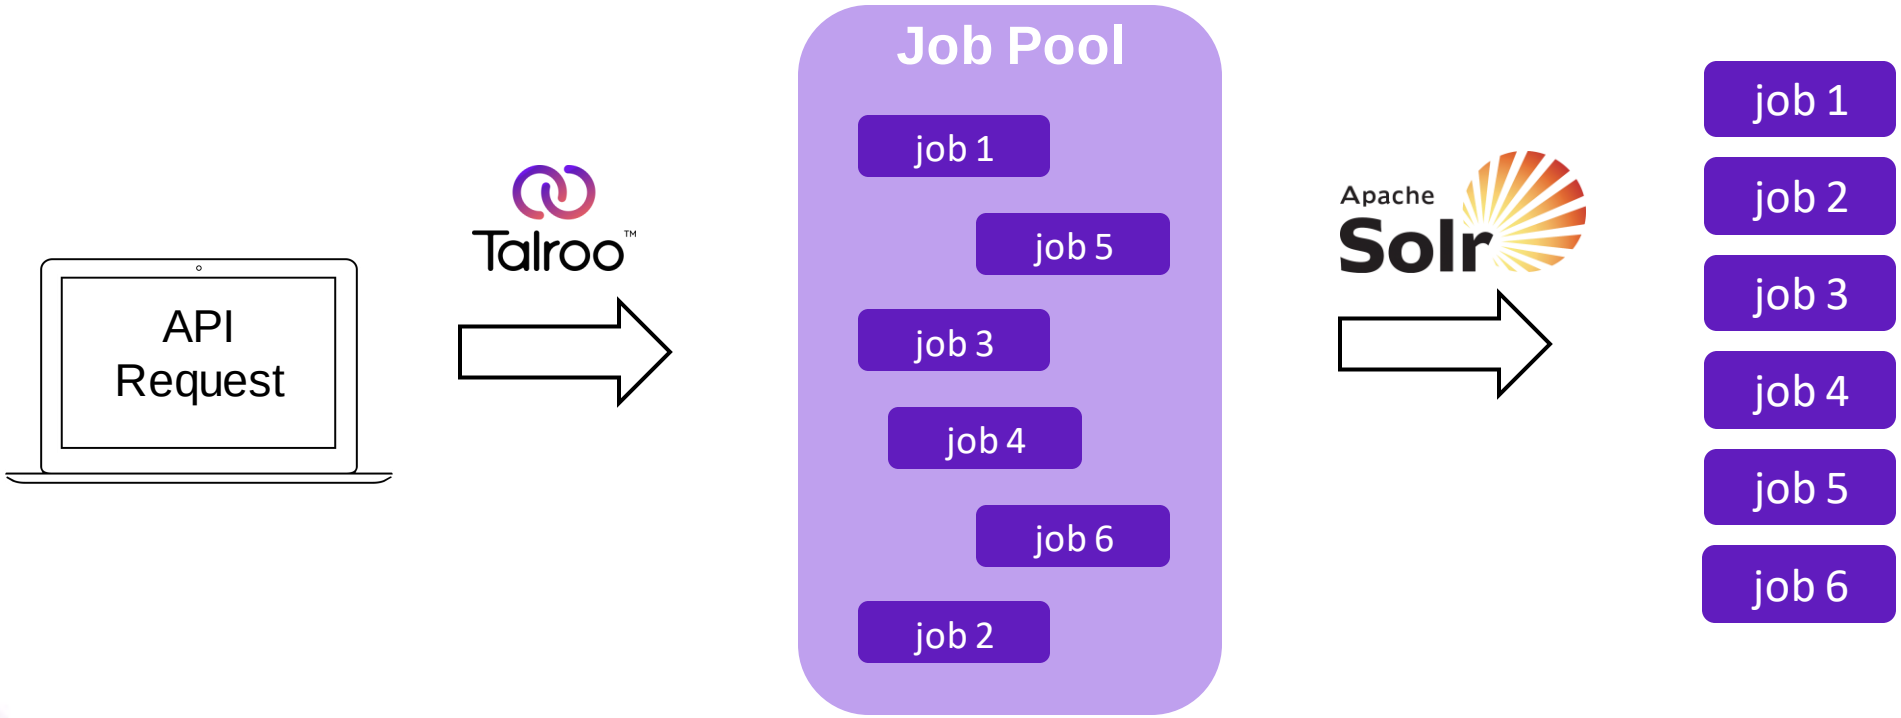

Scoring After - Reranker Solution

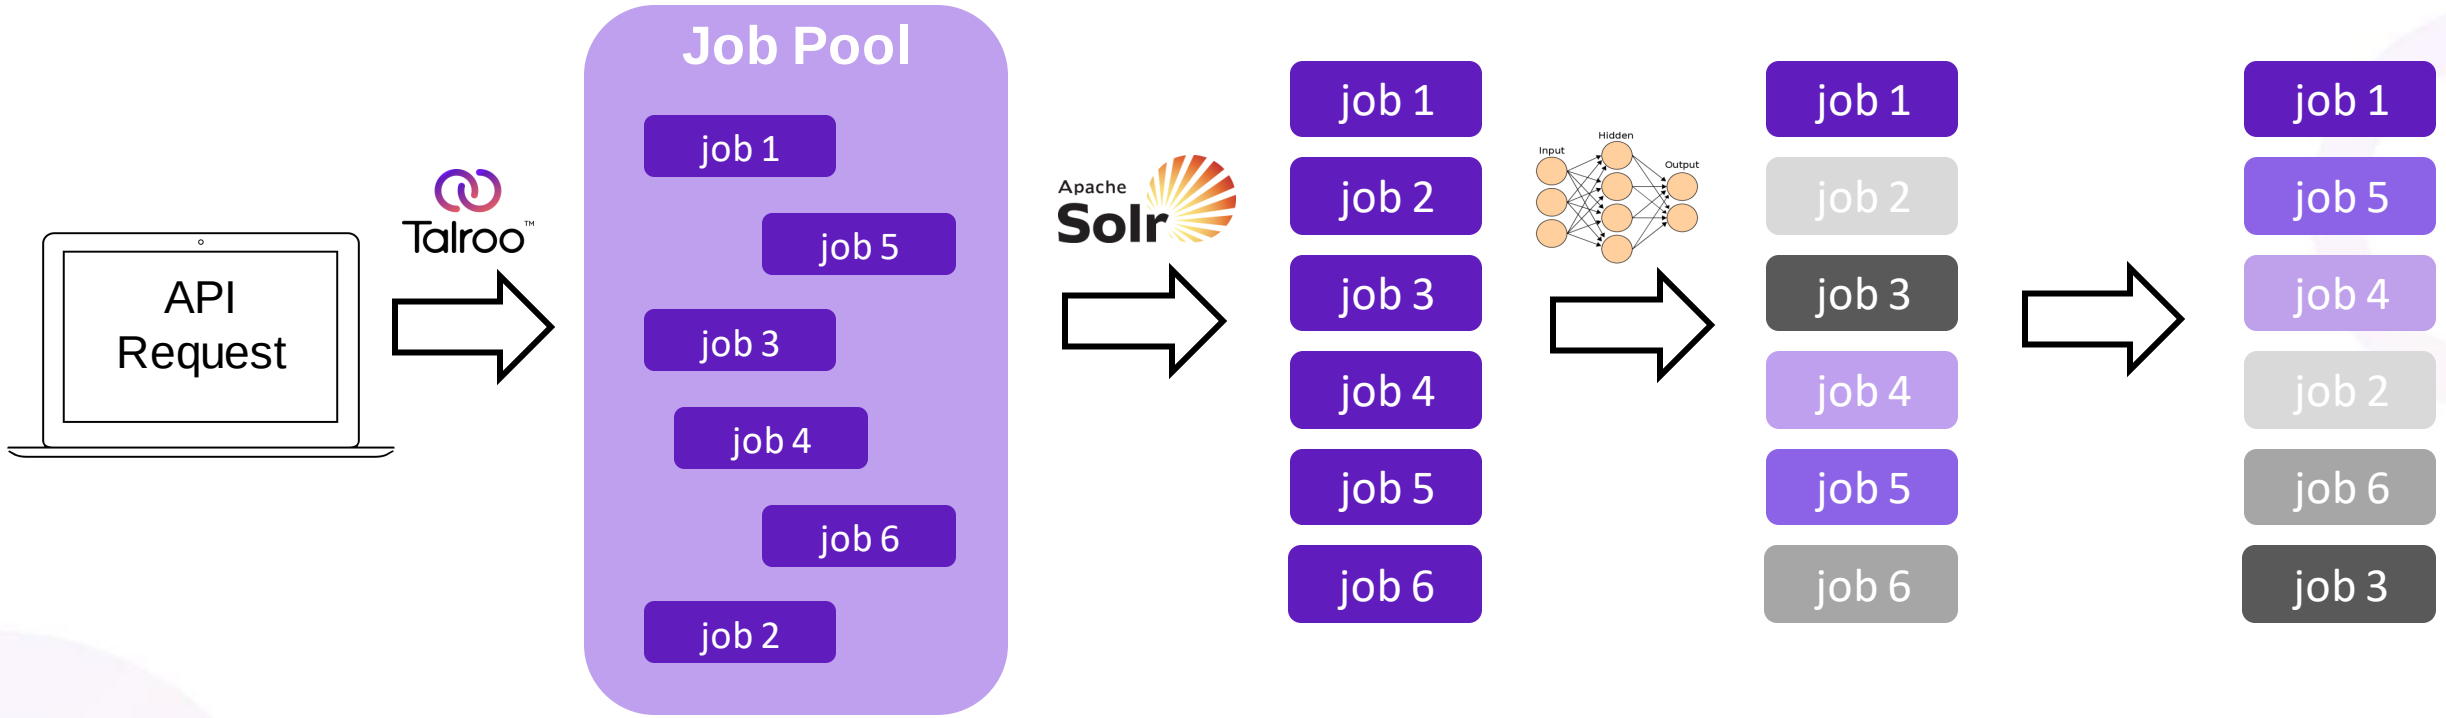

Evaluation Metric – Gold Standard

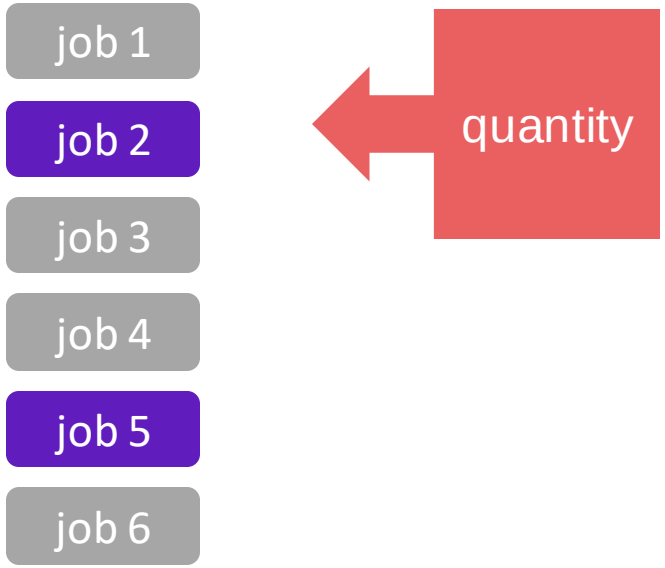

Online

$$\text{CTR} = \frac{\text{clicks}}{\text{impressions}}$$

https://en.wikipedia.org/wiki/Evaluation_measures_(information_retrieval)

