Full Stack Deep Learning

ML Teams

Josh Tobin, Sergey Karayev, Pieter Abbeel

Outline

- The AI talent gap
- ML-Related roles
- ML team structures
- The hiring process
- Exam for this course

Outline

- **The AI talent gap**
- ML-Related roles
- ML team structures
- The hiring process
- Exam for this course

The AI Talent Gap

How many people know how to build AI systems?
5,000 (actively publishing research [Element AI])
10,000 (estimated num people with the right skillset [Element AI])
22,000 (PhD-educated AI researchers [Bloomberg])
90,000 (upper bound on number of people [Element AI])
200,000 - 300,000 (Number of AI researcher / practitioners [Tencent])

3.6M (Number of software developers in the US)
18.2M (Number of software developers in the world)

Sources: The AI Talent Shortage (Nikolai Yakovenko) https://medium.com/@Moscow25/the-ai-talent-shortage-704d8cf0c4cc
Just How Shallow is the Artificial Intelligence Talent Pool (Jeremy Kahn)
https://www.bloomberg.com/news/articles/2018-02-07/just-how-shallow-is-the-artificial-intelligence-talent-pool

The AI talent gap

Fierce competition for AI talent

“Everyone agrees that the competition to hire people who know how to build artificial intelligence systems is intense. It’s turned once-staid academic conferences into frenzied meet markets for corporate recruiters and driven the salaries of the top researchers to seven-figures.”

(Bloomberg)

Sources: Just How Shallow is the Artificial Intelligence Talent Pool (Jeremy Kahn)
https://www.bloomberg.com/news/articles/2018-02-07/just-how-shallow-is-the-artificial-intelligence-talent-pool

The AI talent gap

Fierce competition for AI talent

“Hiring is crazy right now. ML is a young field that got popular very quickly. There’s a ton of demand and not a lot of supply.”

(Computer Vision Engineer at Series C startup)

The AI talent gap

Fierce competition for AI talent

“Hiring for ML is really challenging and takes way more time and effort than we expected. We have someone working on it full-time and we’re still only able to get a few people per quarter”

(Startup Founder)

Outline

- The AI talent gap
- **ML-Related roles**
- ML team structures
- The hiring process
- Exam for this course

Most common ML roles at companies?

- DevOps
- Data engineer
- ML engineer
- ML researcher
- Data scientist

What's the difference?

Breakdown of job function by role

Role	Job Function	Work product	Commonly used tools
DevOps Engineer	Deploy & monitor production systems	Deployed product	AWS, etc.

Breakdown of job function by role

Role	Job Function	Work product	Commonly used tools
DevOps Engineer	Deploy & monitor production systems	Deployed product	AWS, etc.
Data engineer	Build data pipelines, aggregation, storage, monitoring	Distributed system	Hadoop, Kafka, Airflow

Breakdown of job function by role

Role	Job Function	Work product	Commonly used tools
DevOps Engineer	Deploy & monitor production systems	Deployed product	AWS, etc.
Data engineer	Build data pipelines, aggregation, storage, monitoring	Distributed system	Hadoop, Kafka, Airflow
ML Engineer	Train & deploy prediction models	Prediction system running on real data (often in production)	Tensorflow, Docker

Breakdown of job function by role

Role	Job Function	Work product	Commonly used tools
DevOps Engineer	Deploy & monitor production systems	Deployed product	AWS, etc.
Data engineer	Build data pipelines, aggregation, storage, monitoring	Distributed system	Hadoop, Kafka, Airflow
ML Engineer	Train & deploy prediction models	Prediction system running on real data (often in production)	Tensorflow, Docker
ML Researcher	Train prediction models (often forward looking or not production-critical)	Prediction model & report describing it	Tensorflow, pytorch, Jupyter

Breakdown of job function by role

Role	Job Function	Work product	Commonly used tools
DevOps Engineer	Deploy & monitor production systems	Deployed product	AWS, etc.
Data engineer	Build data pipelines, aggregation, storage, monitoring	Distributed system	Hadoop, Kafka, Airflow
ML Engineer	Train & deploy prediction models	Prediction system running on real data (often in production)	Tensorflow, Docker
ML Researcher	Train prediction models (often forward looking or not production-critical)	Prediction model & report describing it	Tensorflow, pytorch, Jupyter
Data Scientist	Blanket term used to describe all of the above. In some orgs, means answering business questions using analytics	Prediction model or report	SQL, Excel, Jupyter, Pandas, SKLearn, Tensorflow

What skills are needed for the roles?

