

# AI: The LinkedIn Way



**Deepak Agarwal**

VP of LinkedIn AI

September 19, 2018

Let people know you're open

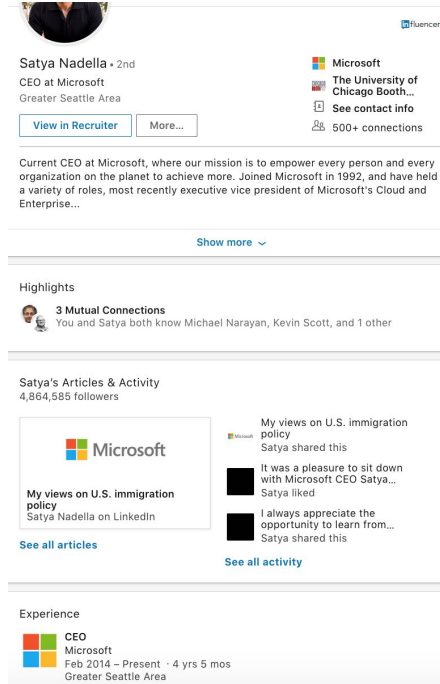


**We are all seeking new professional opportunities**

To Advance our Careers

# LinkedIn operates the largest professional network on the Internet

Tell your  
Story



Satya Nadella • 2nd  
CEO at Microsoft  
Greater Seattle Area

Microsoft  
The University of Chicago Booth...  
See contact info  
500+ connections

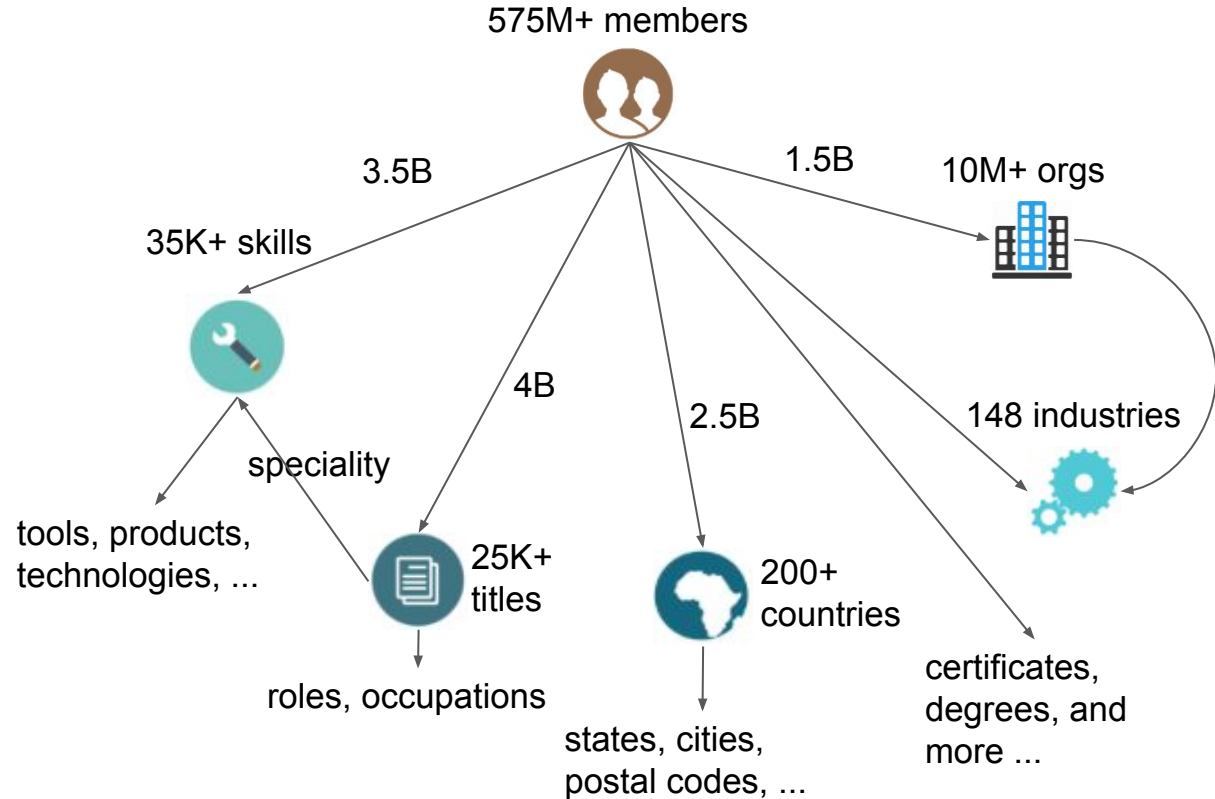
Current CEO at Microsoft, where our mission is to empower every person and every organization on the planet to achieve more. Joined Microsoft in 1992, and have held a variety of roles, most recently executive vice president of Microsoft's Cloud and Enterprise...

Highlights  
3 Mutual Connections  
You and Satya both know Michael Narayan, Kevin Scott, and 1 other

Satya's Articles & Activity  
4,864,585 followers

My views on U.S. immigration policy  
Satya shared this  
It was a pleasure to sit down with Microsoft CEO Satya...  
Satya liked  
I always appreciate the opportunity to learn from...  
Satya shared this

Experience  
CEO  
Microsoft  
Feb 2014 – Present · 4 yrs 5 mos  
Greater Seattle Area



# Artificial Intelligence at LinkedIn

Every professional interaction with LinkedIn is personal, relevant and helps in providing opportunity



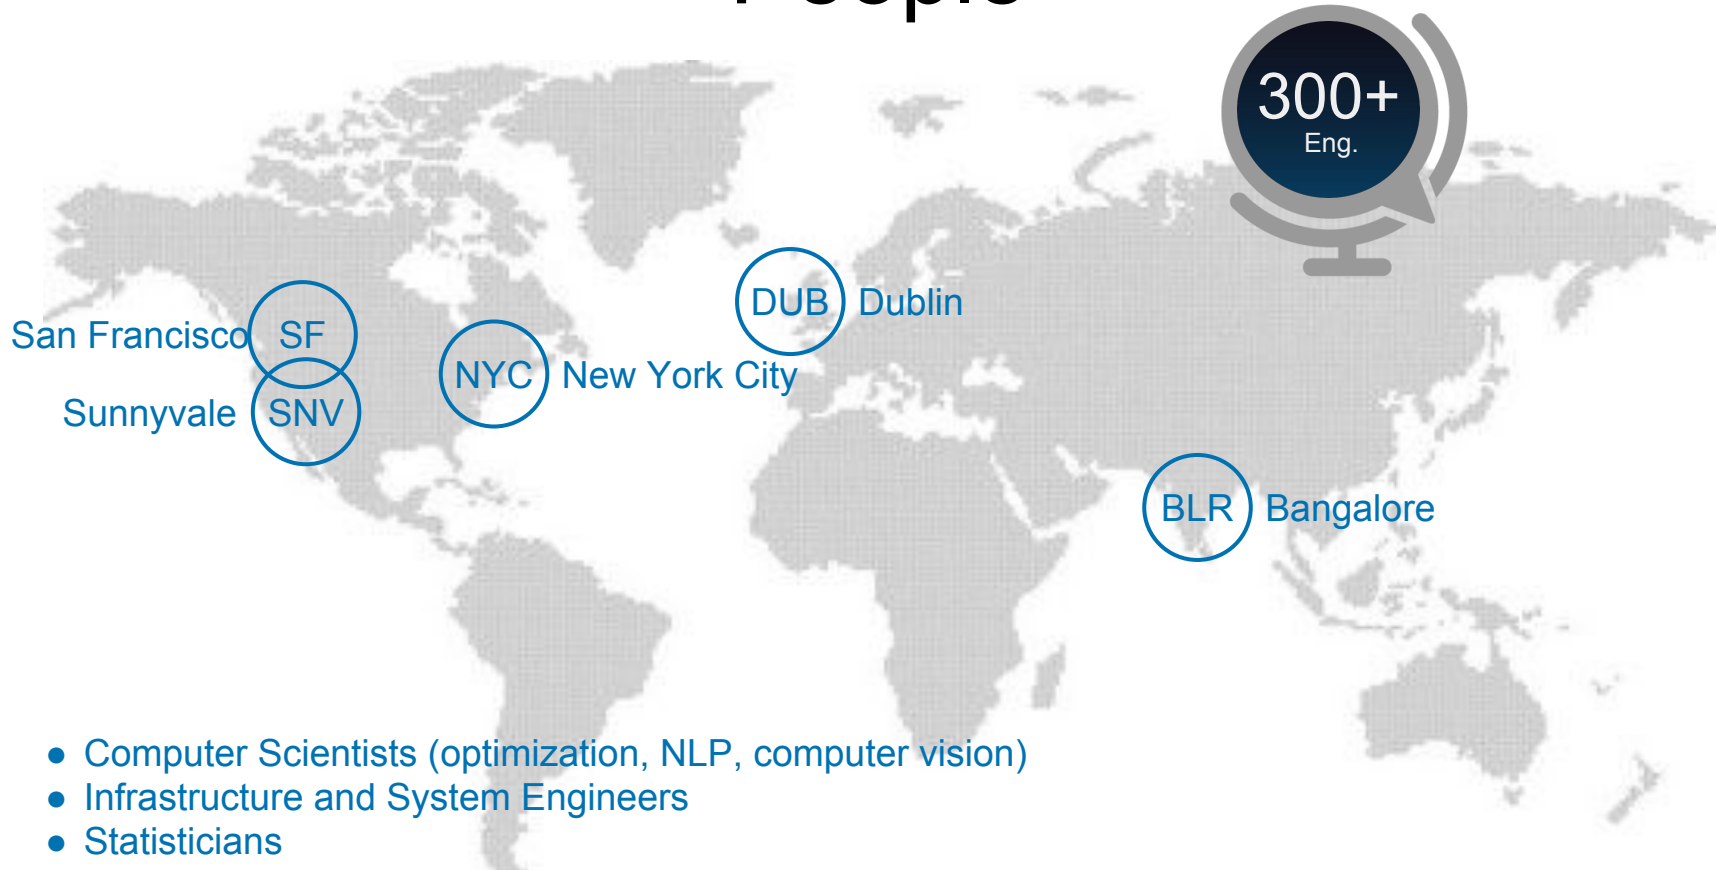
Automatically and optimally deliver the right information to the right user at the right time through the right channel



Accelerator of LinkedIn value propositions: engagement on the platform, customer ROI



# People



- Computer Scientists (optimization, NLP, computer vision)
- Infrastructure and System Engineers
- Statisticians

## People You May Know

12:30

in People, jobs, posts and more...


670 Connections


5 of your batchmates from Indian Institute of Management (2010-2012)

Suresh Kiran Venk... Zishan Nikita

VIEW ALL CONNECT (5)

People you may know

 **Soumya Yadav**  
Founder Director & Practice Head at Ele...  
27 mutual connections  
[CONNECT](#)

 **Kirthi Venkatesh**  
Creative Head at Tree Design  
8 mutual connections  
[CONNECT](#)

Home Jobs Network Profile


## Feed

9:41 AM 100%

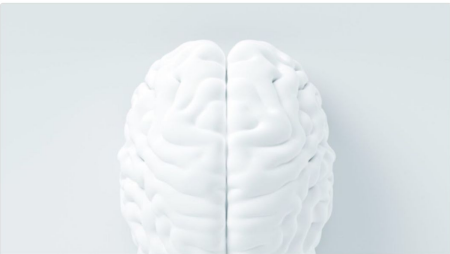
Search for people, jobs, posts and m...

Feed Featured

Share an article, photo or idea

 **Tomer Cohen**  
Head of Content, Search & Discovery Products  
1 h

Congrats [Lior Ron](#) on the acquisition! Drinks are on you next time we meet.



**Artificial intelligence is more artificial than intelligent**

wired.com Artificial Intelligence Google

1.2k Likes · 117 Comments

Like Comment Share

Home My Network Messaging Notifications Jobs

## Jobs



### Principal Data Scientist

Microsoft · Bangalore, IN

Posted 3 days ago · 1,142 views



Save

Apply

#### Job

- 147 applicants
- Full-time

#### Company

- 10001 employees
- Computer Software

#### Connections

- 32 connections
- 100 company alumni



**32 connections** can refer you

Get referred to **increase your chances** of landing an interview.

[Ask for a referral](#)

#### Job description

We are currently building the next generation of cloud services for partner monetization, user acquisition, engagement & membership platform. These services have a huge global footprint of over 240 markets and process millions of transactions daily, with loads growing linearly as Microsoft moves to a "cloud first", "mobile first" strategy. The platform powers all of Microsoft's key services - Windows App Store, Windows Phone, XBOX, Bing Ads, Office 365, Microsoft Azure to name just a few. This endeavor offers big opportunities for data science and machine learning .