Evaluation Metric - Quantity

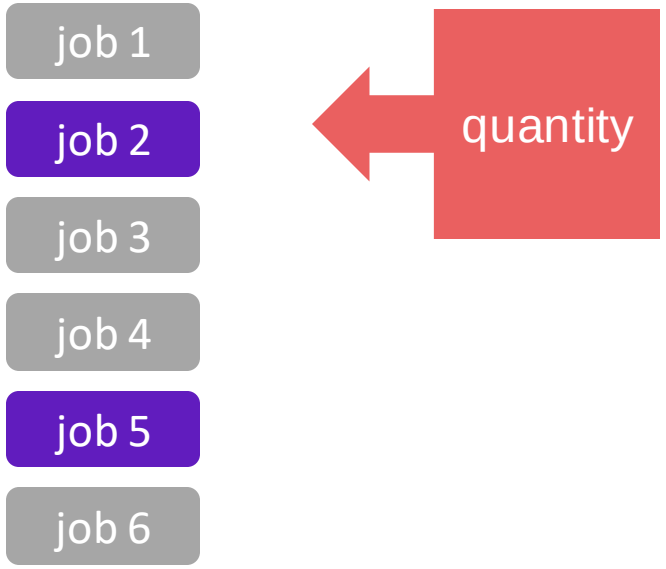

Online

$$\text{CTR} = \frac{\text{clicks}}{\text{impressions}}$$

Offline

$$\text{Precision} = \frac{\text{tp}}{\text{tp} + \text{fp}}$$

https://en.wikipedia.org/wiki/Evaluation_measures_(information_retrieval)