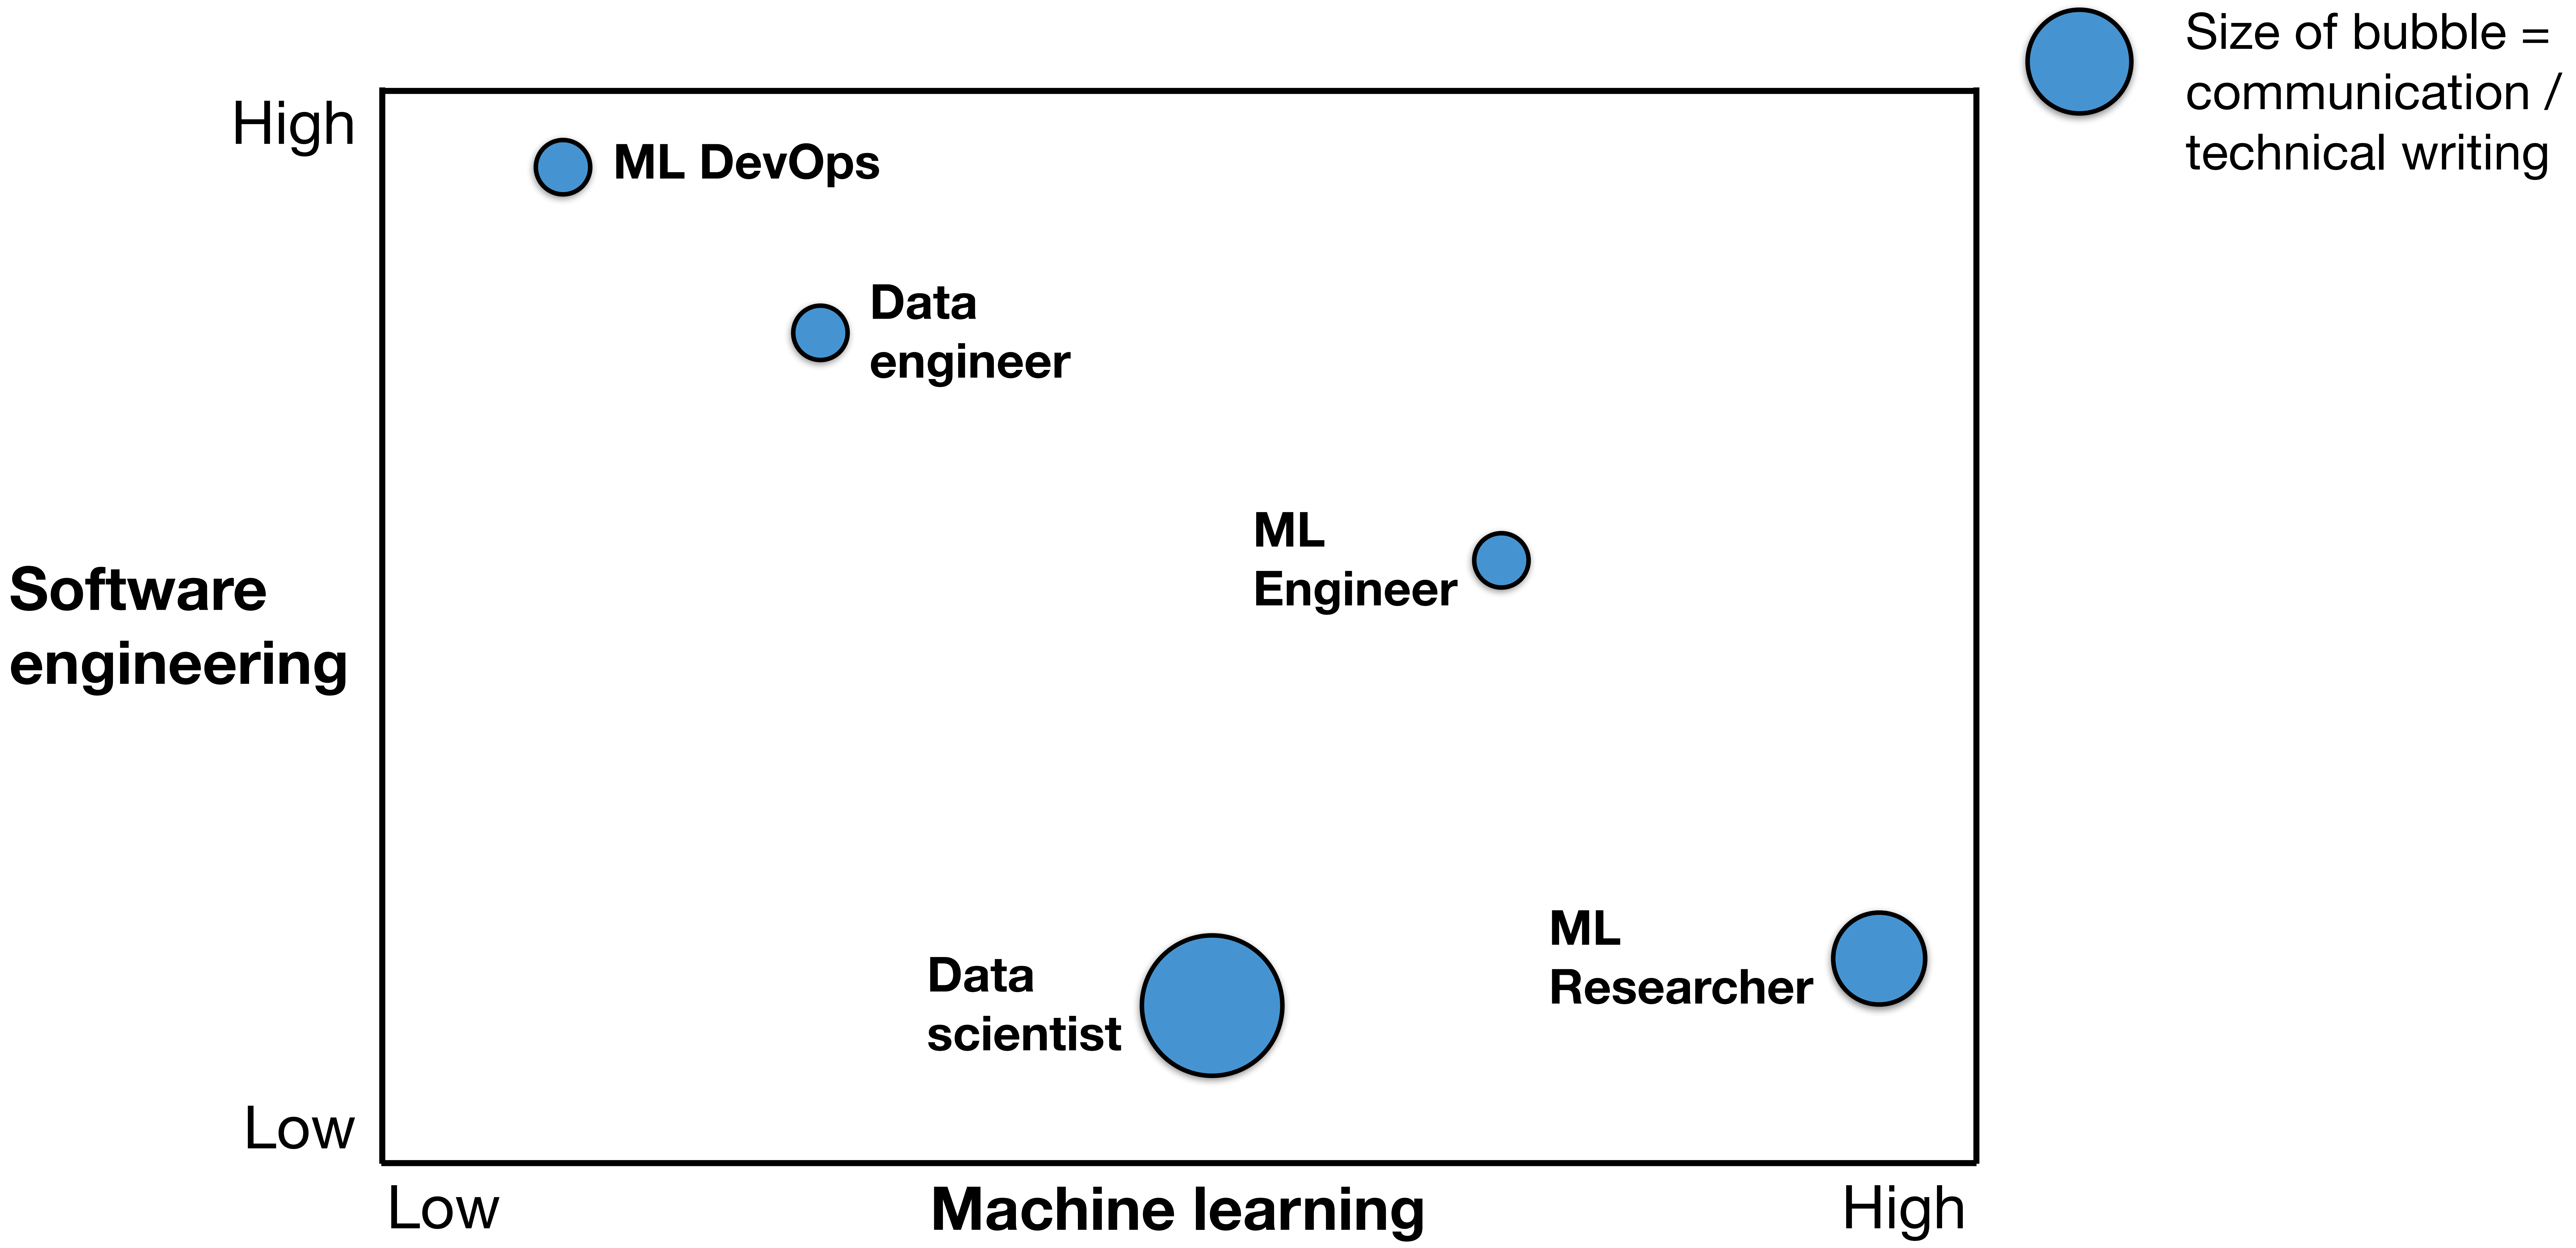

What skills are needed for the roles?