**Seniority Level**  
Not Applicable

**Industry**  
Computer Hardware ,  
Computer Software ,  
Information Technology and Services

## Learning

**PREMIUM**

Deepak, ace your Salt interview - learn Machine Learning skills

**Building and Deploying Deep Learning Applications with TensorFlow**  
2,636 viewers

[See my courses](#)

#### Courses related to Deepak's skills



**Data Science Foundations: Python Scientific Stack**  
Viewers: 6,033



**Pandas for Data Science**  
Viewers: 10,197



**Building and Deploying Deep Learning Applications with TensorFlow**  
Viewers: 2,636

[See more courses](#)

# Recruiter Search

The screenshot shows the LinkedIn Recruiter interface. At the top, there are navigation tabs: PROJECTS, CLIPBOARD, JOBS, REPORTS, and MORE. On the right, there are icons for messages, notifications, and a profile. The main header displays statistics: 9K total candidates, 463 have company connections, 201 engaged with your talent brand, and 27 past applicants. A search bar contains 'Project Manager' and a location filter for 'Greater Chicago Area'. A dropdown menu is open, showing suggestions for skills: Trending skills for this title, Project plans, Business strategy, Management, Microsoft Project, and PMP. Below the search bar, there are 'Updates' sections: 'Follow candidates to get updated on their profile changes. Some people you may want to follow', 'Your saved search, Project Managers, has 6 new results!', 'Salva Larin is having a job anniversary: 5 years as Independent HTML/CSS coder at oDesk. Send Message', and 'Britny Paulson is having a job anniversary: 5 years as User Experience Designer at Google. Send Message'. On the right side, there are sections for 'Start searching for candidates and sending messages to receive recommendations on people you may want to hire.', 'Project activity Software Engineer: 0 profiles • 0 contacted • 0 reviewed', 'Architect: 78 profiles • 4 contacted • 2 reviewed', and 'UX Designer: 78 profiles • 4 contacted • 15 reviewed'.

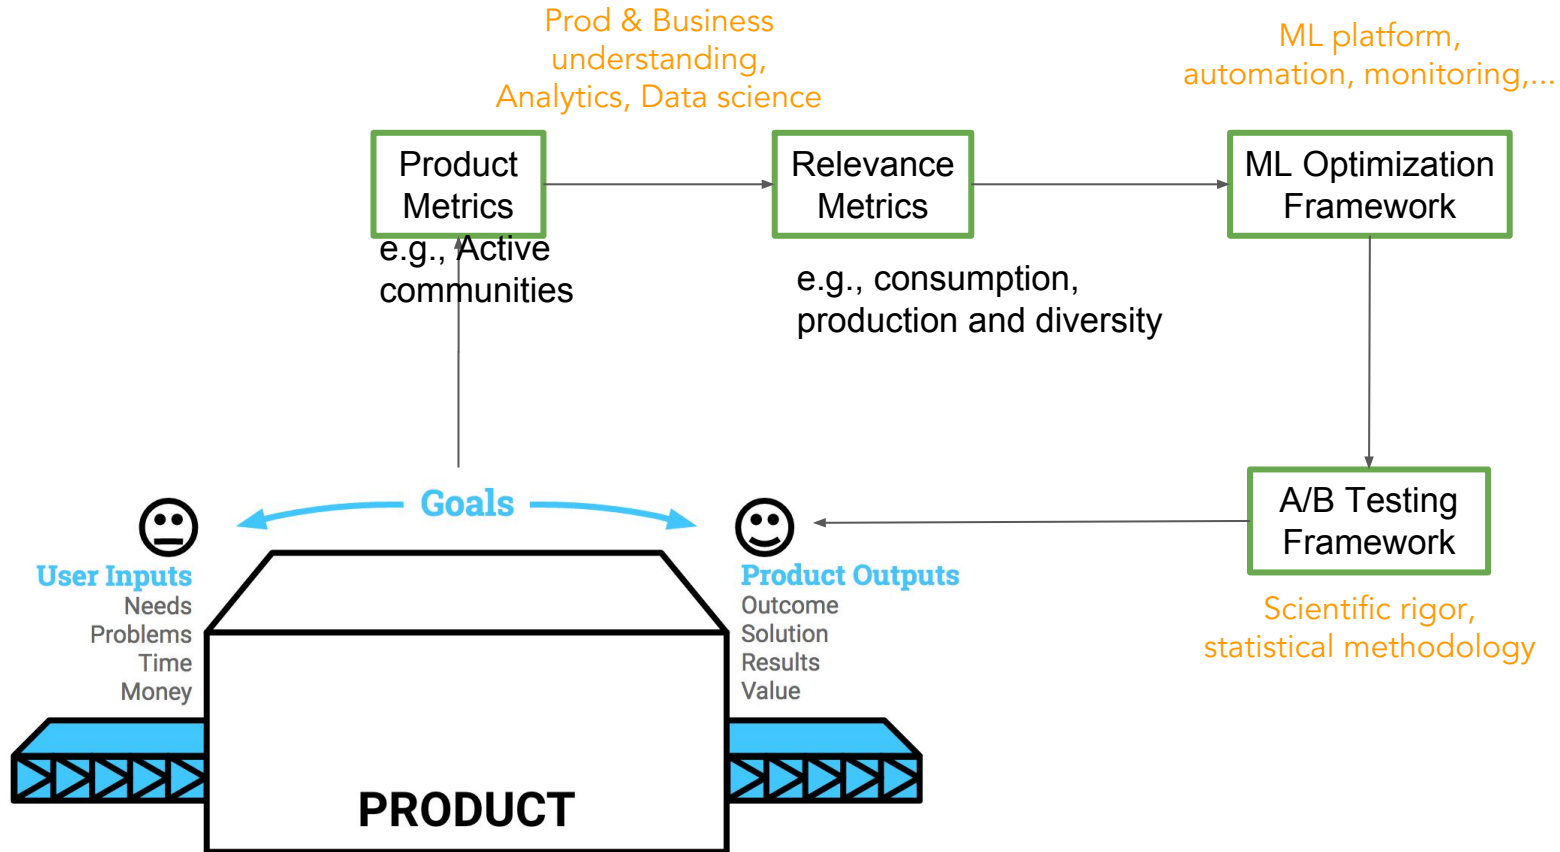
# Sales

The screenshot shows the LinkedIn Sales Navigator interface. At the top, there are navigation tabs: HOME, LISTS, DISCOVER, and REPORTS. On the right, there are icons for messages, notifications, and a profile. The main header displays a search bar with 'Search for leads or accounts' and a 'Saved Search' button. Below the search bar, there are filters for 'Current quarter' and '8 deals - \$950K'. A table of deals is displayed with columns: Account, Deal Name, Buyer circle, Amount, Close date, Stage, Probability, Next steps, and a menu icon. The table contains the following data:

Account	Deal Name	Buyer circle	Amount	Close date	Stage	Probability	Next steps
Fixdex	Fixdex opportunity	[Buyer Circle]	\$85,000	07-03-2017	4. Legal	90%	Met with the legal team to finalize the paperwork.
Trustway 70k	Trustway 70k	[Buyer Circle]	\$70,000	09-07-2017	3. Proposal	75%	Set up meeting with Alice Craig and Jon T...
Zoomjax acc	Zoomjax acc	[Buyer Circle]	\$10,000	07-27-2017	3. Proposal	80%	The account was very impressed with the pres...
Techcore new bus	Techcore new bus	[Buyer Circle]	\$5,000	06-14-2017	1. Discovery	—	Jacob seems really interested, looking for mo...
Fixdex opportunity	Fixdex opportunity	[Buyer Circle]	\$10,000	02-28-2017	2. Lead generation	50%	Finding solid leads. Will update once I make mo...
W&C opp	W&C opp	[Buyer Circle]	\$50,000	12-04-2017	2. Lead generation	50%	Continuing the search after a call with Arnold during our...
Ostia Account	Ostia Account	[Buyer Circle]	—	07-29-2017	0. Qualification	—	Working with SD to hone in on this one.
Lamcane 20k	Lamcane 20k	[Buyer Circle]	\$20,000	12-10-2017	2. Lead generation	25%	Still searching for more leads



# AI/ML process



App 1: Member  
engagement

AI Platform

# Contributors on LinkedIn

- John contributes by sharing



Share this with

- Email
- Facebook
- Messenger
- Twitter
- Pinterest
- LinkedIn**

## umans in poker game


beaten four human players in a  
0 days. Libratus, an artificial  
developed at Carnegie Mellon University,  
ant of the game known as no-limit

*Messaging, commenting, publishing are other ways to contribute on LinkedIn*

# Receiving Response to Contributions

in Search Home My Network Jobs Messages

**PREMIUM**




**Bee-Chung Chen**  
Principal Staff Engineer & Applied Researcher at LinkedIn

494  
Who's viewed your profile

719  
Connections  
[Manage your network](#)

**John Smith**  
1h

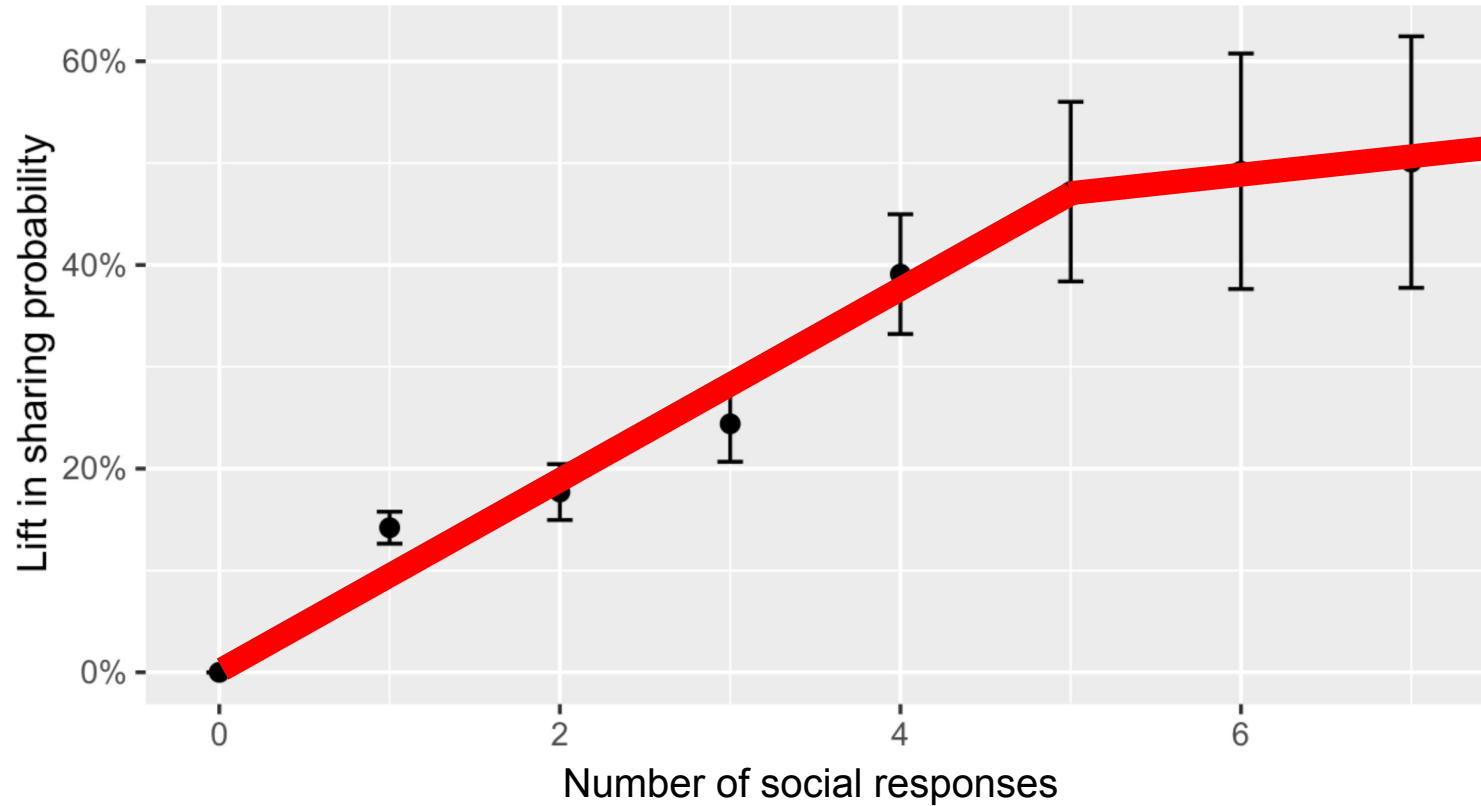
"The bot gets better and better every day. It's like a tougher version of us" ...see more