Evaluation Metric - Quality

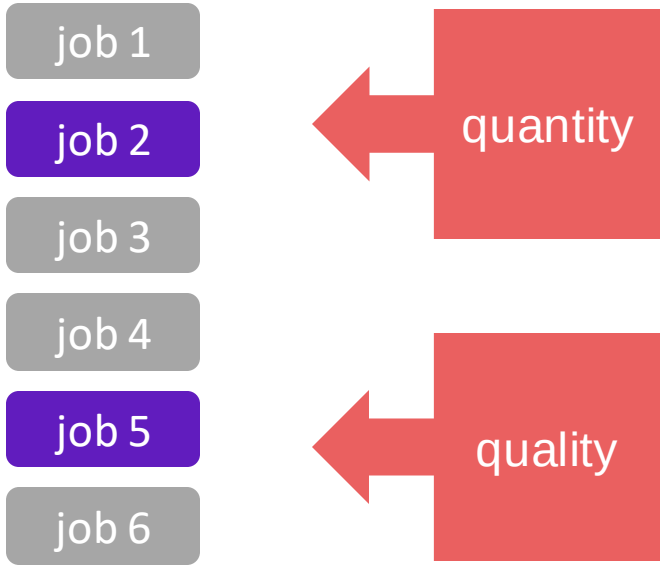

Online

$$\text{CTR} = \frac{\text{clicks}}{\text{impressions}}$$

Offline

$$\text{Precision} = \frac{\text{tp}}{\text{tp} + \text{fp}}$$

$$\text{MRR} = \frac{1}{\text{Q}} \sum_{\text{i}=1}^{\text{Q}} \frac{1}{\text{rank}_\text{i}}$$

https://en.wikipedia.org/wiki/Evaluation_measures_(information_retrieval)