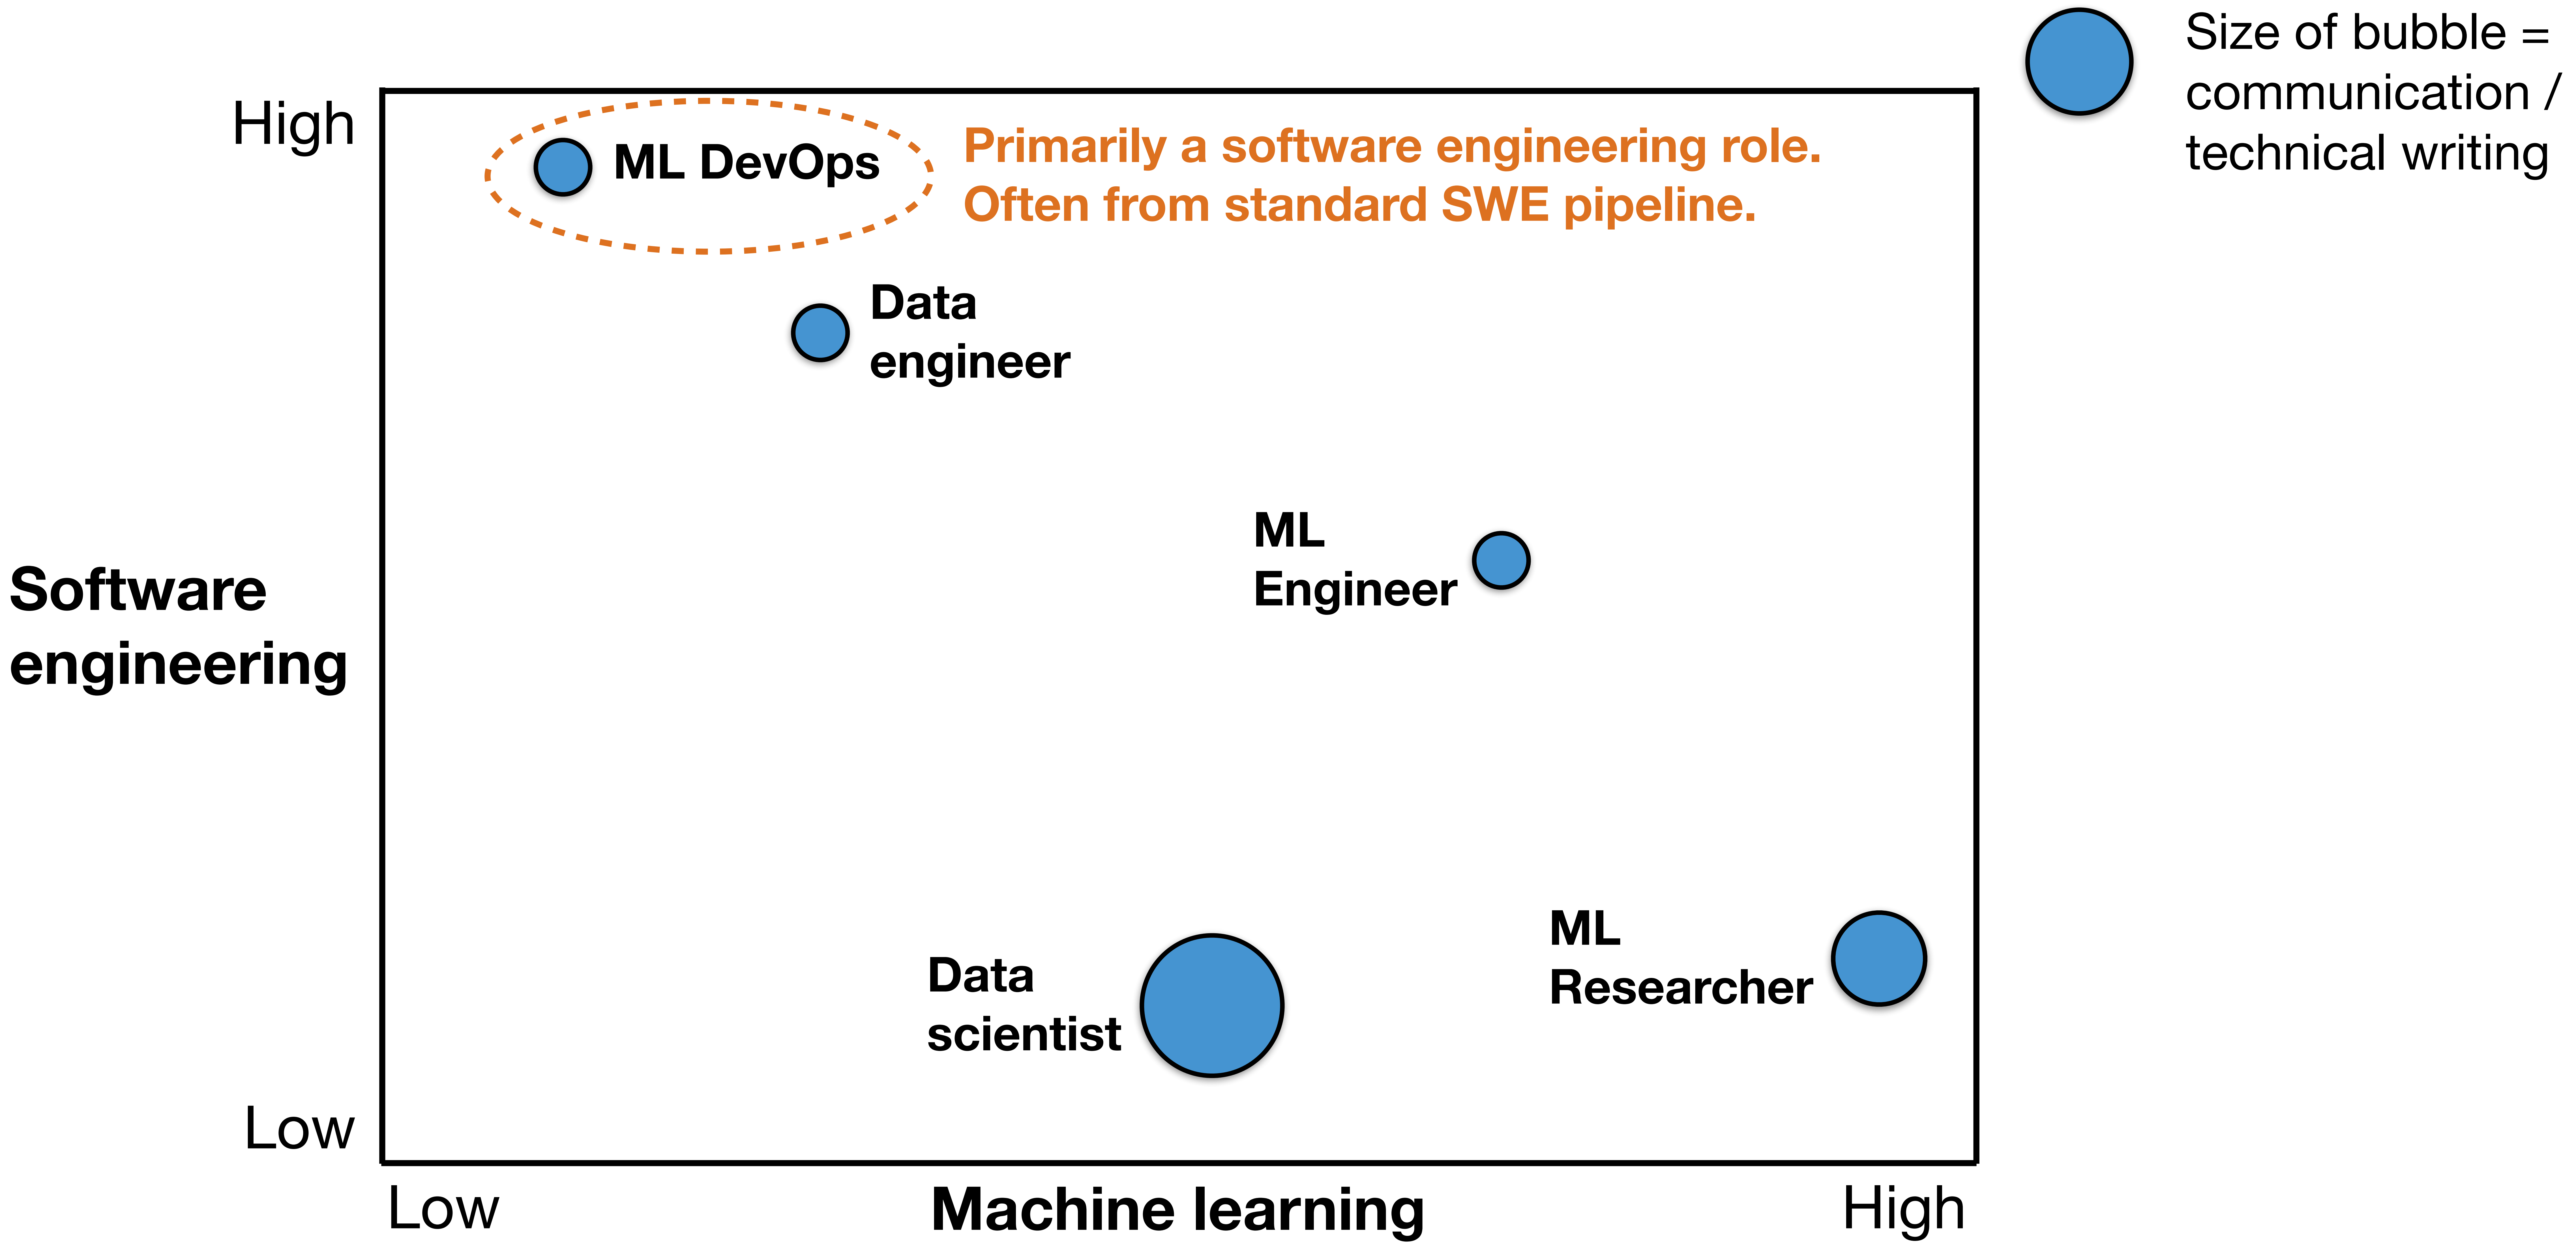

What skills are needed for the roles?