**AI program beats humans in poker game - BBC News**  
bbc.com

Like Comment Share

## Response to contributions leads to repeat contributions

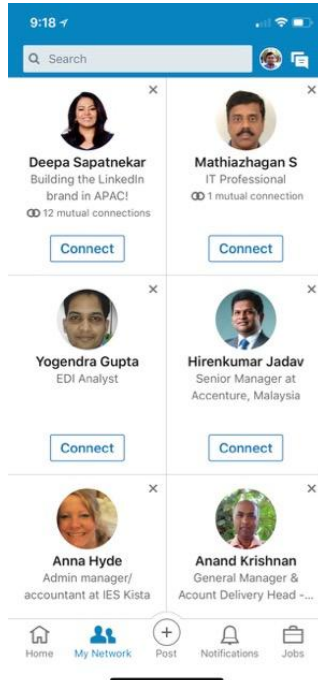


# Contributions and Daily Active Uniques

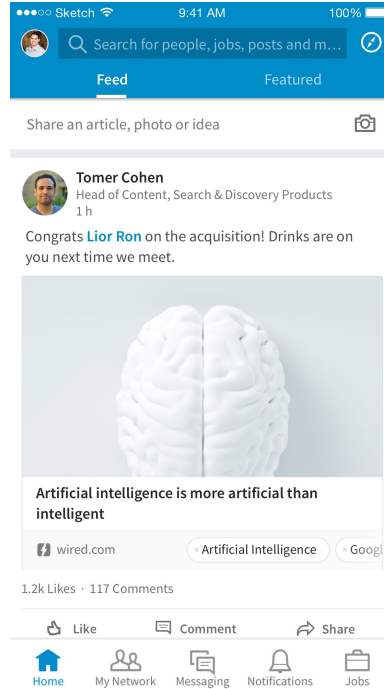
- 1 unit increase in contributors leads to 0.4-0.7 increase in Daily Active Users- Network A/B testing (Eckles et al, Journal of Causal Inference, 2016)
- **Summary:** Increase in frequent contributors improves daily active uniques and engagement

# Standard approach: Local optimization

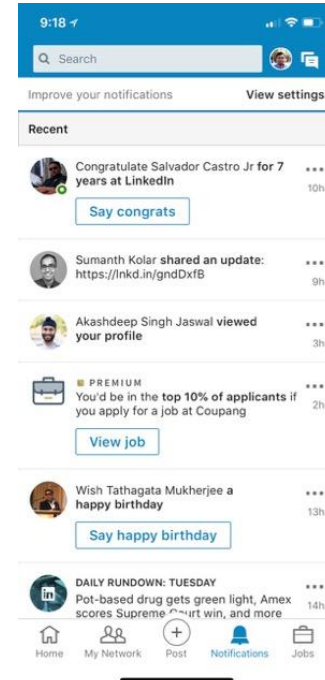
## PYMK



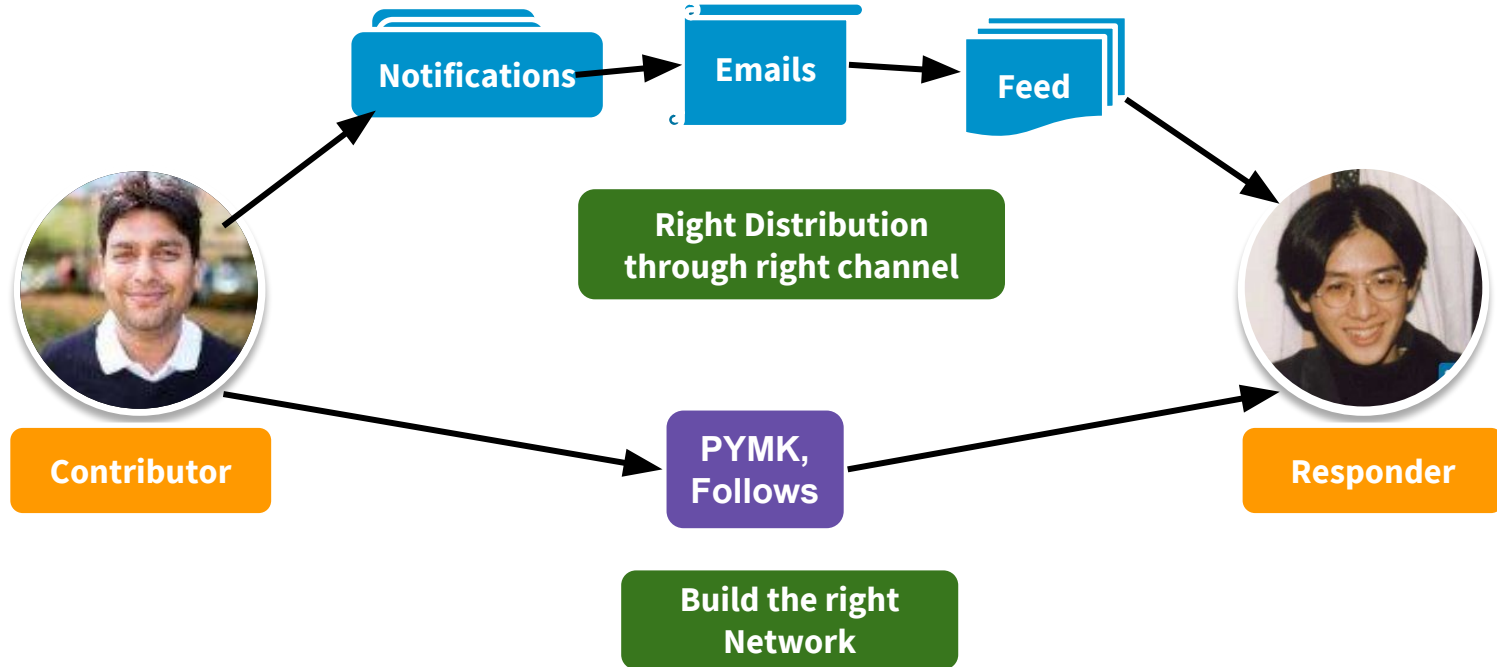
## Feed



## Notifications



# Optimize Eco-system





# Problem Formulation

- Create the right network that is value based, not volume based
  - Incorporate future downstream value in recommending new edges
- Optimize distribution of contributions to network across channels to encourage repeat contributions
- Guard against “high concentration” when distributing contributions
  - Better to have more contributors receiving some feedback instead of a few receiving a lot

# Optimizing right network (PYMK)

Build connections that are likely to drive conversations:

$$\textit{score} = P(\textit{connection}) + \alpha P(\textit{conversation})$$

*connection*: same school? same company? ...

*conversation*: messaging(private) or comment/like/share (public)

# Content Distribution

$$P(\textit{click}) + \alpha P(\textit{viral})E(\#\textit{DSI}) + \beta P(\textit{resp})\exp(-\gamma E(\#\textit{creator response}))$$