Evaluation Metric - Cutoffs

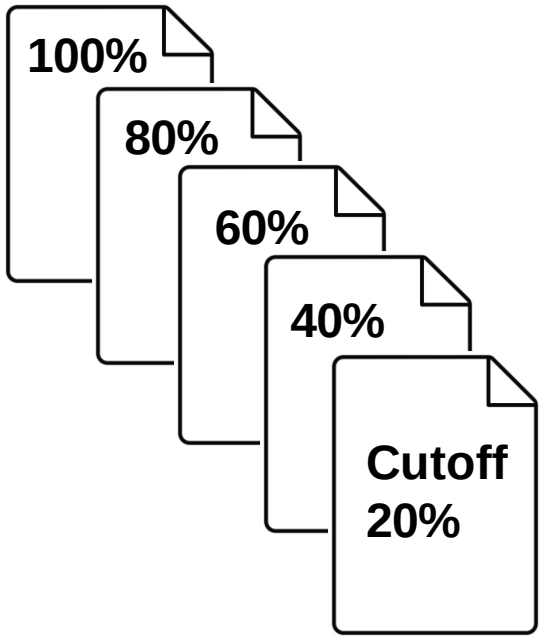

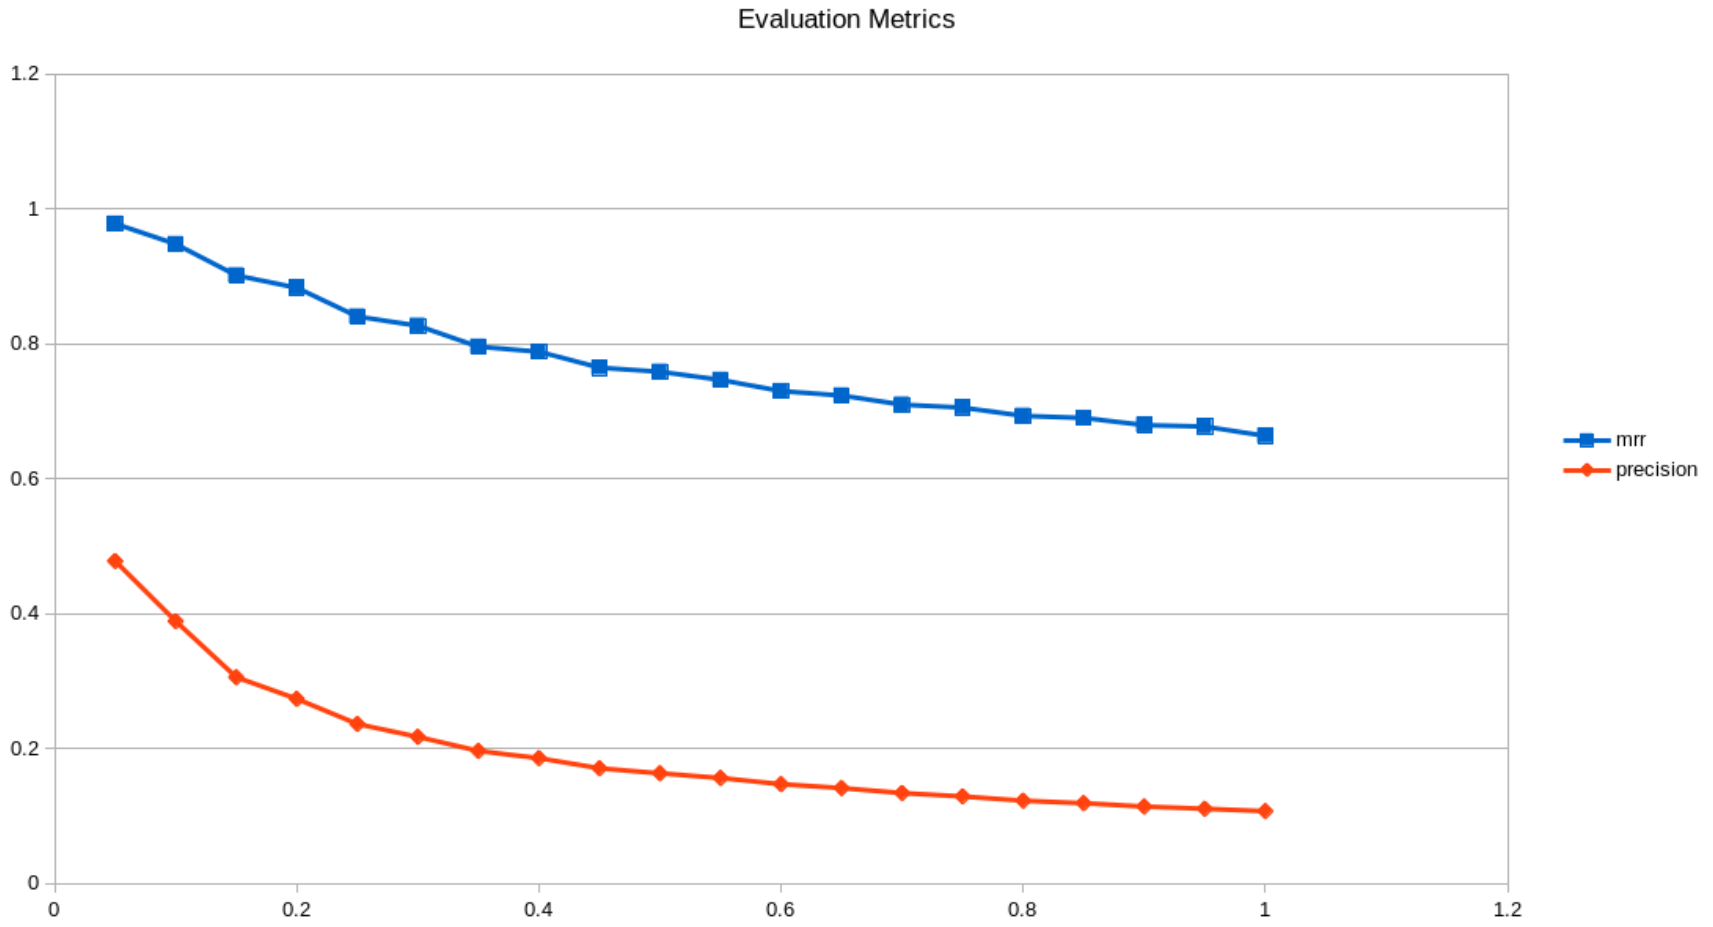

Resume Search Solution

Agenda

- Talroo Overview
- Job Search Challenges
- Resume Search Opportunity
- Analytics Zoo and BigDL Overview
- Resume Search Analytics Zoo Solution
- Lessons Learned

Intel BigDL

Talroo™

High Performance Deep Learning for **Apache Spark*** on **CPU Infrastructure**

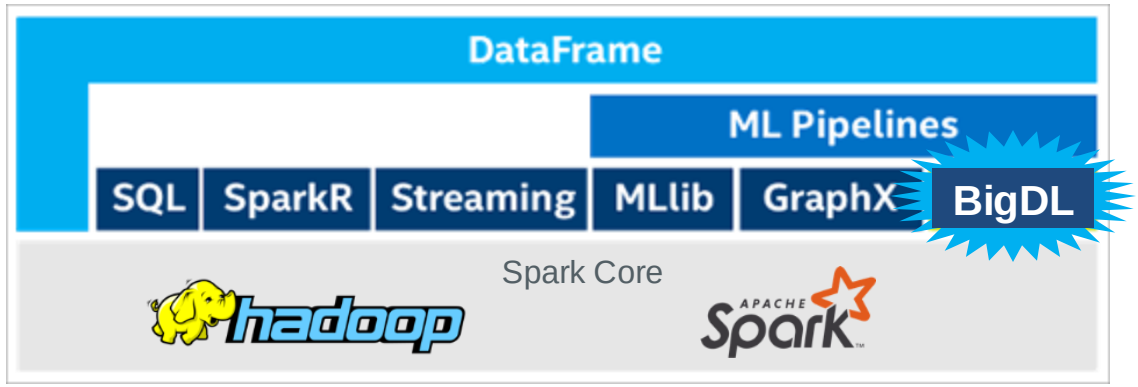

BigDL is an **open-source** distributed deep learning library for Apache Spark* that can run directly on top of existing Spark or Apache Hadoop* clusters

Ideal for DL Models TRAINING and INFERENCE

Designed and Optimized for Intel® Xeon®

No need to deploy costly accelerators, duplicate data, or suffer through scaling headaches!

Feature Parity & Model Exchange with TensorFlow*, Caffe*, Keras, Torch*

Lower TCO and improved ease of use with existing infrastructure

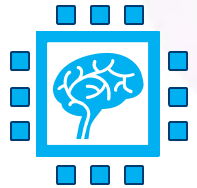

Deep Learning on Big Data Platform, Enabling **Efficient Scale-Out**

Powered by Intel® MKL and multi-threaded programming

https://github.com/intel-analytics/BigDL

https://bigdl-project.github.io/

Intel Analytics Zoo

Talroo™

Build and Productionize **Deep Learning Apps** for **Big Data at Scale**

Reference Use Cases	• Anomaly detection • Sentiment analysis • Fraud detection • Chatbot, sequence prediction, etc.
Built-In Deep Learning Models	• Image classification • Object detection • Text classification • Recommendations • Sequence-to-sequence, GAN, etc.
Feature Engineering	Feature transformations for • Image, text, 3D imaging, time series, speech, etc.
High-Level Pipeline APIs	• Native deep learning support in Spark DataFrames and ML Pipelines • Autograd, Keras and transfer learning APIs for model definition • Support for model serving/inference pipelines
Backends	Spark, BigDL, TensorFlow, etc.

https://github.com/intel-analytics/analytics-zoo/

https://analytics-zoo.github.io/

Agenda

- Talroo Overview
- Job Search Challenges
- Resume Search Opportunity
- Analytics Zoo and BigDL Overview
- Resume Search Analytics Zoo Solution
- Lessons Learned

Analytics Zoo solution

- Retrieve document vectors (resume and job) using Glove vectors
- Analytics Zoo recommender model
- End-to-End flow(offline training)
- Evaluation Results

Word Embeddings and GloVe Vectors

- Words or phrases from the vocabulary are mapped to **vectors of real numbers**.
- **Global log-bilinear regression model for the unsupervised** learning algorithm.
- Training is performed on aggregated global **word-word co-occurrence statistics** from a Wikipedia.
- Vector representations showcase meaningful **linear substructures** of the word vector space.