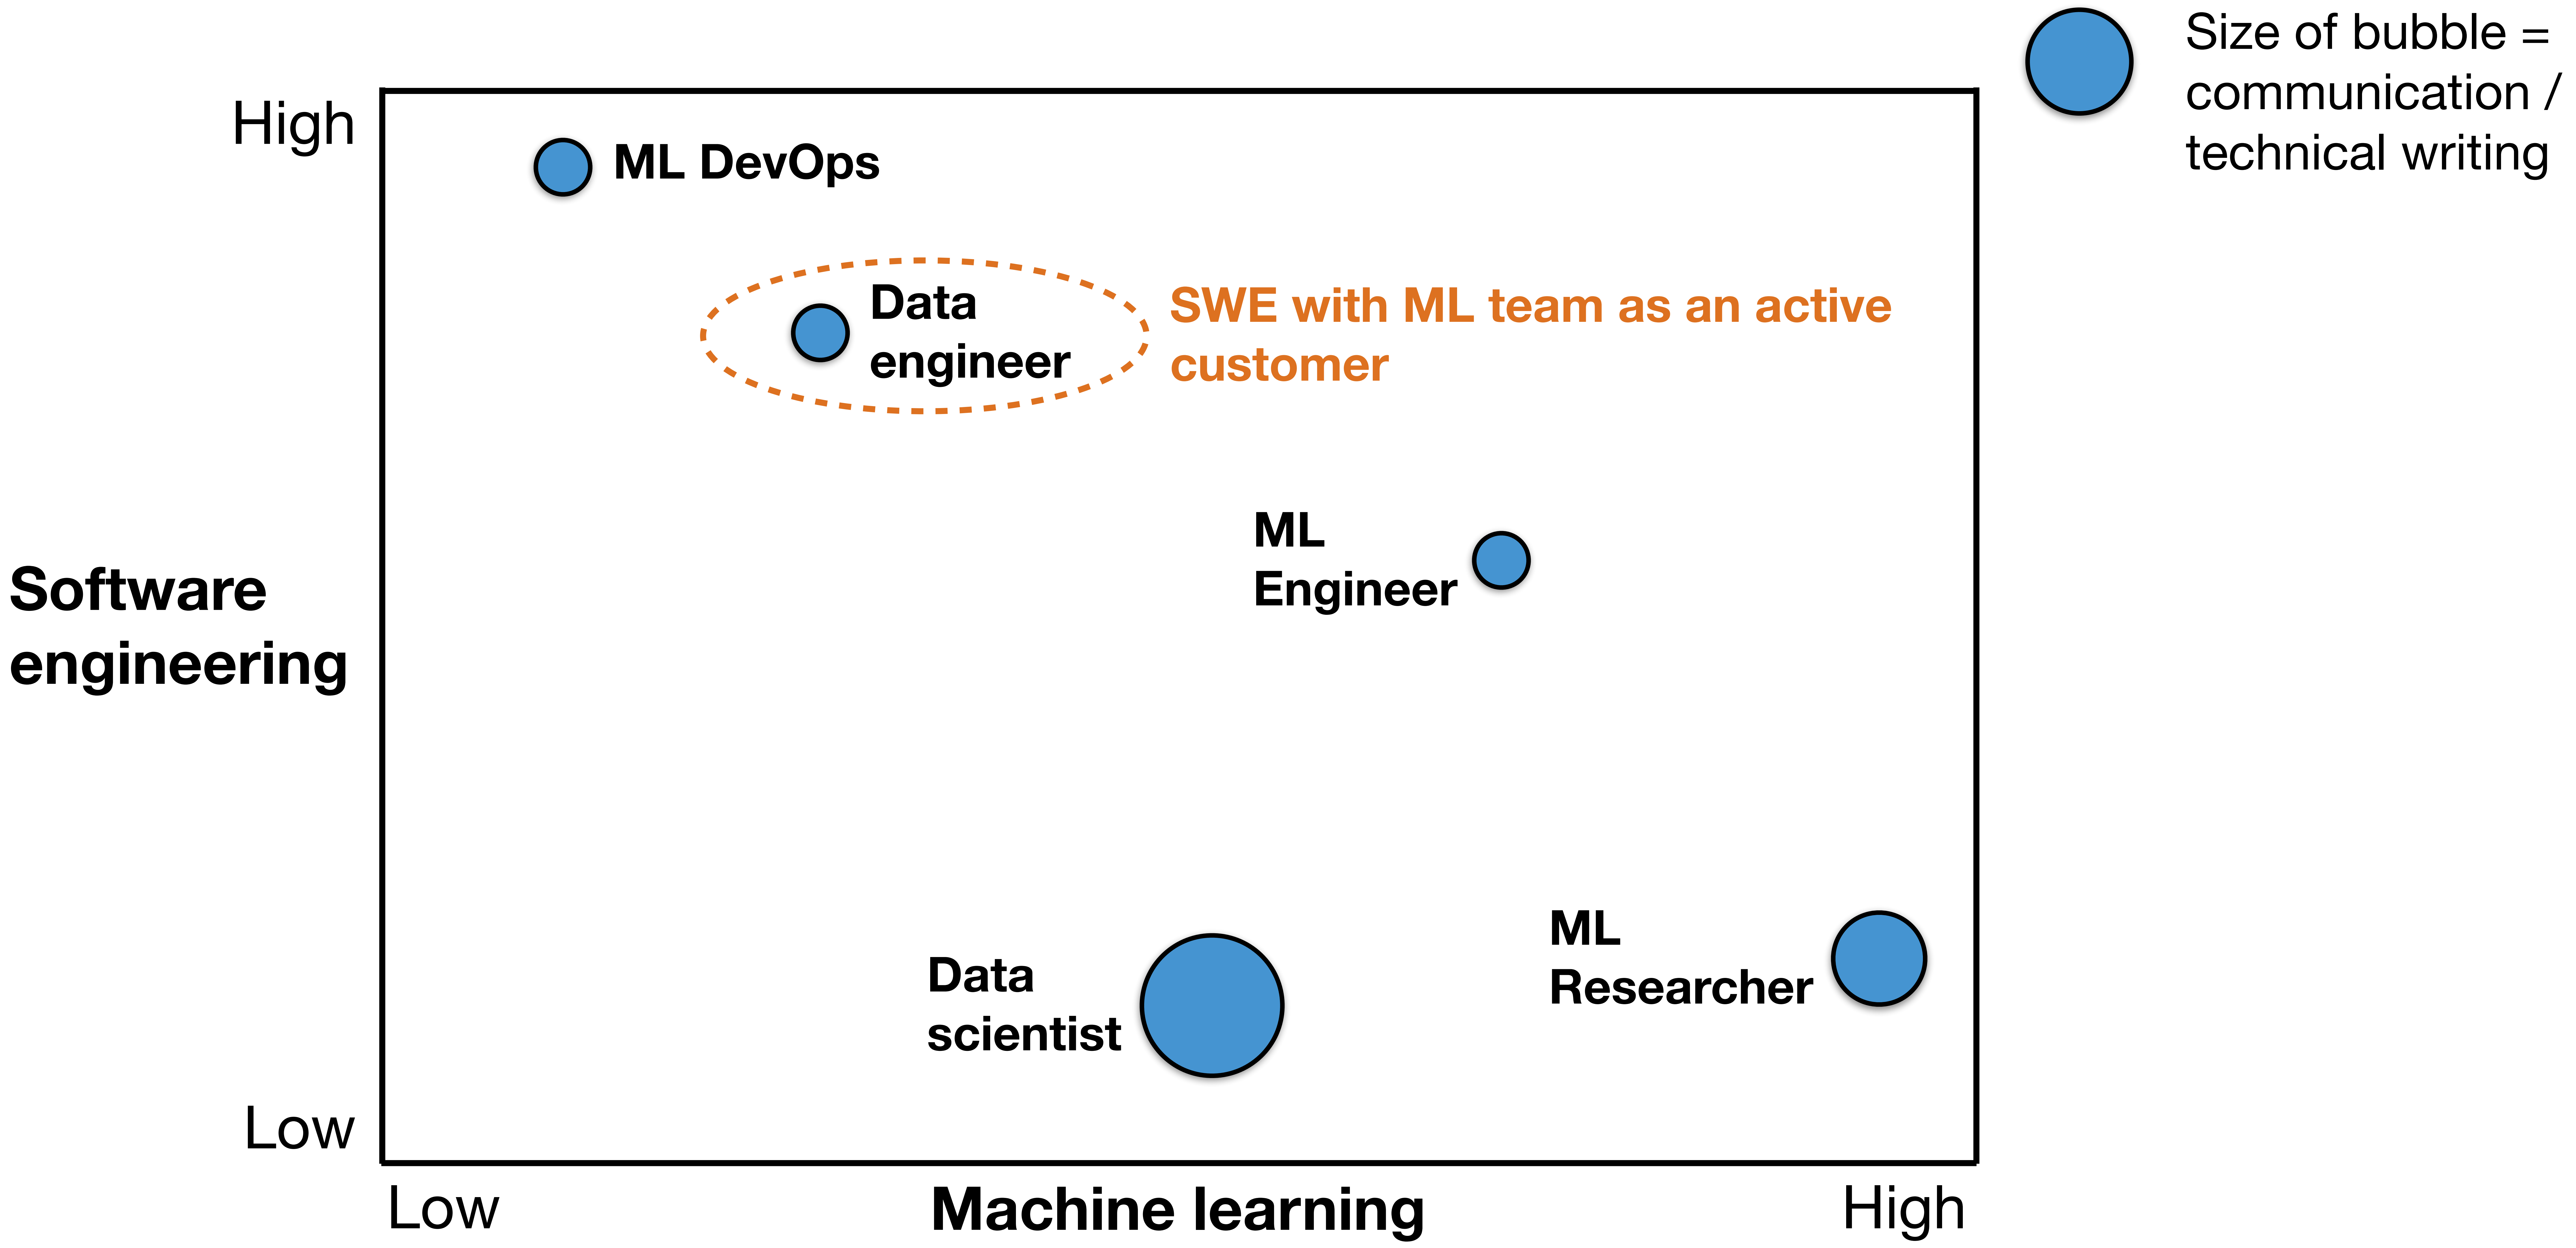

What skills are needed for the roles?