*DSI* = downstream impact

*viral* = share, like, comment

*resp* = response to the creator via comment, like, re-share, messaging

$\gamma$  = temperature

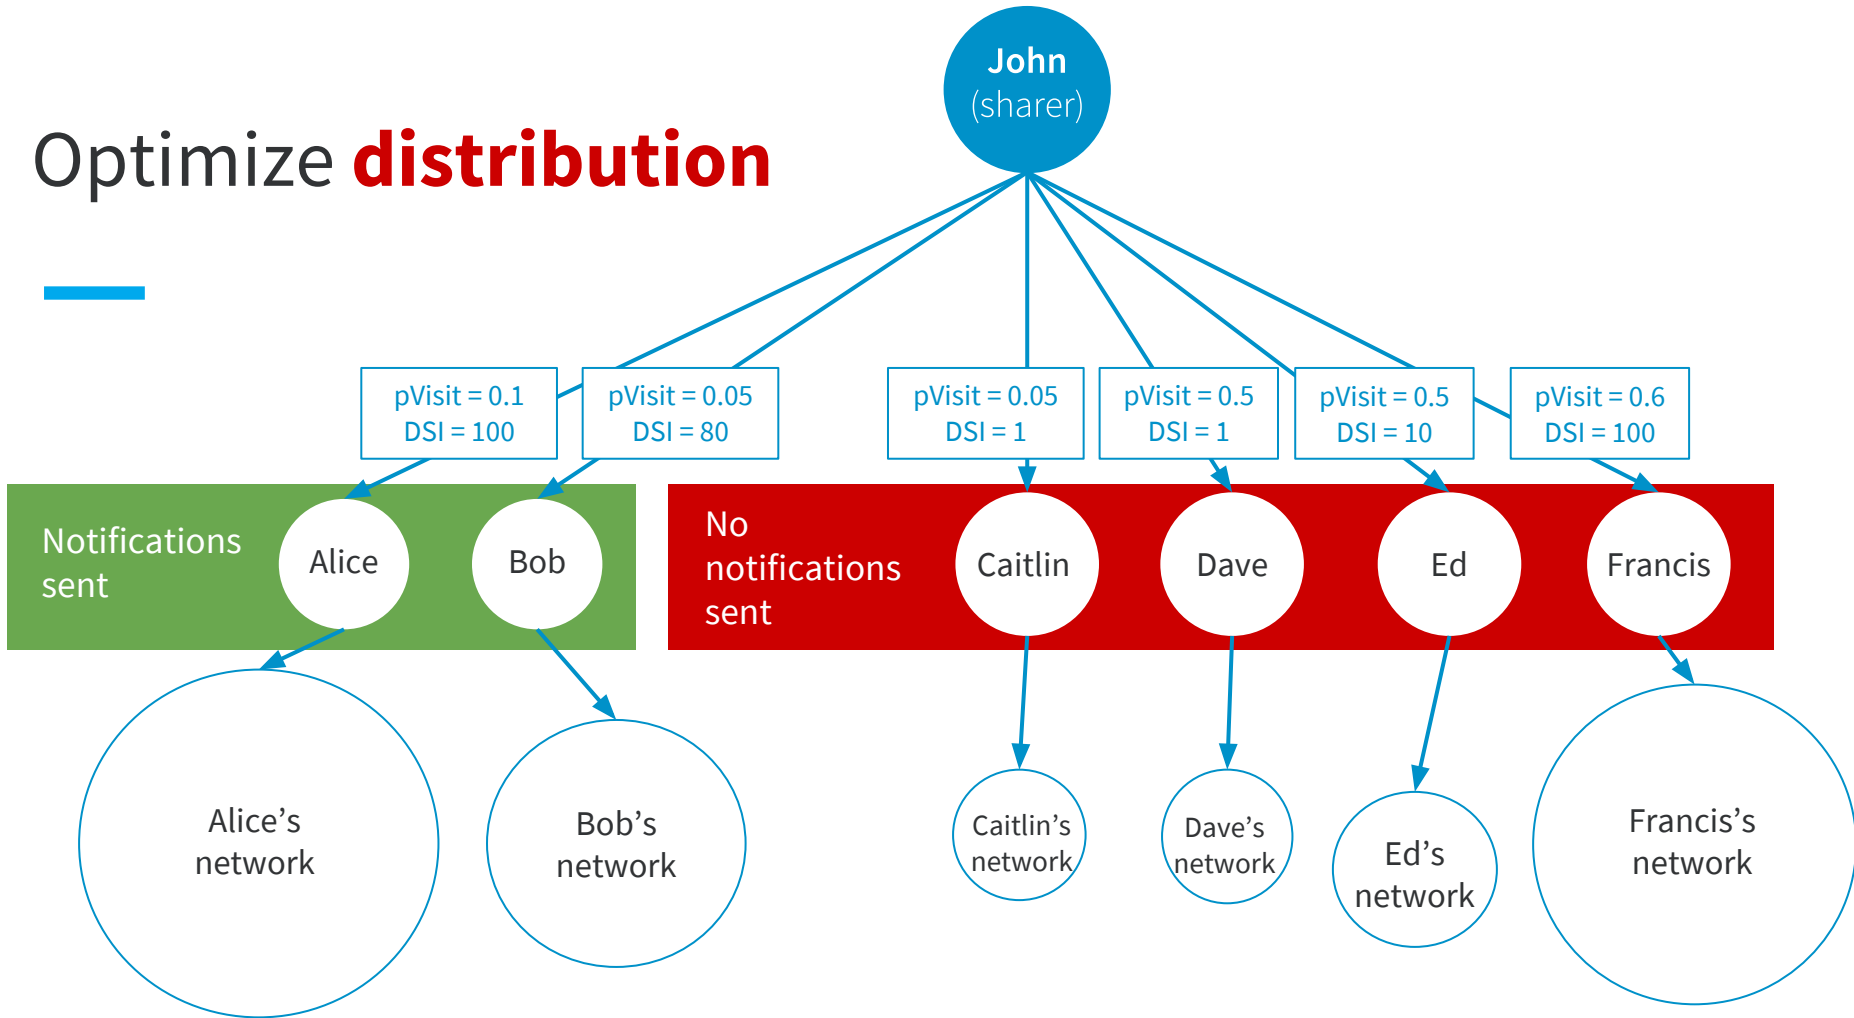
$\alpha, \beta$  are weights



Get Notified

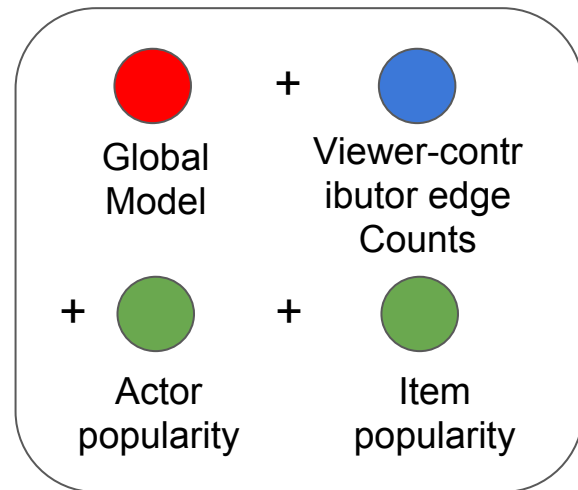
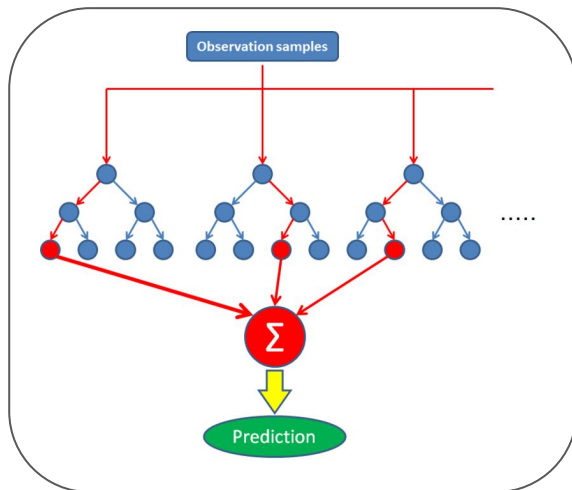

$$p\textit{Visit} = \begin{cases} \textit{small} & : \textit{notify viewer} \\ \textit{big} & : \textit{wait to show on feed} \end{cases}$$

# Optimize **distribution**

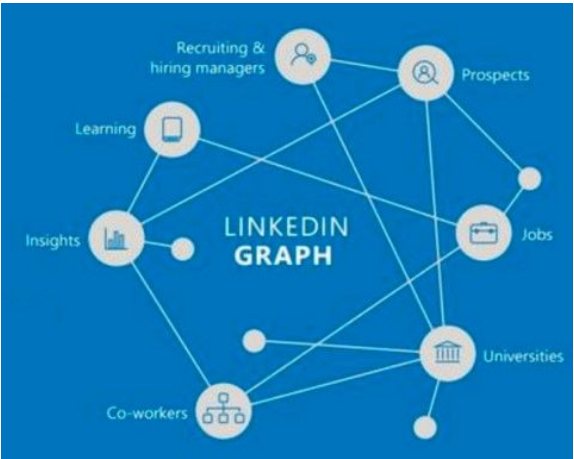
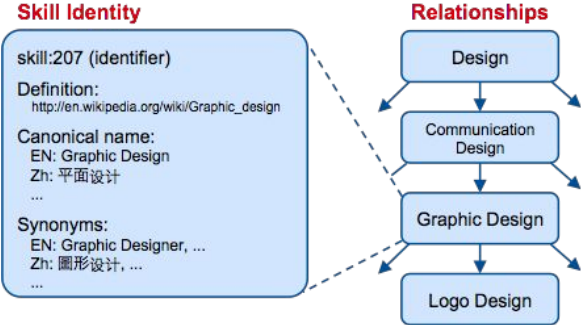


# What Kind of Model?


A combination of deep learning and trees models and real time “corrections” at fine granularities



# Features in ML model: Knowledge Graph, Learned Representation

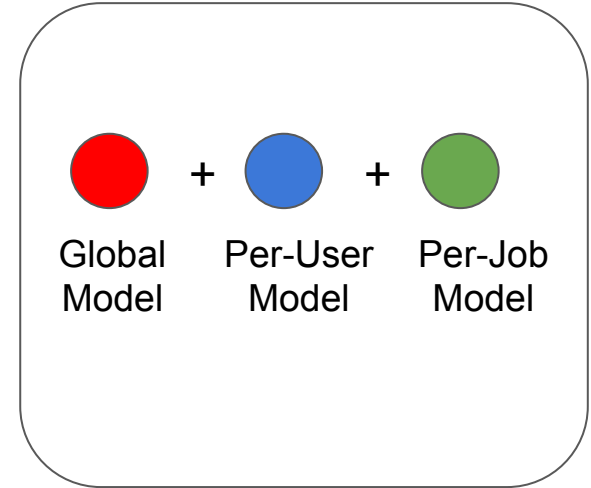
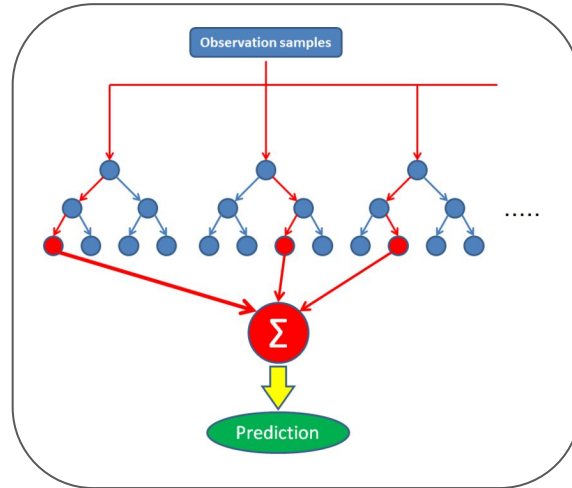



Information extraction  
 Human curation  
 Deep Learning



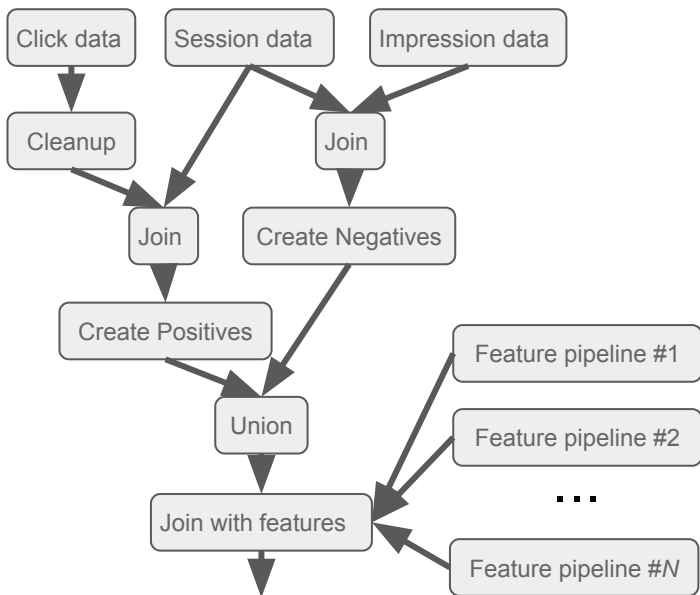
# Challenges

- Jointly train a combination of different classes of models
- Automatically determine features, model structure & hyperparameters

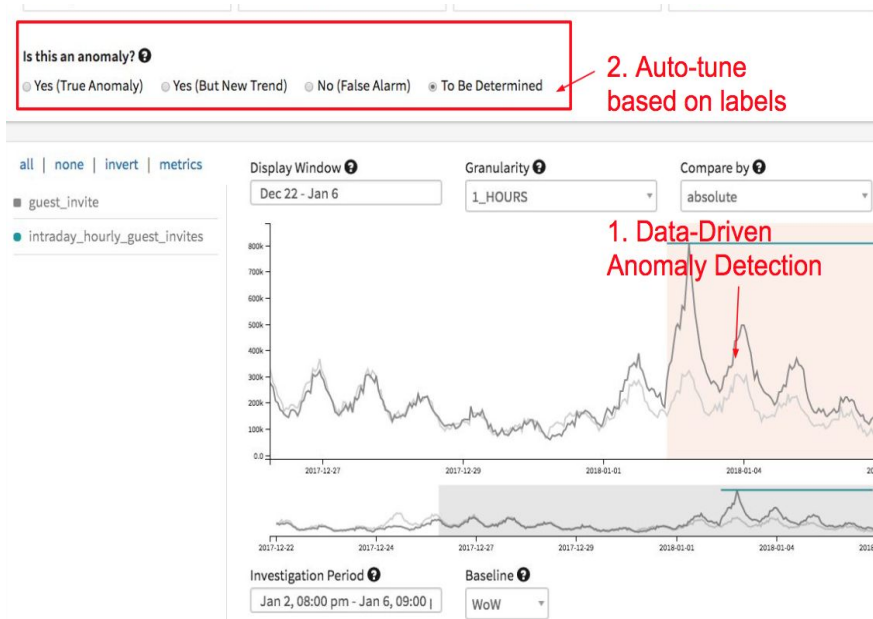


# Challenges

Manage data/modeling pipelines with complex dependencies

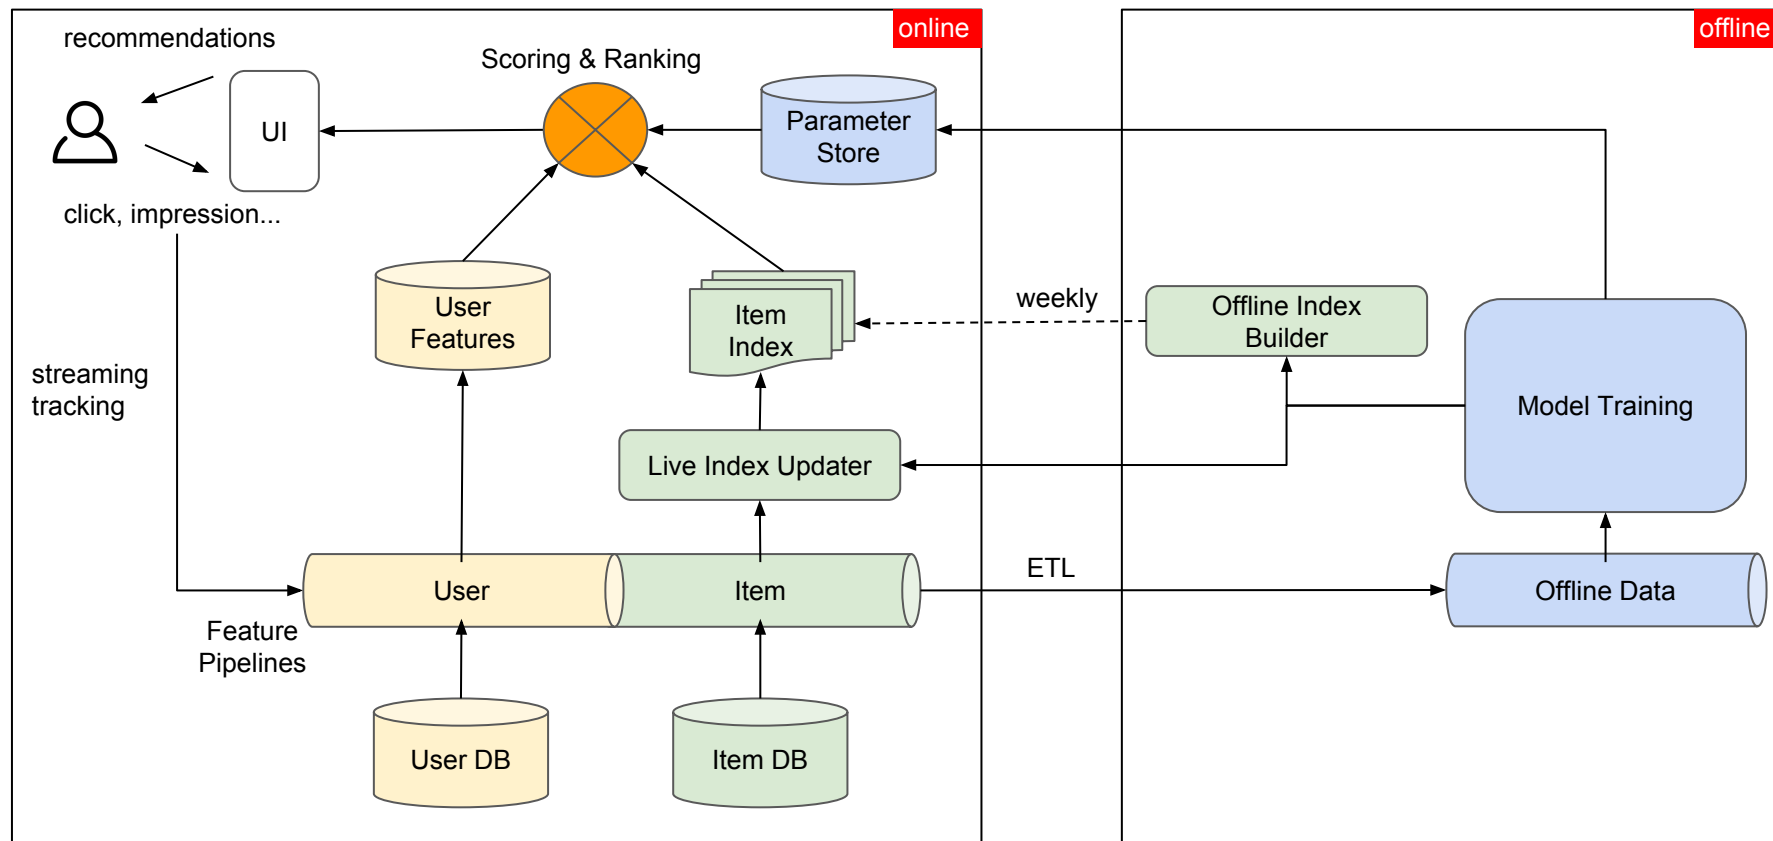


Automatically detect issues and gracefully degrade model performance





- Challenges: 1) deploy large models to different places in production  
2) manage complex model/data deployment processes