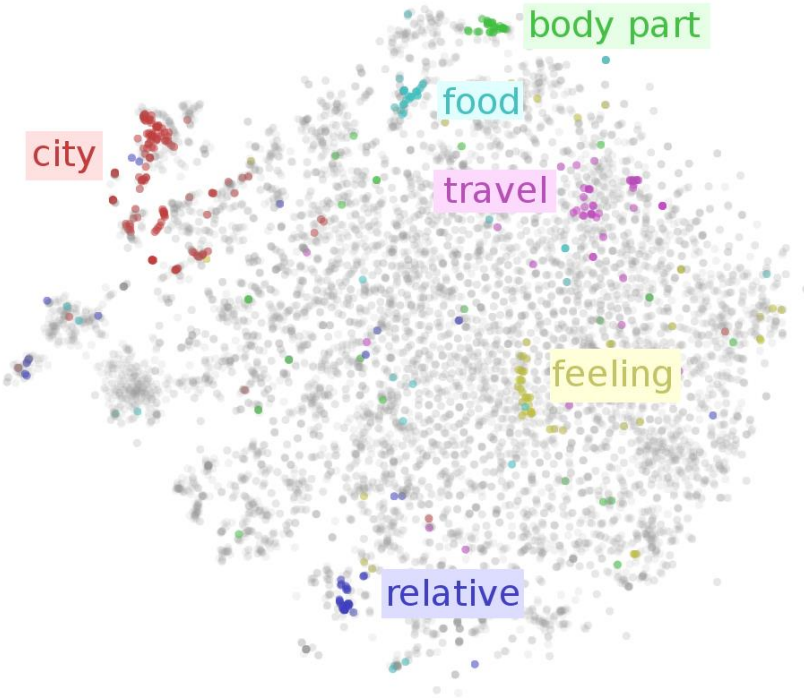

https://nlp.stanford.edu/projects/glove/

Analytics Zoo Recommender Model

- Neural collaborative filtering, Wide and Deep
- Answer the question using classification methodologies
- Implicit feedback and explicit feedback
- APIs
  - recommendForUser
  - recommendForItem
  - predictUserItemPair