App 1: Member  
engagement

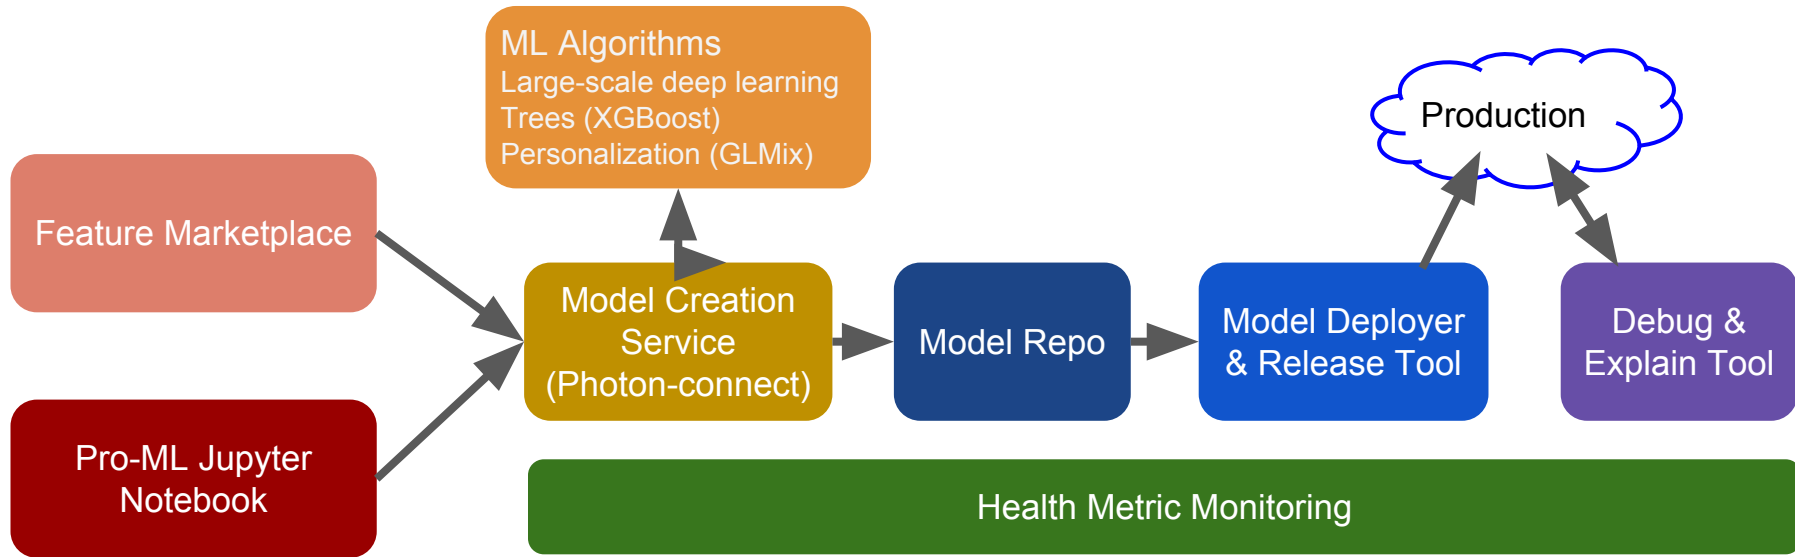
AI Platform

# Pro-ML: Productive Machine Learning

We make the end-to-end process of running and iterating on large ML workflows easy, robust and almost automated



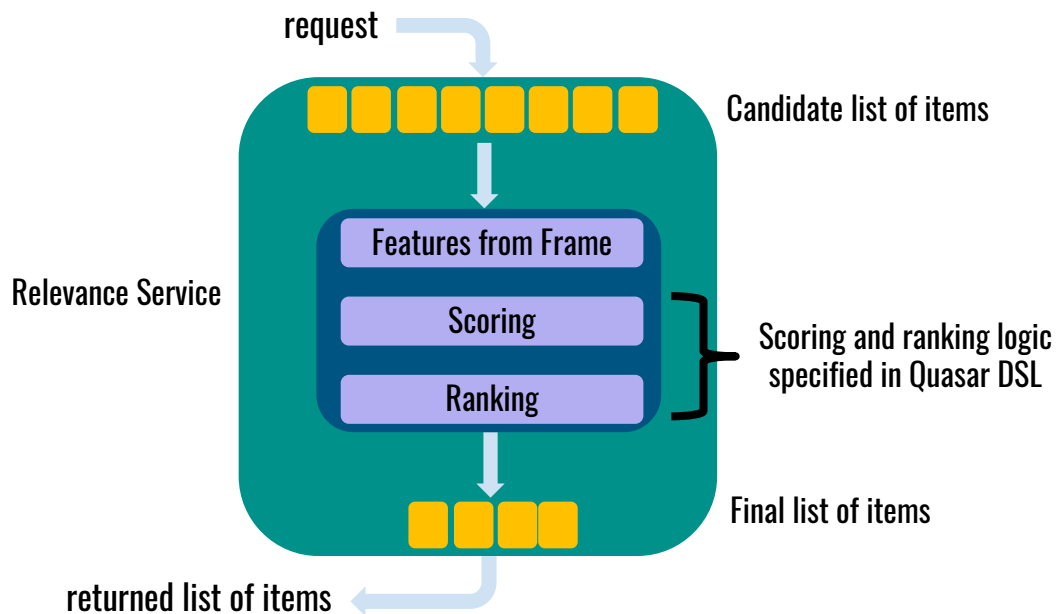
# Pro-ML Key Components





# Quasar: Unified Inference Engine @ LinkedIn

Provides a DSL to define the scoring and ranking computation logic of a model, and a high-performance execution engine to run the model



# Democratize ML at LinkedIn

---

**Focus:** Using LinkedIn tools and infra, train all engineers to use basic supervised ML in their day-to-day job

Curriculum:

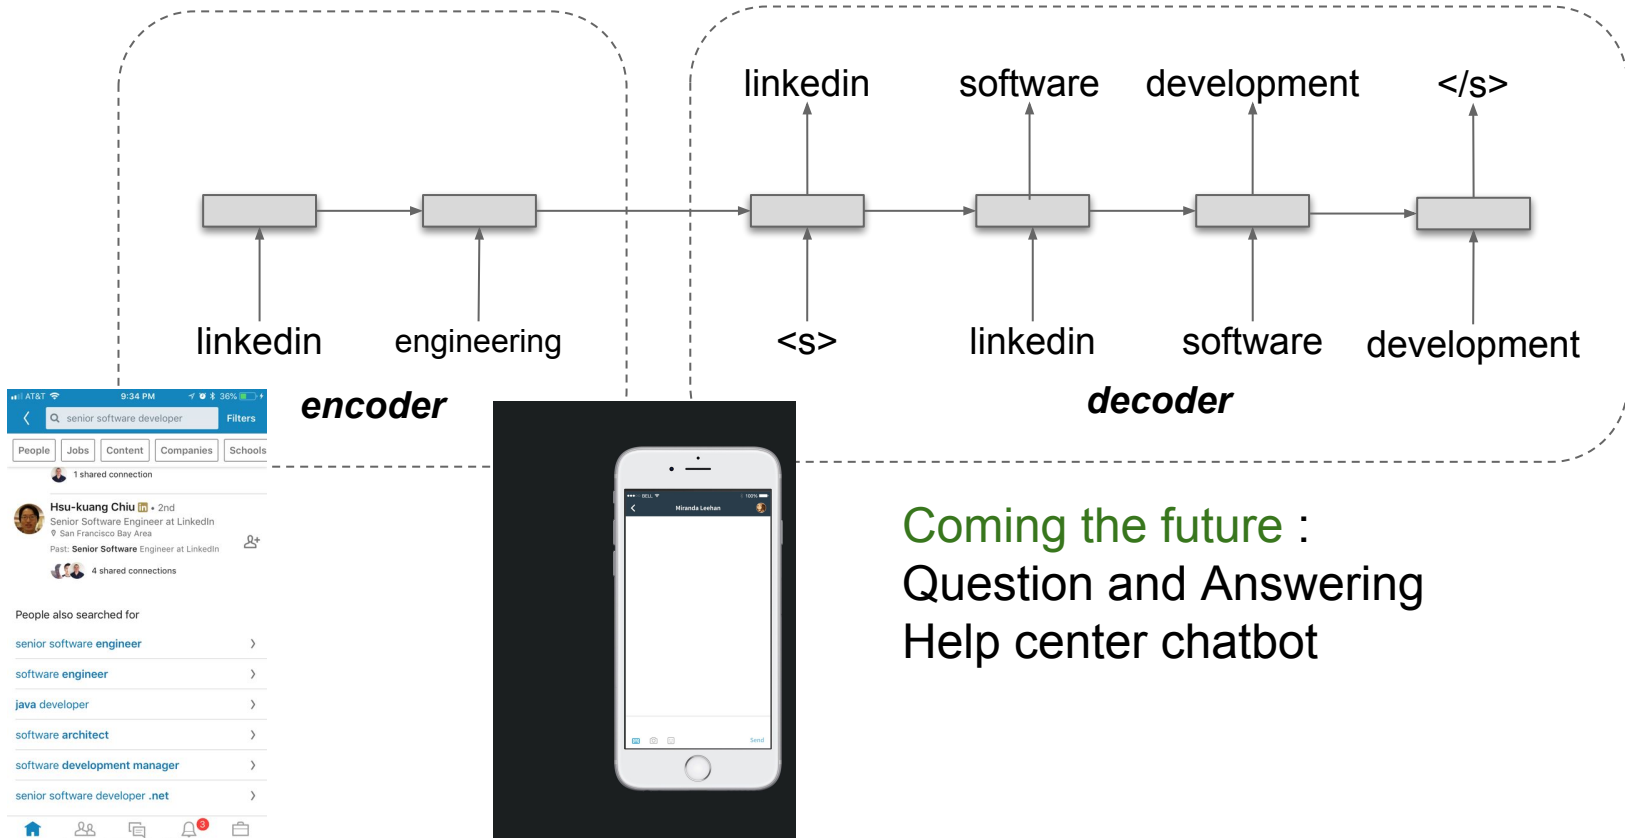
- AI Boot Camp: 5 full days (aimed at ICs)
- One-day course: Aimed at managers

Continuing education via LinkedIn Learning



**LinkedIn**  
**AI Academy**

# Conversational AI: seq2seq model.





# More Methodological Research

---

Richer class of non-parametric models for personalization

Transfer learning

Active learning to collect human labels more effectively

Large scale & real-time Graph computations

Fairness, Transparency and Privacy

Efficient network A/B testing

.....

# Summary

---

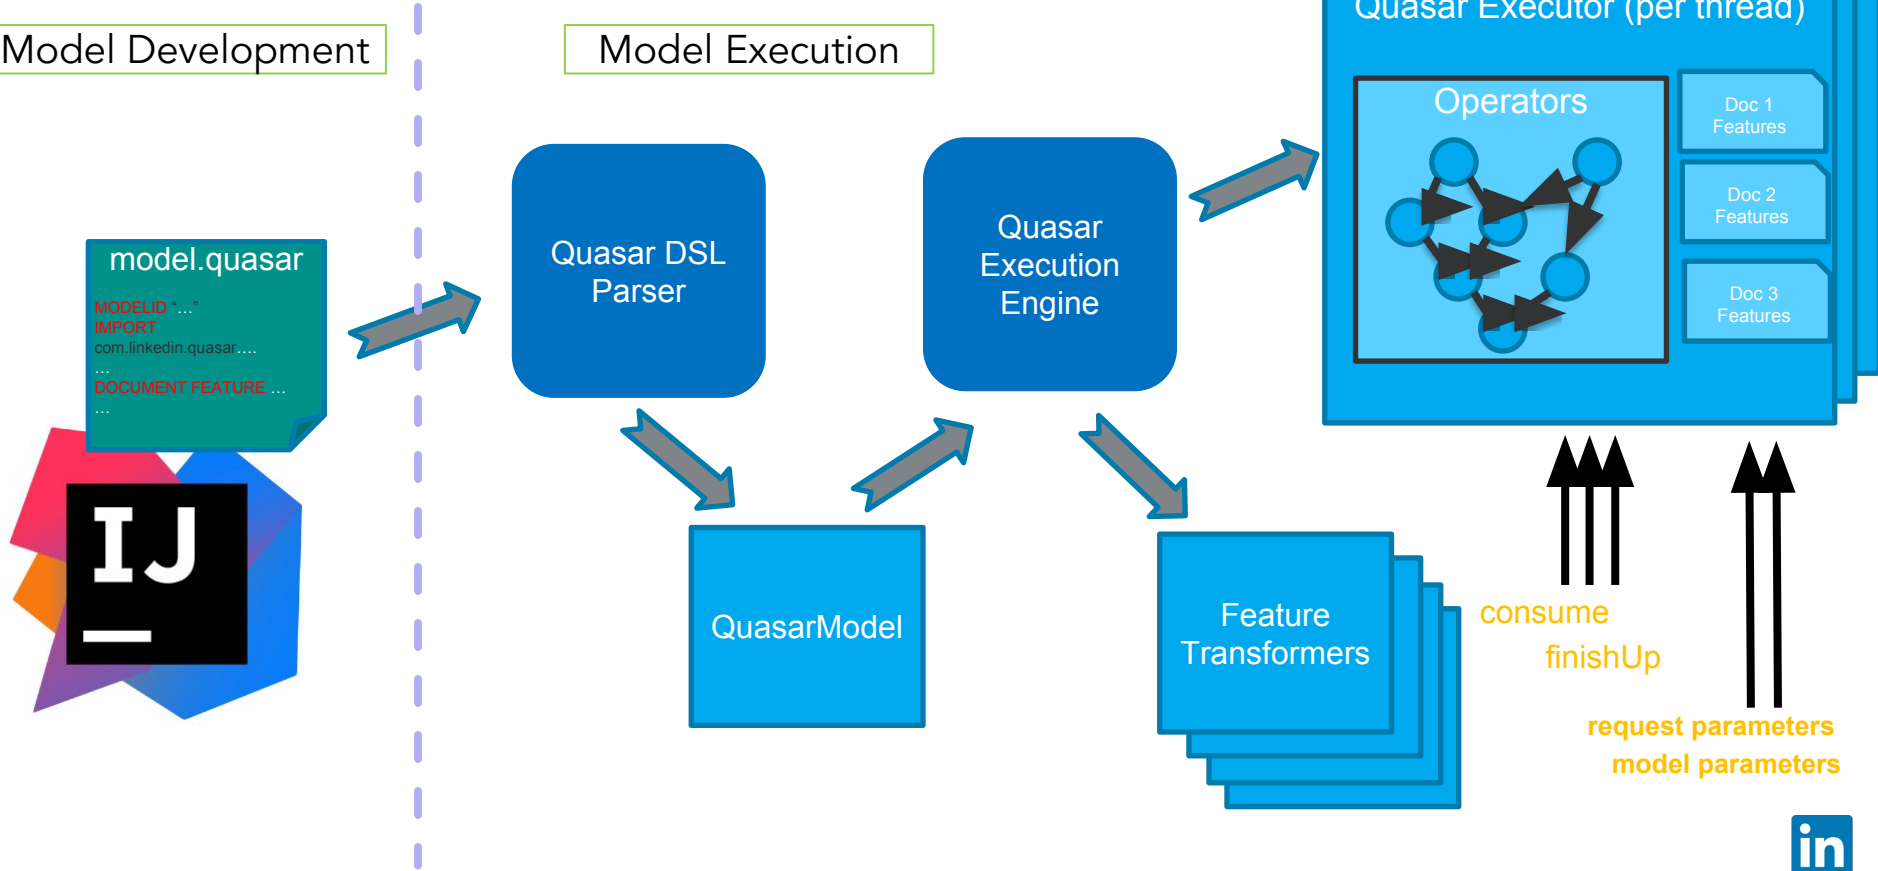
- ML/AI is “oxygen” of LinkedIn product
- Converting ill-posed business/product problem to well posed ML problem requires effort, rigor and collaboration
- Scaling ML process for large workflows challenging both from infra and organizational perspective

# Appendix

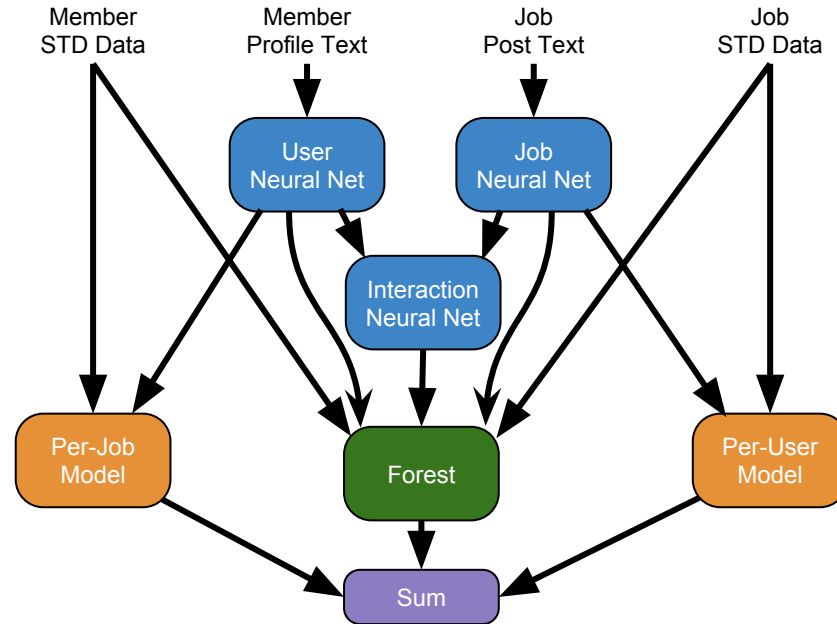
# Quasar Components

Model Development

Model Execution



Trains a DAG of different classes of ML models specified using Quasar DSL with features specified using Frame config



# Applied Research

---




1

Focused on LinkedIn  
Ecosystem - create  
member delight



2

Accelerator of Member  
Value



3

Centrally Managed but  
Deeply Aligned within  
Product Development

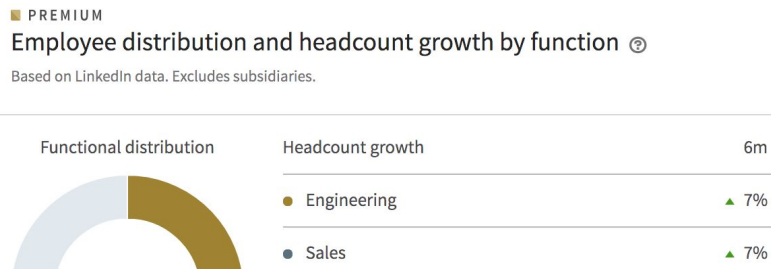
# Knowledge Marketplace

## Mission


Create a knowledge marketplace where any member and enterprise customer can ask a question in natural language and we provide an answer almost instantaneously, eliciting more information if necessary.

Hey, we found you're skilled in {#skills}id, {/skills} technologies and hence you seem to be a good match for a role at {company}

## Enrich conversations

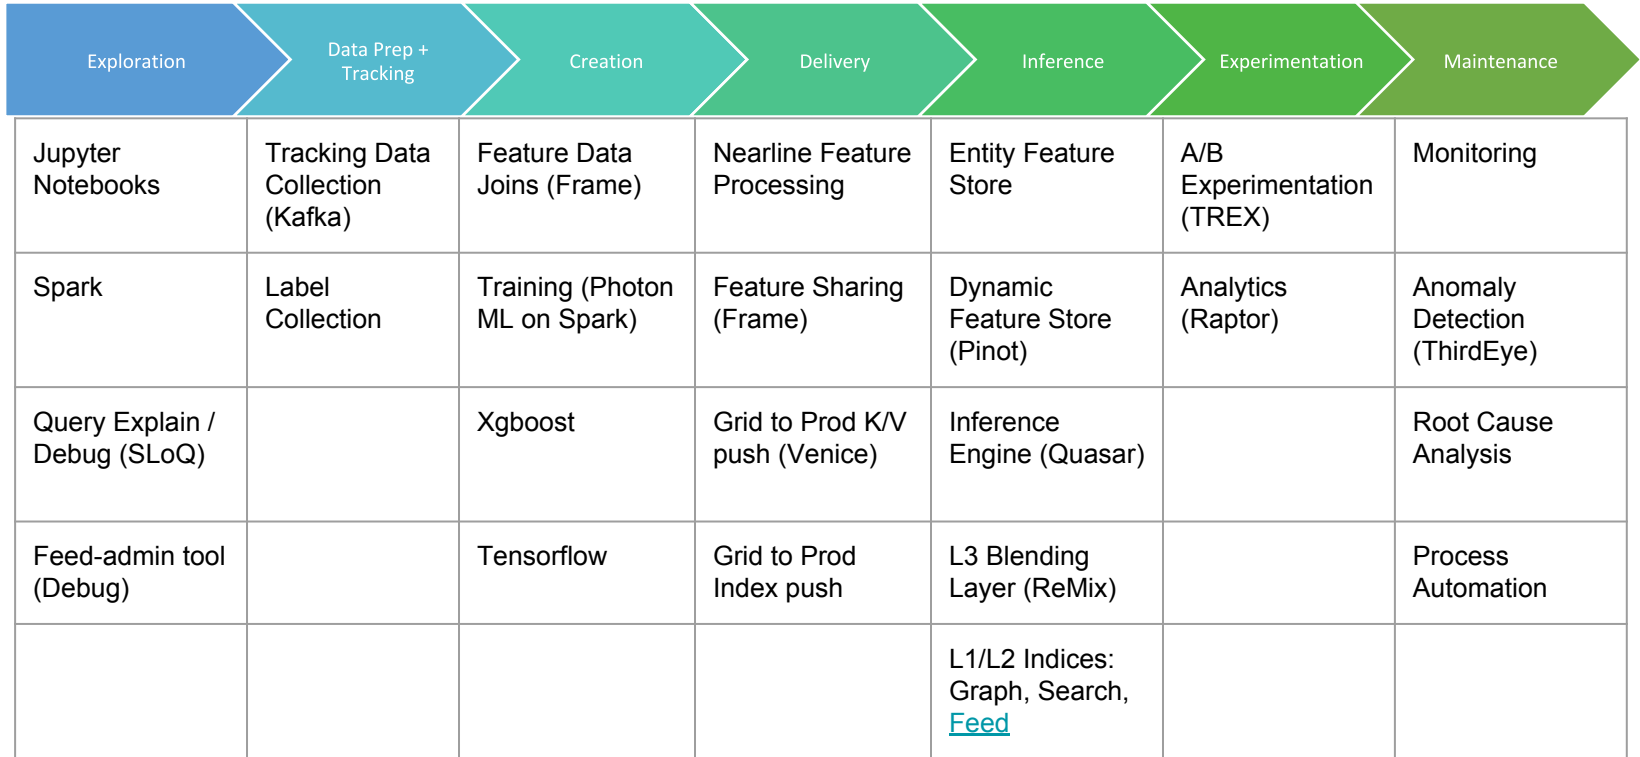


## Surface premium value



## Data access

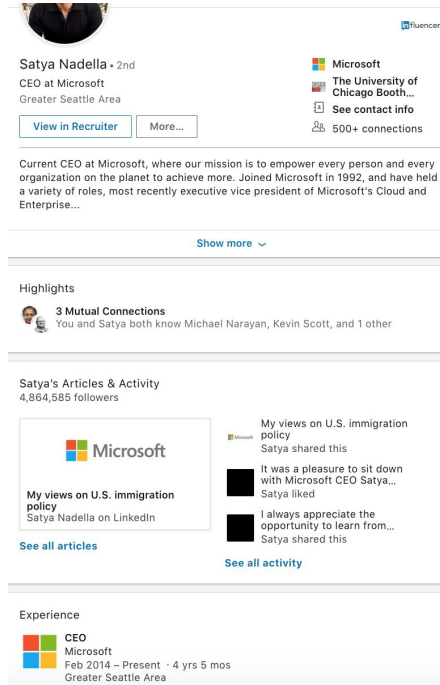
# ML Infrastructure Ecosystem





# LinkedIn operates the largest professional network on the Internet

Tell your Story



Satya Nadella • 2nd  
CEO at Microsoft  
Greater Seattle Area

Microsoft  
The University of Chicago Booth...  
See contact info  
500+ connections