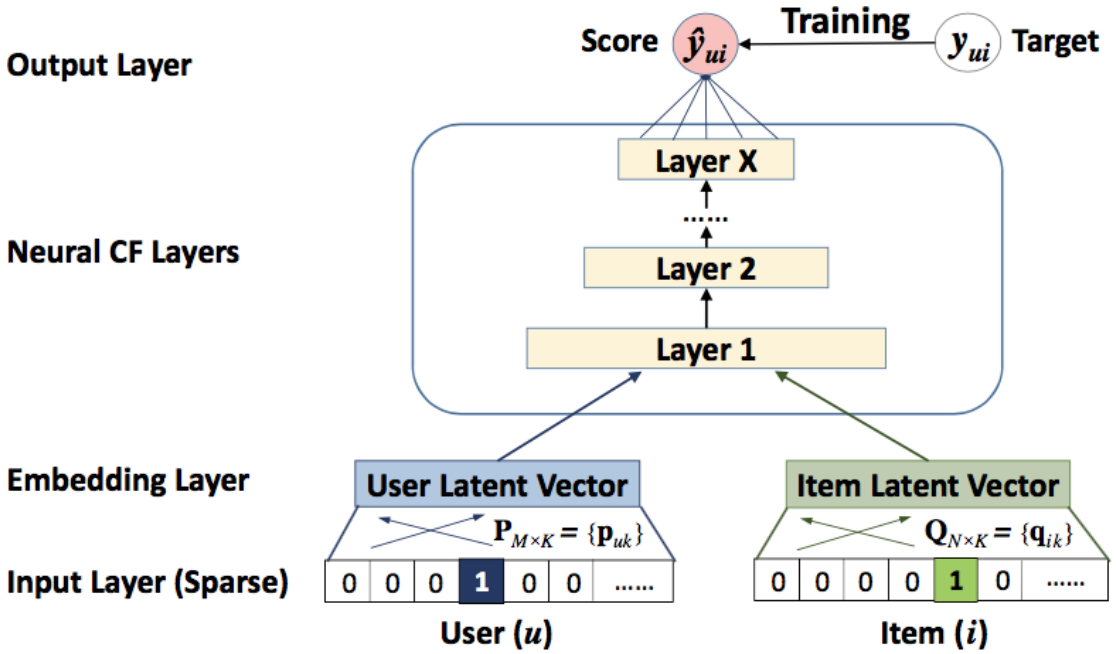

He, 2015

https://github.com/intel-analytics/analytics-zoo/

https://analytics-zoo.github.io/

Recommender model

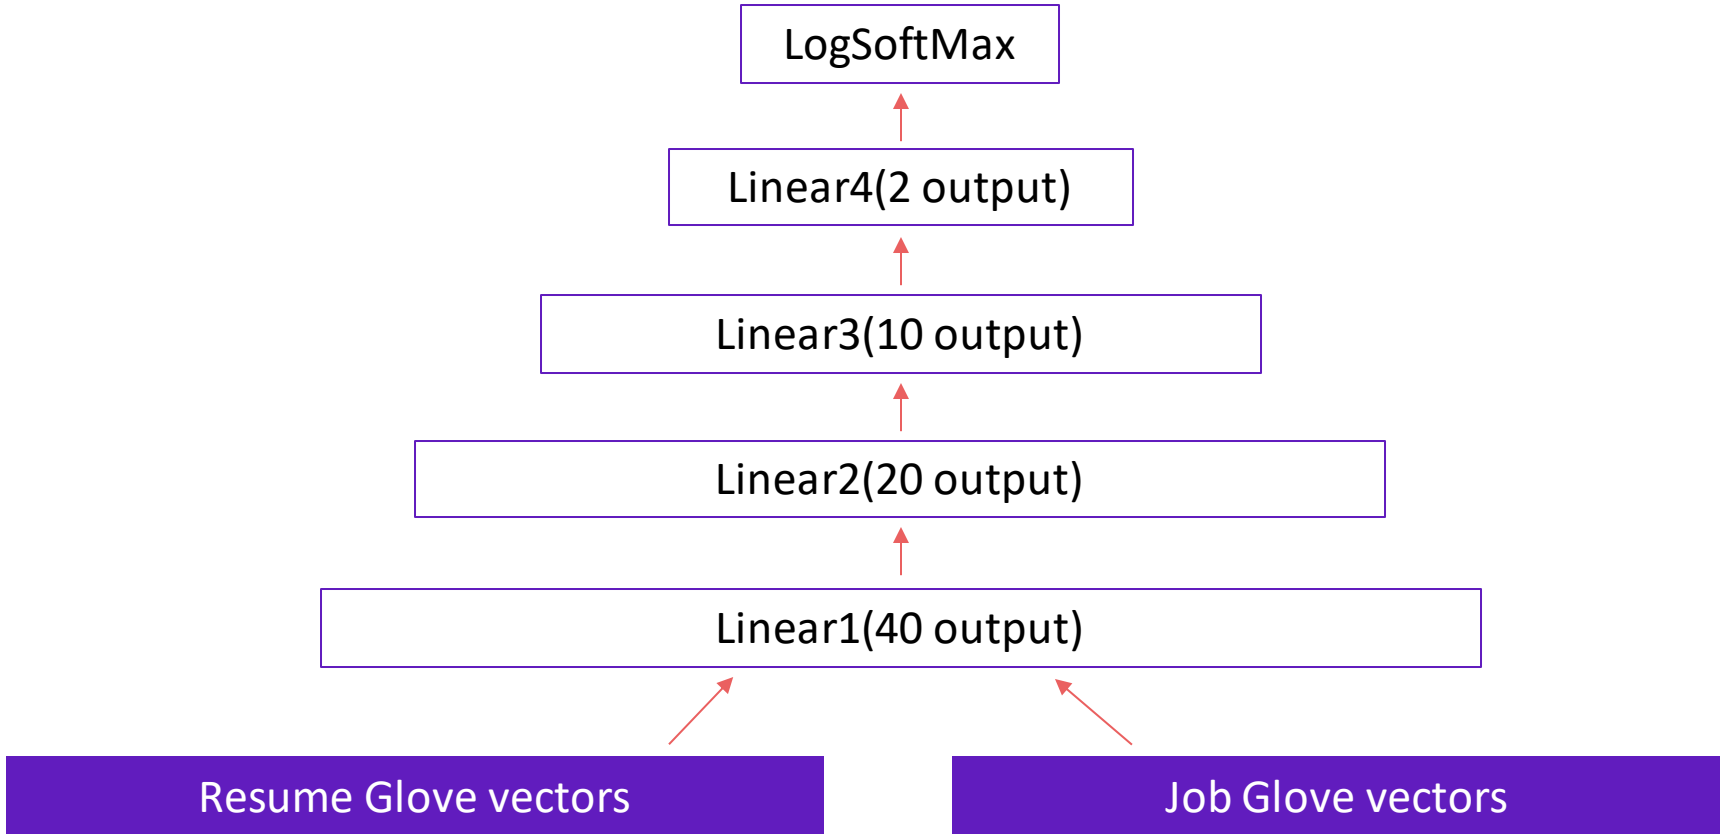

End to End Flow

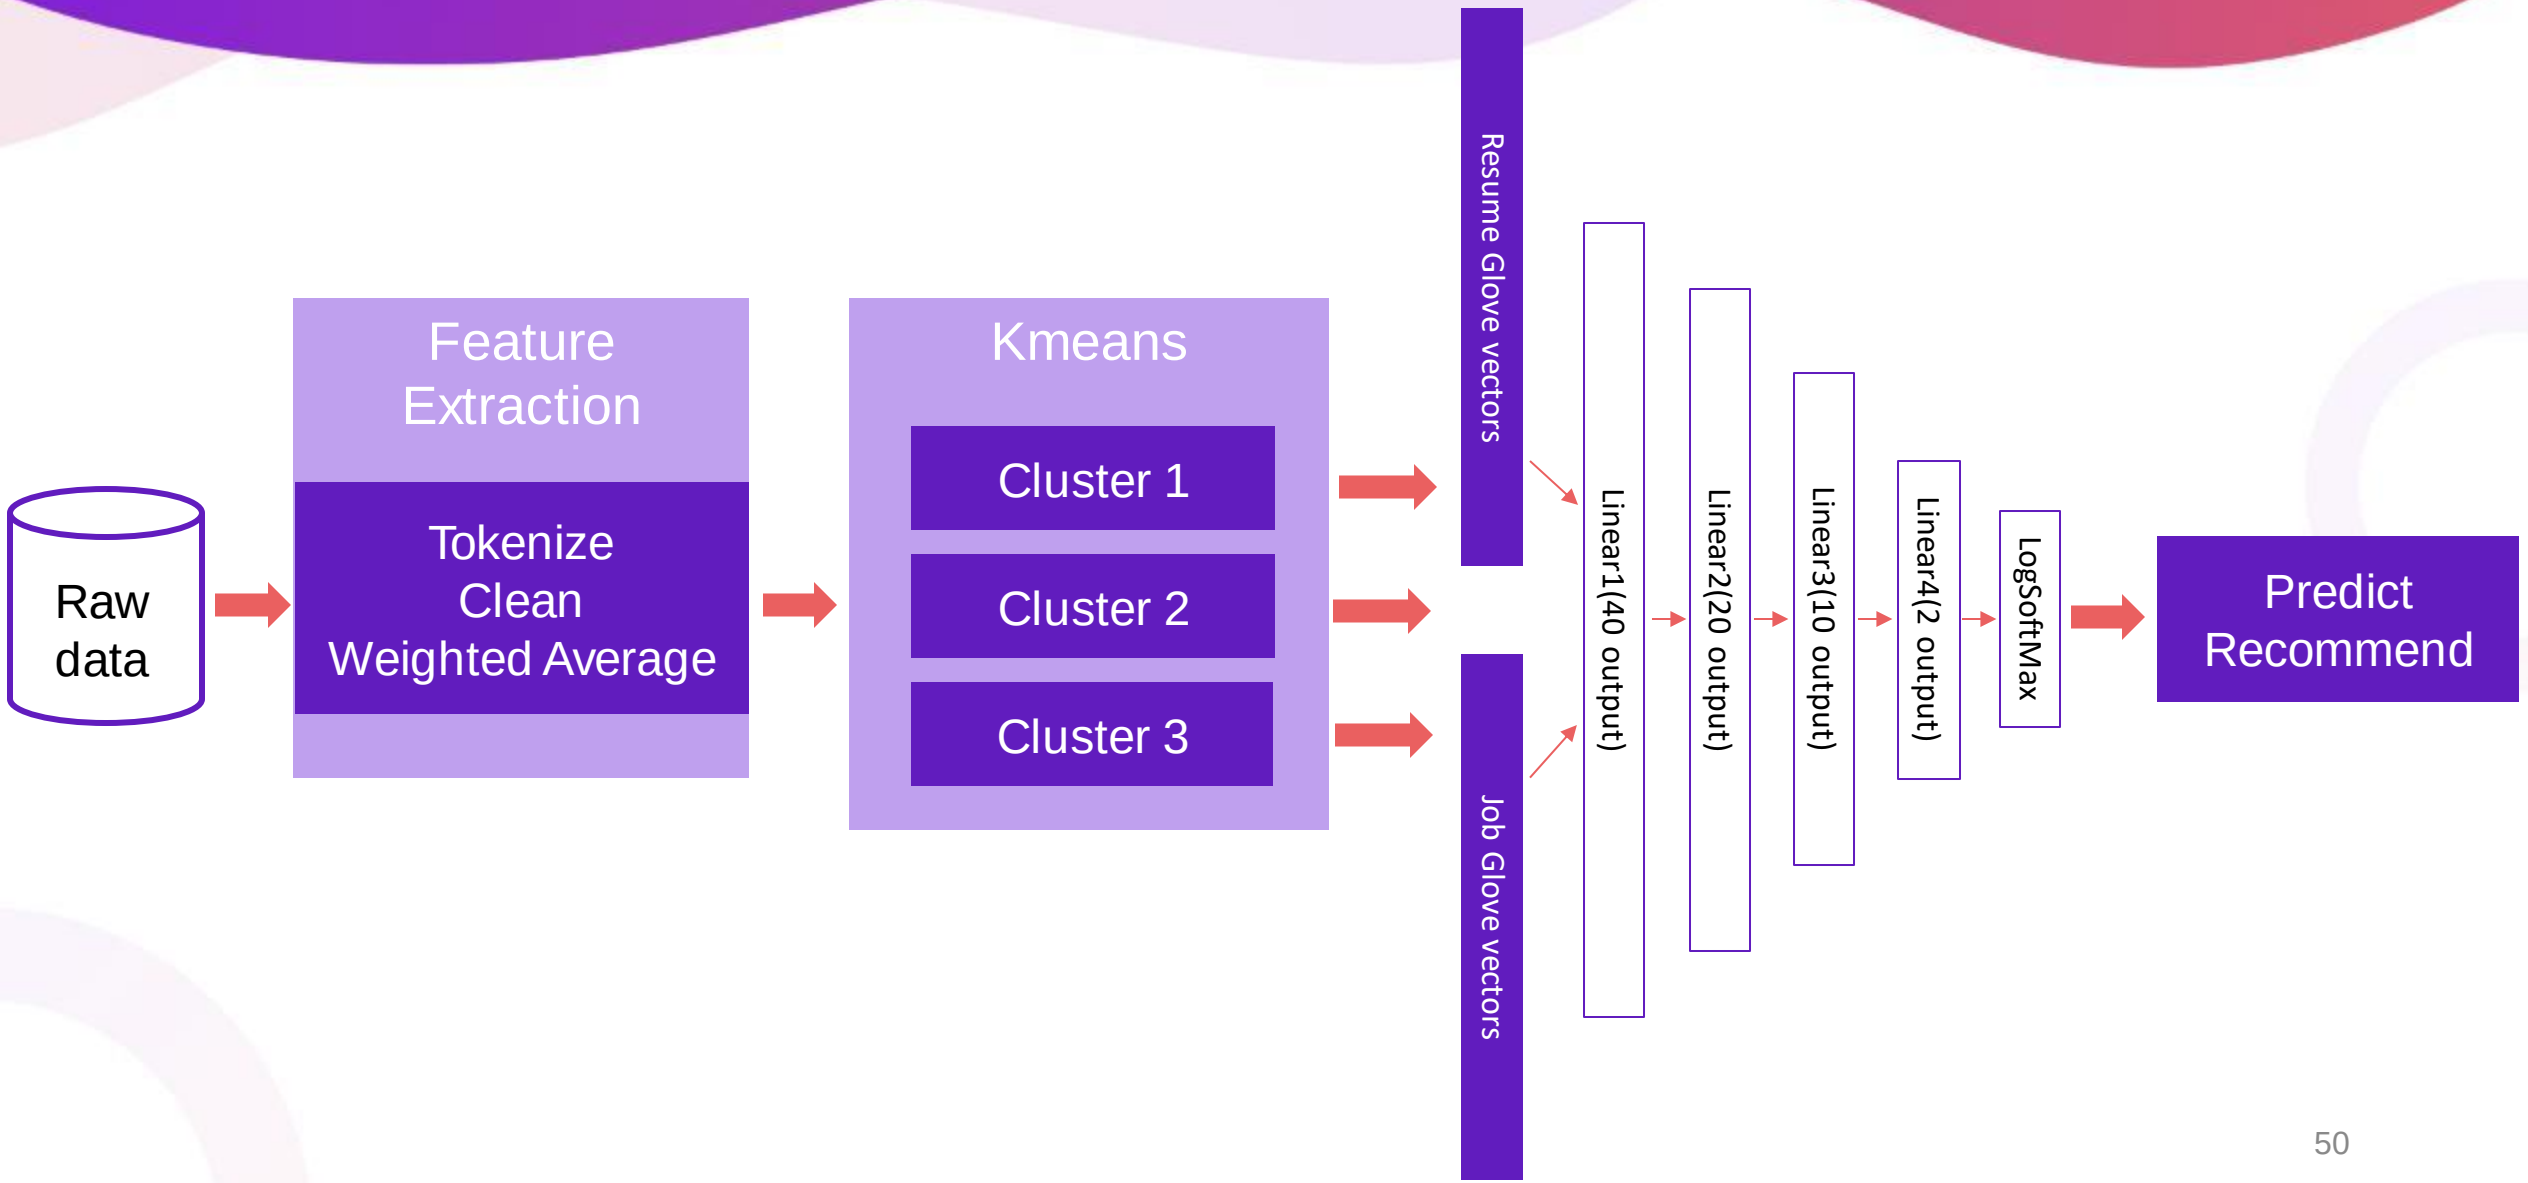

Evaluation Results

Talroo™

Precision

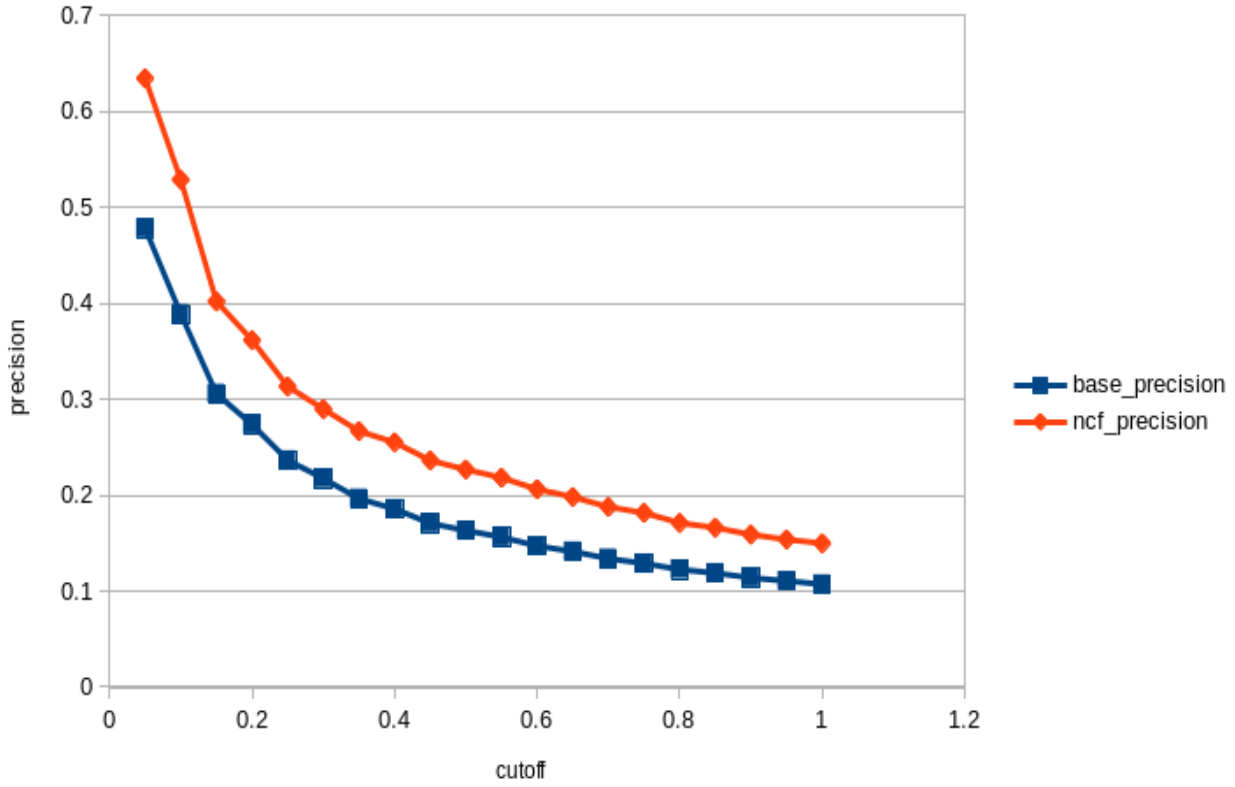

MRR

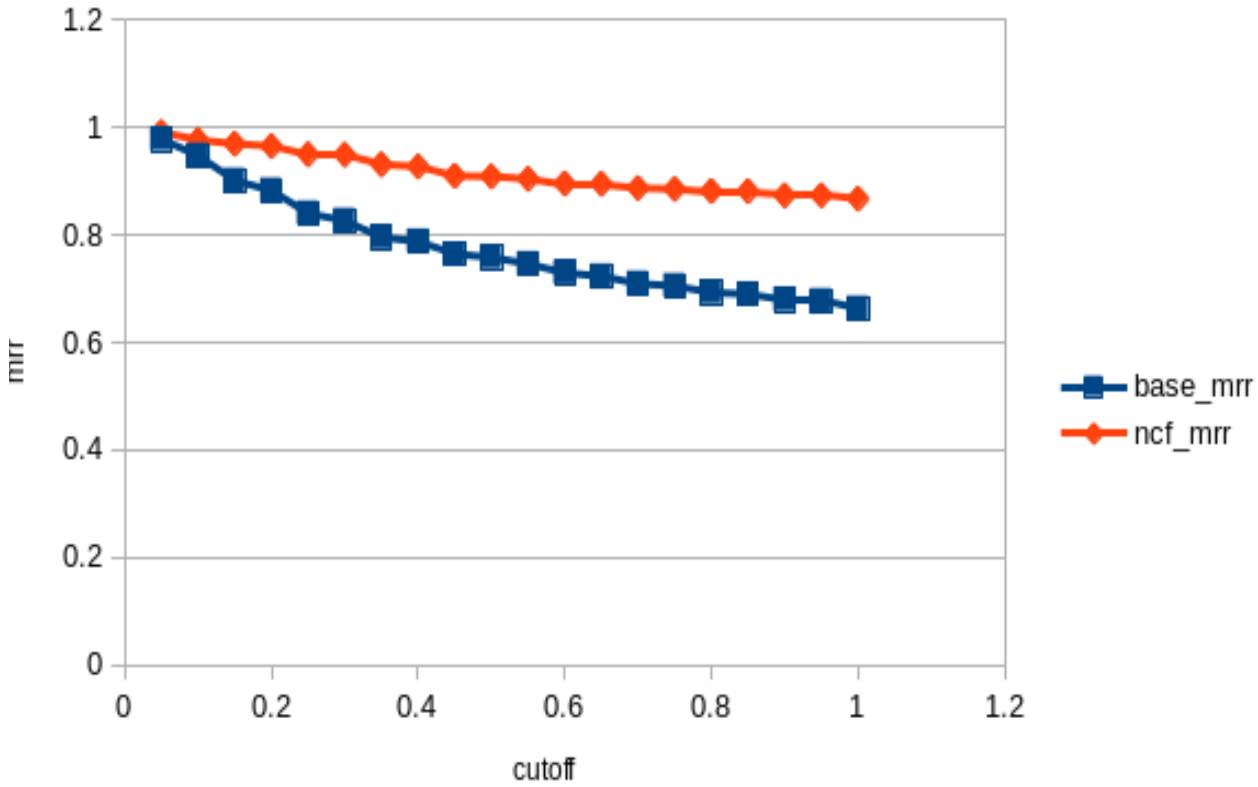

Agenda

- Talroo Overview
- Job Search Challenges
- Resume Search Opportunity
- Analytics Zoo and BigDL Overview
- Resume Search Analytics Zoo Solution
- Lessons Learned

Lessons learned

- Analytics Zoo/BigDL integrates well into existing AWS Databricks Spark ETL and machine learning platform
- Analytics Zoo/BigDL scales with our data and business
- Jobs and resumes can be effectively modeled and processed through embeddings
- Ensembling multiple models and glove embedding feature embedding proved to be very effective for rich content
- More information available at https://analytics-zoo.github.io/

Legal Disclaimer

- Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.
- No computer system can be absolutely secure.
- Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit **http://www.intel.com/performance**.

Questions?

Wenjing Zhan
Data Scientist, Talroo
wenjing@talroo.com

Guoqiong Song
Software Engineer, Intel
guoqiong.song@intel.com

Jacob Eisinger
Director of Data, Talroo
jacob@talroo.com