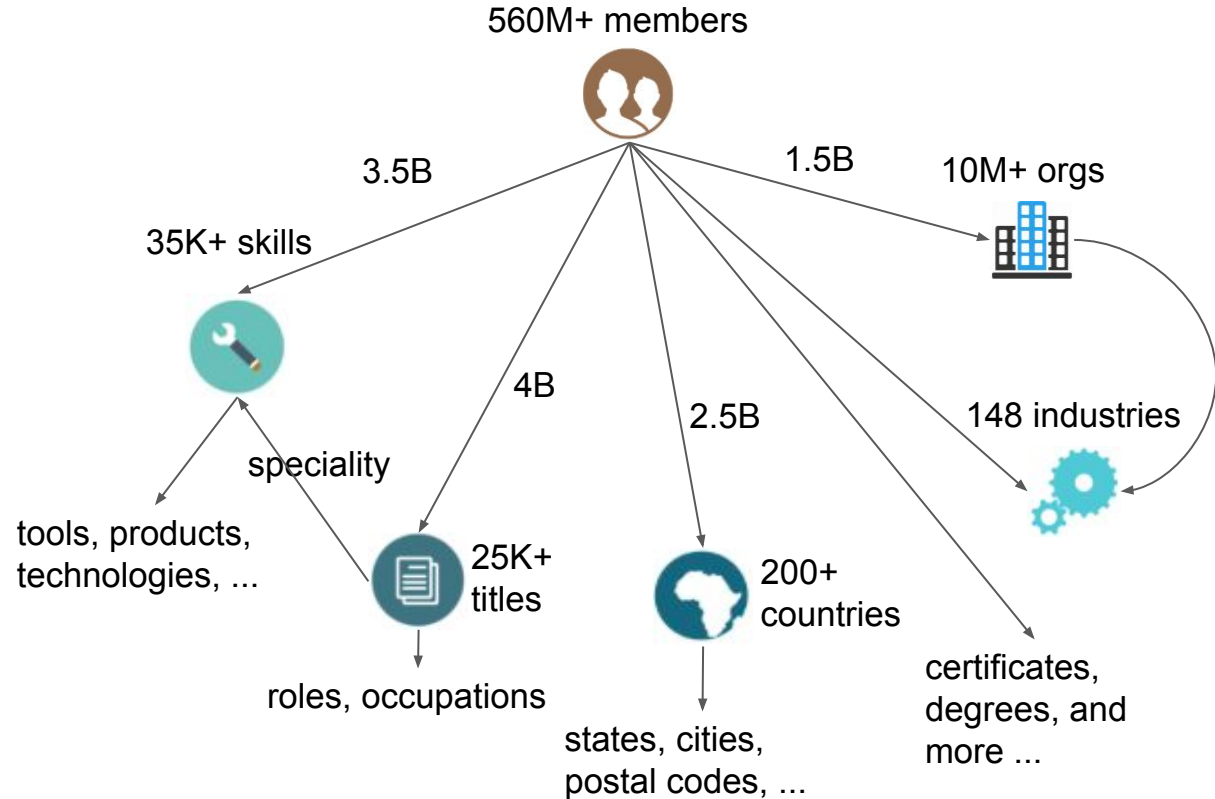
Current CEO at Microsoft, where our mission is to empower every person and every organization on the planet to achieve more. Joined Microsoft in 1992, and have held a variety of roles, most recently executive vice president of Microsoft's Cloud and Enterprise...

Highlights  
3 Mutual Connections  
You and Satya both know Michael Narayan, Kevin Scott, and 1 other

Satya's Articles & Activity  
4,864,585 followers

My views on U.S. immigration policy  
Satya shared this  
It was a pleasure to sit down with Microsoft CEO Satya...  
Satya liked  
I always appreciate the opportunity to learn from...  
Satya shared this


Experience  
CEO  
Microsoft  
Feb 2014 – Present · 4 yrs 5 mos  
Greater Seattle Area




# LinkedIn operates the largest professional network on the Internet

## Make your Connections


Invitations (123) Manage all

 **Danwei (Chloe) Zhang**  
Managing Partner @UpHonest Capital  
👤 Lutz Finger and 4 others

Ignore [Accept](#)

 **Vineet Vinayak**  
Graduate Student Researcher at Georgia  
Institute of Technology | Expertise in...  
👤 Ben Shahshahani and 45 others

Ignore [Accept](#)


 **Faisal Farooq**  
General Chair at KDD 2018 - Leader,  
Innovator, Strategist, Data Scientist  
👤 Deepak Agarwal and 25 others

Ignore [Accept](#)


[Show more](#)

---


People you may know

 **Janardhanan**  
Lifelong Learner  
👤 Brad Ciraulo and 60 others

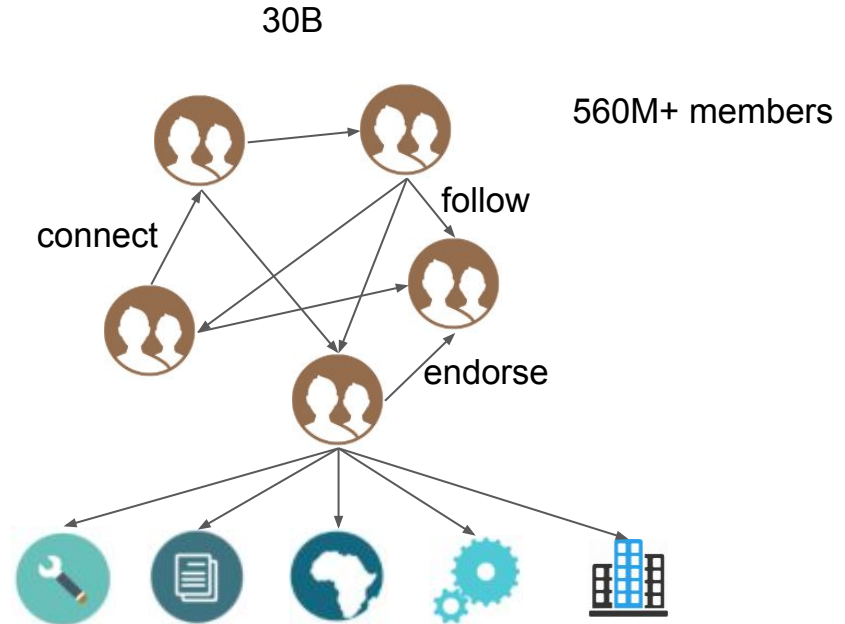
[Connect](#)

 **Jui-Ting Huang**  
Software Engineer in Machine  
Learning  
👤 Badrul Sarwar and 34 others

[Connect](#)

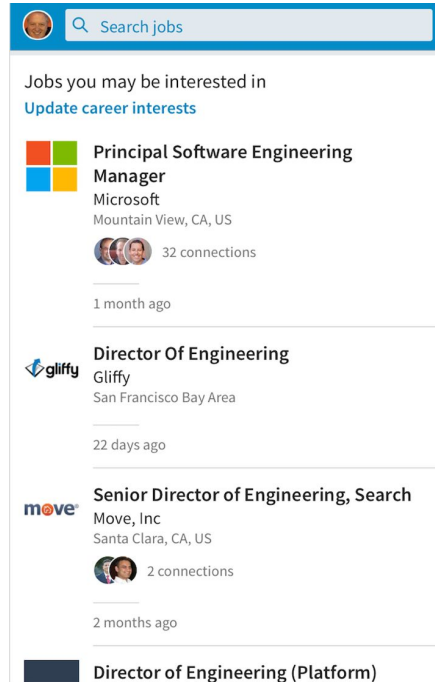
 **Suresh A V L**  
Recruiter at Amazon  
👤 Muthusamy Chelliah  
and 2 others

[Connect](#)



# LinkedIn operates the largest professional network on the Internet

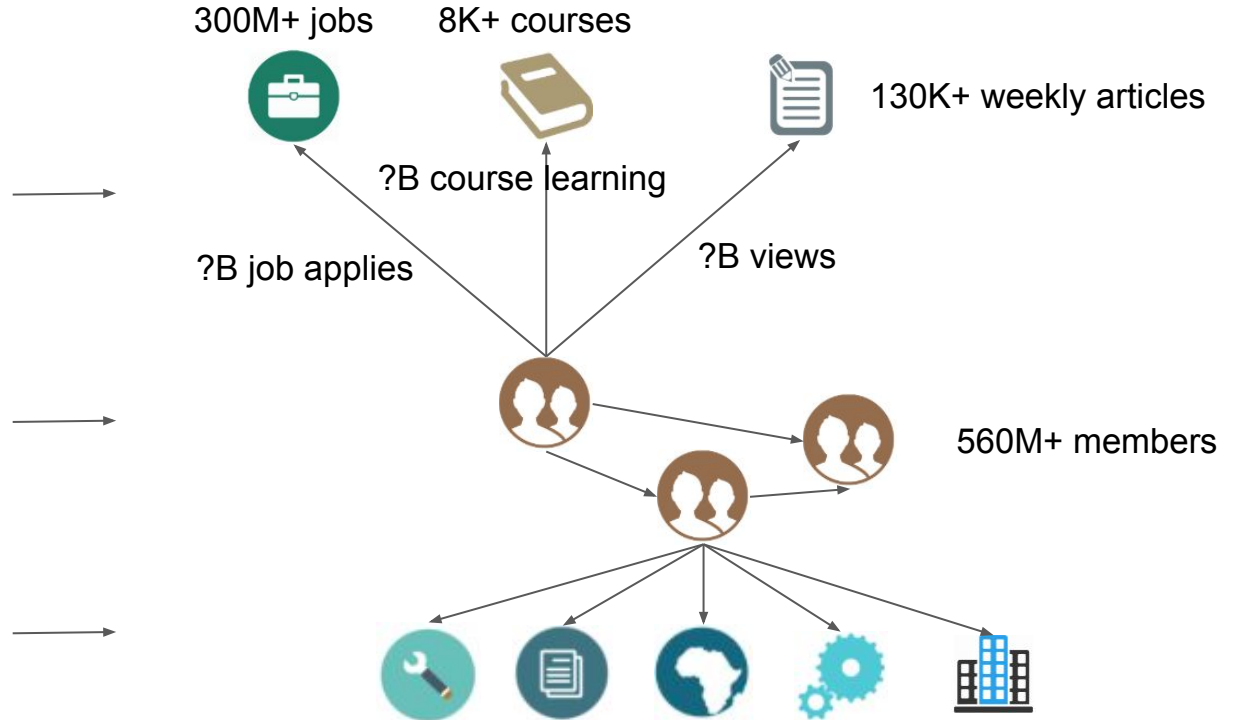
## Explore your Opportunities




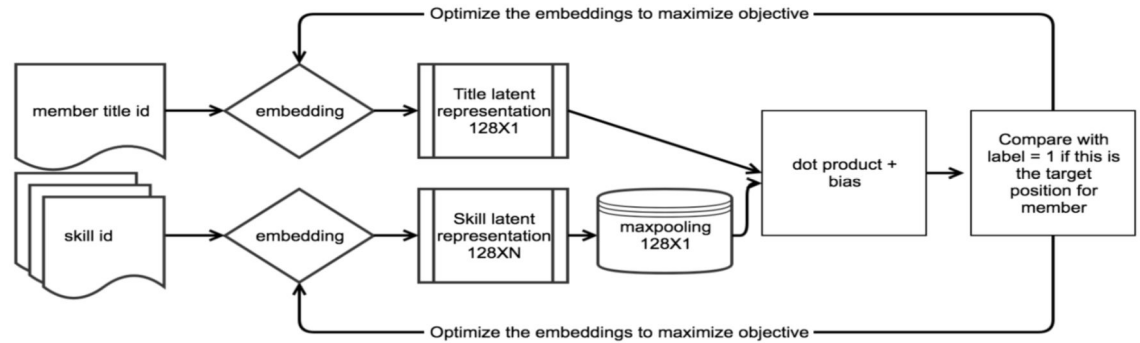
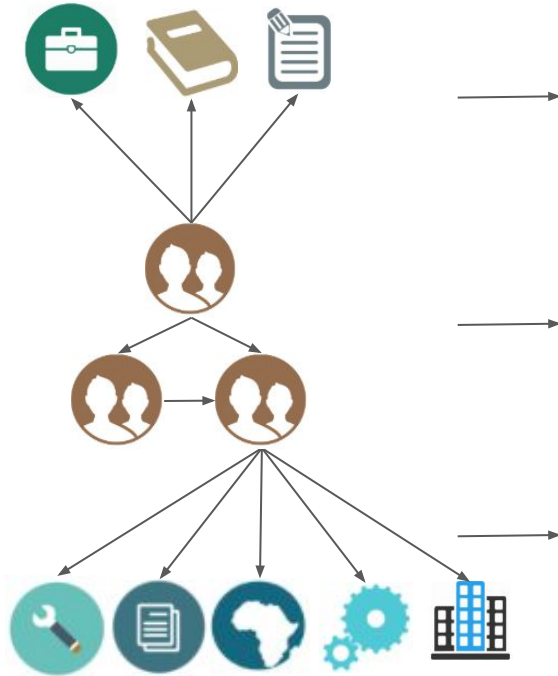
A screenshot of the LinkedIn job search interface. At the top, there is a search bar with the text "Search jobs". Below it, the section "Jobs you may be interested in" is displayed. Three job listings are visible:

- Principal Software Engineering Manager** at Microsoft, Mountain View, CA, US. It shows 32 connections and was posted 1 month ago.
- Director Of Engineering** at Gliffy, San Francisco Bay Area. It was posted 22 days ago.
- Senior Director of Engineering, Search** at Move, Inc, Santa Clara, CA, US. It shows 2 connections and was posted 2 months ago.

At the bottom, a partial listing for "Director of Engineering (Platform)" is visible.



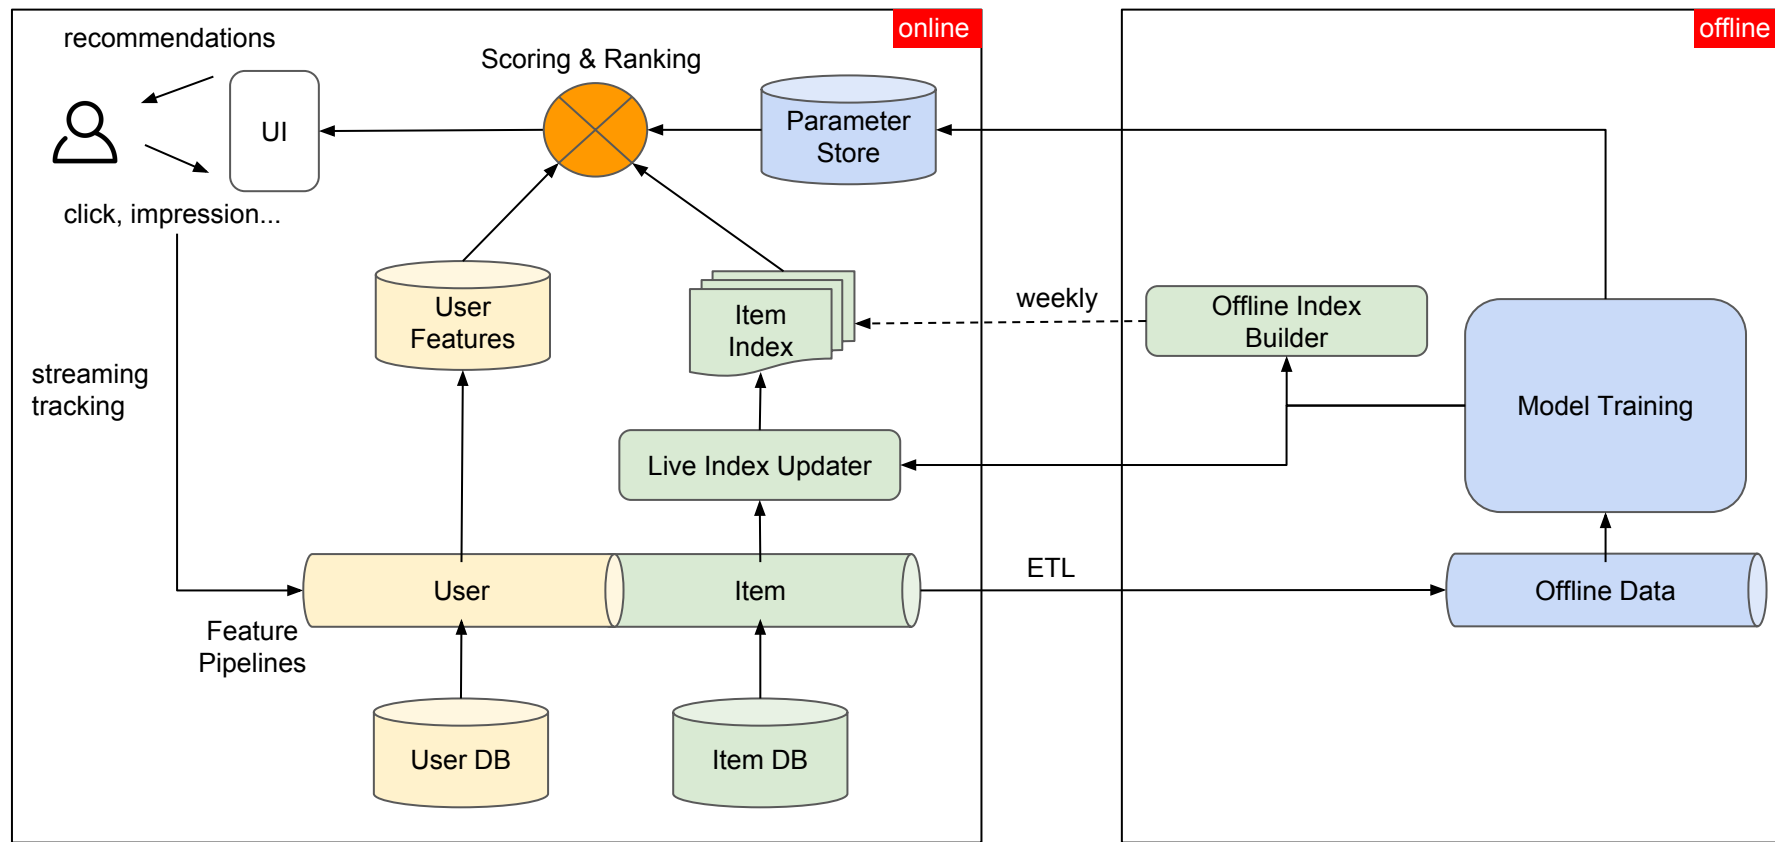
# Network embedding



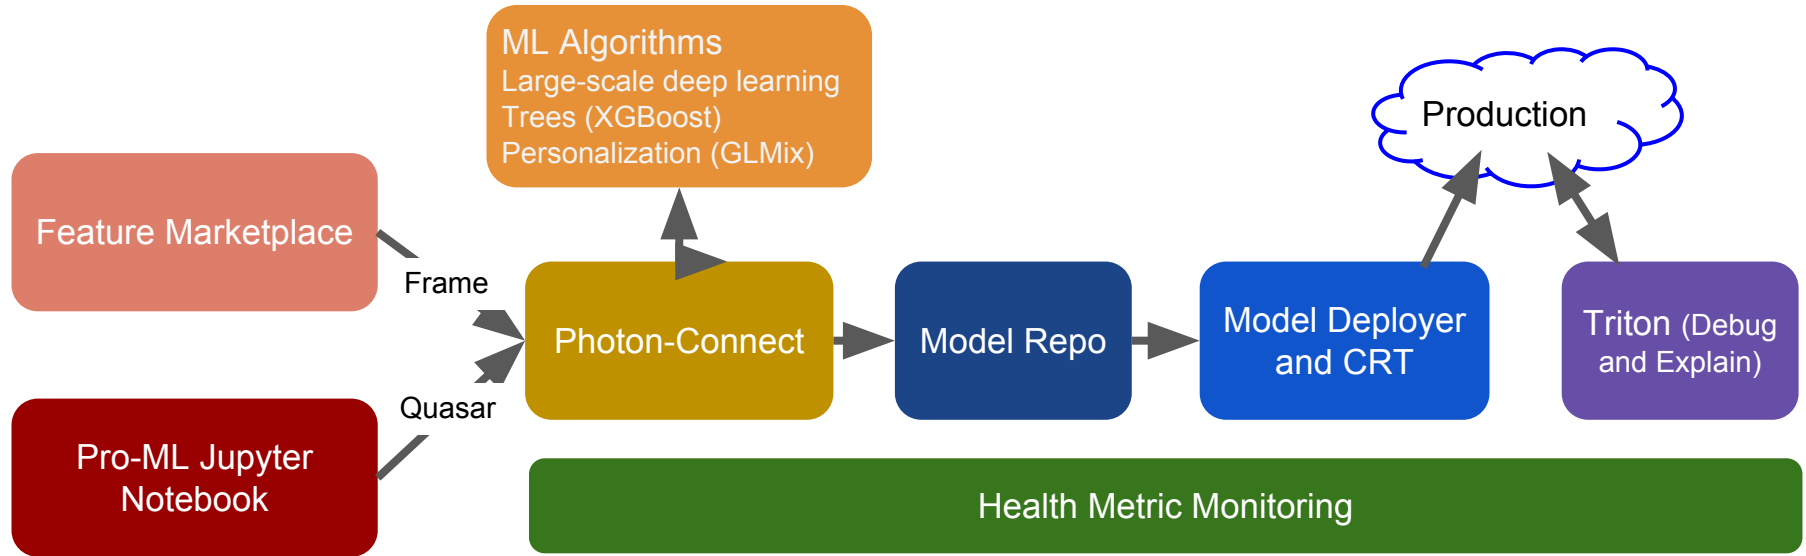
Use the embedding:

- connect with other data, e.g., natural language
- ML model features

# Challenge: Complex Model Deployment Process



# Pro-ML Key Components



## Provide

- A variety of common-used features
- A mechanism to share features across applications
- A feature data access abstraction layer

Enable leverage of features and simplified configuration of complex feature access patterns

## Provide

- A model repo to record all ML models
- A deployer to deploy ML models and the associated data to different destinations in the production environment
- A central release tool to manage the deployment process

Enable ML models and the associated data to be first-class citizens in LinkedIn deployment tooling



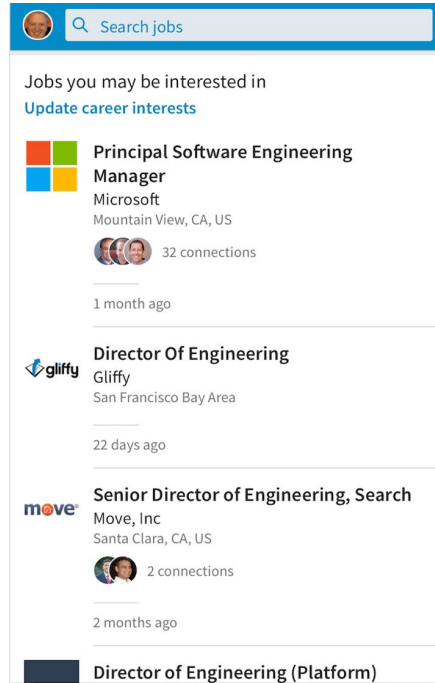
## Provide

- Monitoring of health metrics of ML models, feature data and workflows
- Mechanisms for issue detection and model graceful degradation
- Tools to debug and explain ML models

Enable automatic detection of issues and reduce developers' time on issue investigation

# LinkedIn operates the largest professional network on the Internet

## Explore your Opportunities



A screenshot of the LinkedIn job search interface. At the top, there is a search bar with the text "Search jobs". Below it, a section titled "Jobs you may be interested in" includes a link to "Update career interests". Three job listings are shown:

- Principal Software Engineering Manager** at Microsoft, Mountain View, CA, US. 32 connections. Posted 1 month ago.
- Director Of Engineering** at Gliffy, San Francisco Bay Area. Posted 22 days ago.
- Senior Director of Engineering, Search** at Move, Inc, Santa Clara, CA, US. 2 connections. Posted 2 months ago.

At the bottom, a partial listing for "Director of Engineering (Platform)" is visible.