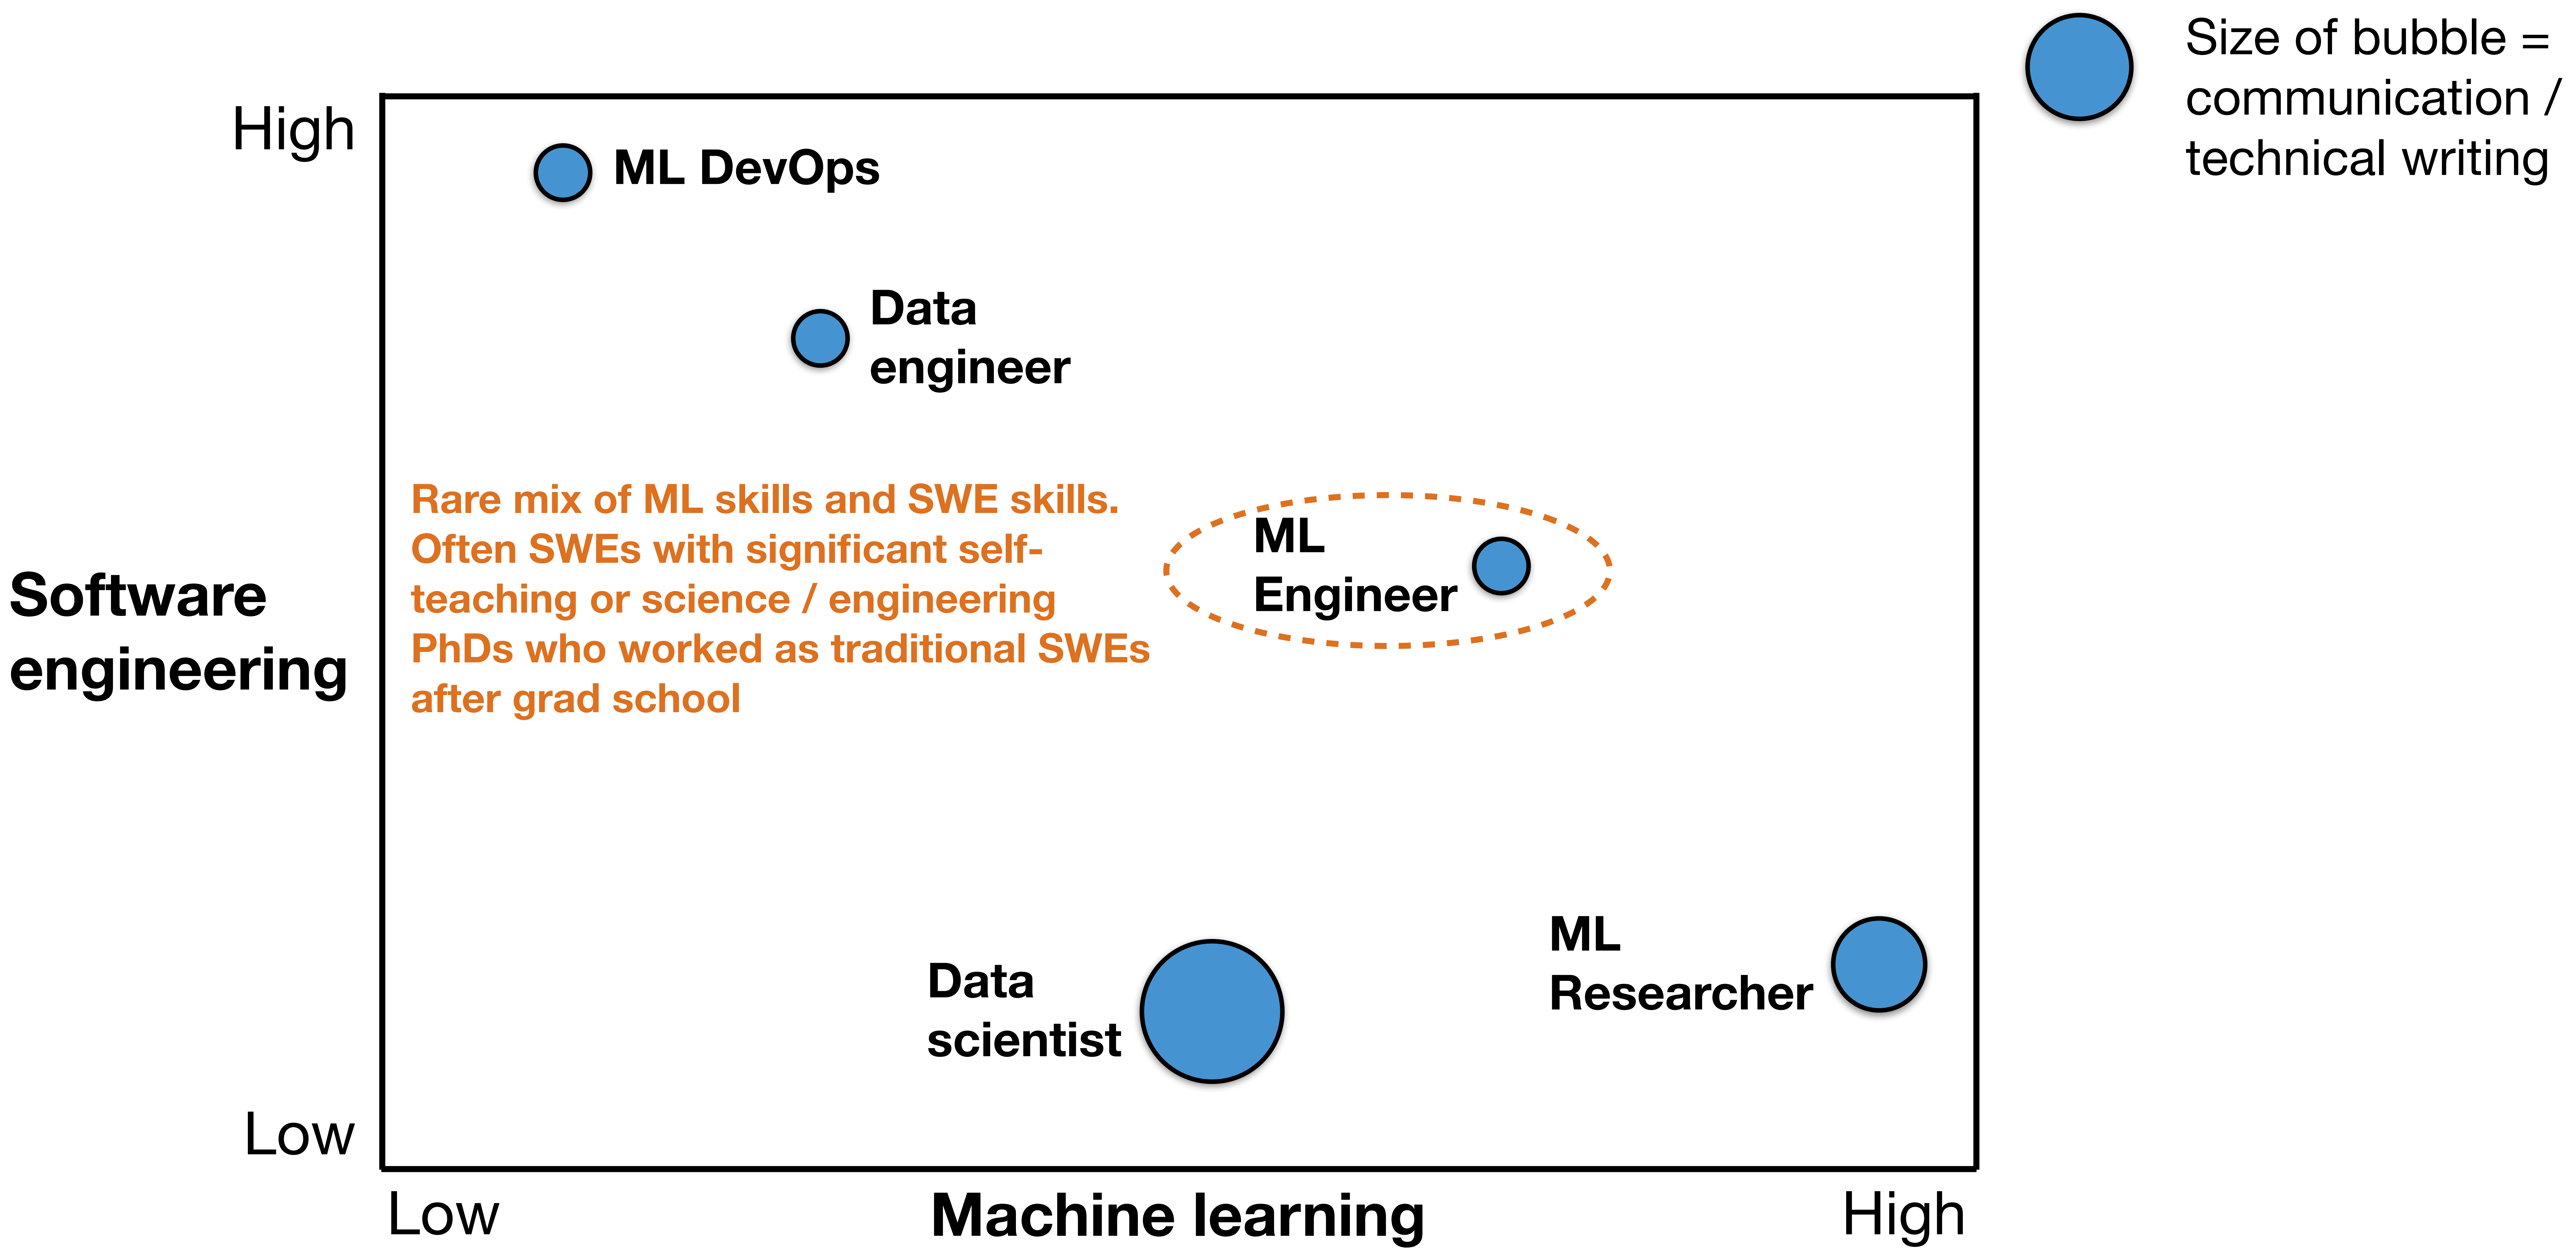

What skills are needed for the roles?

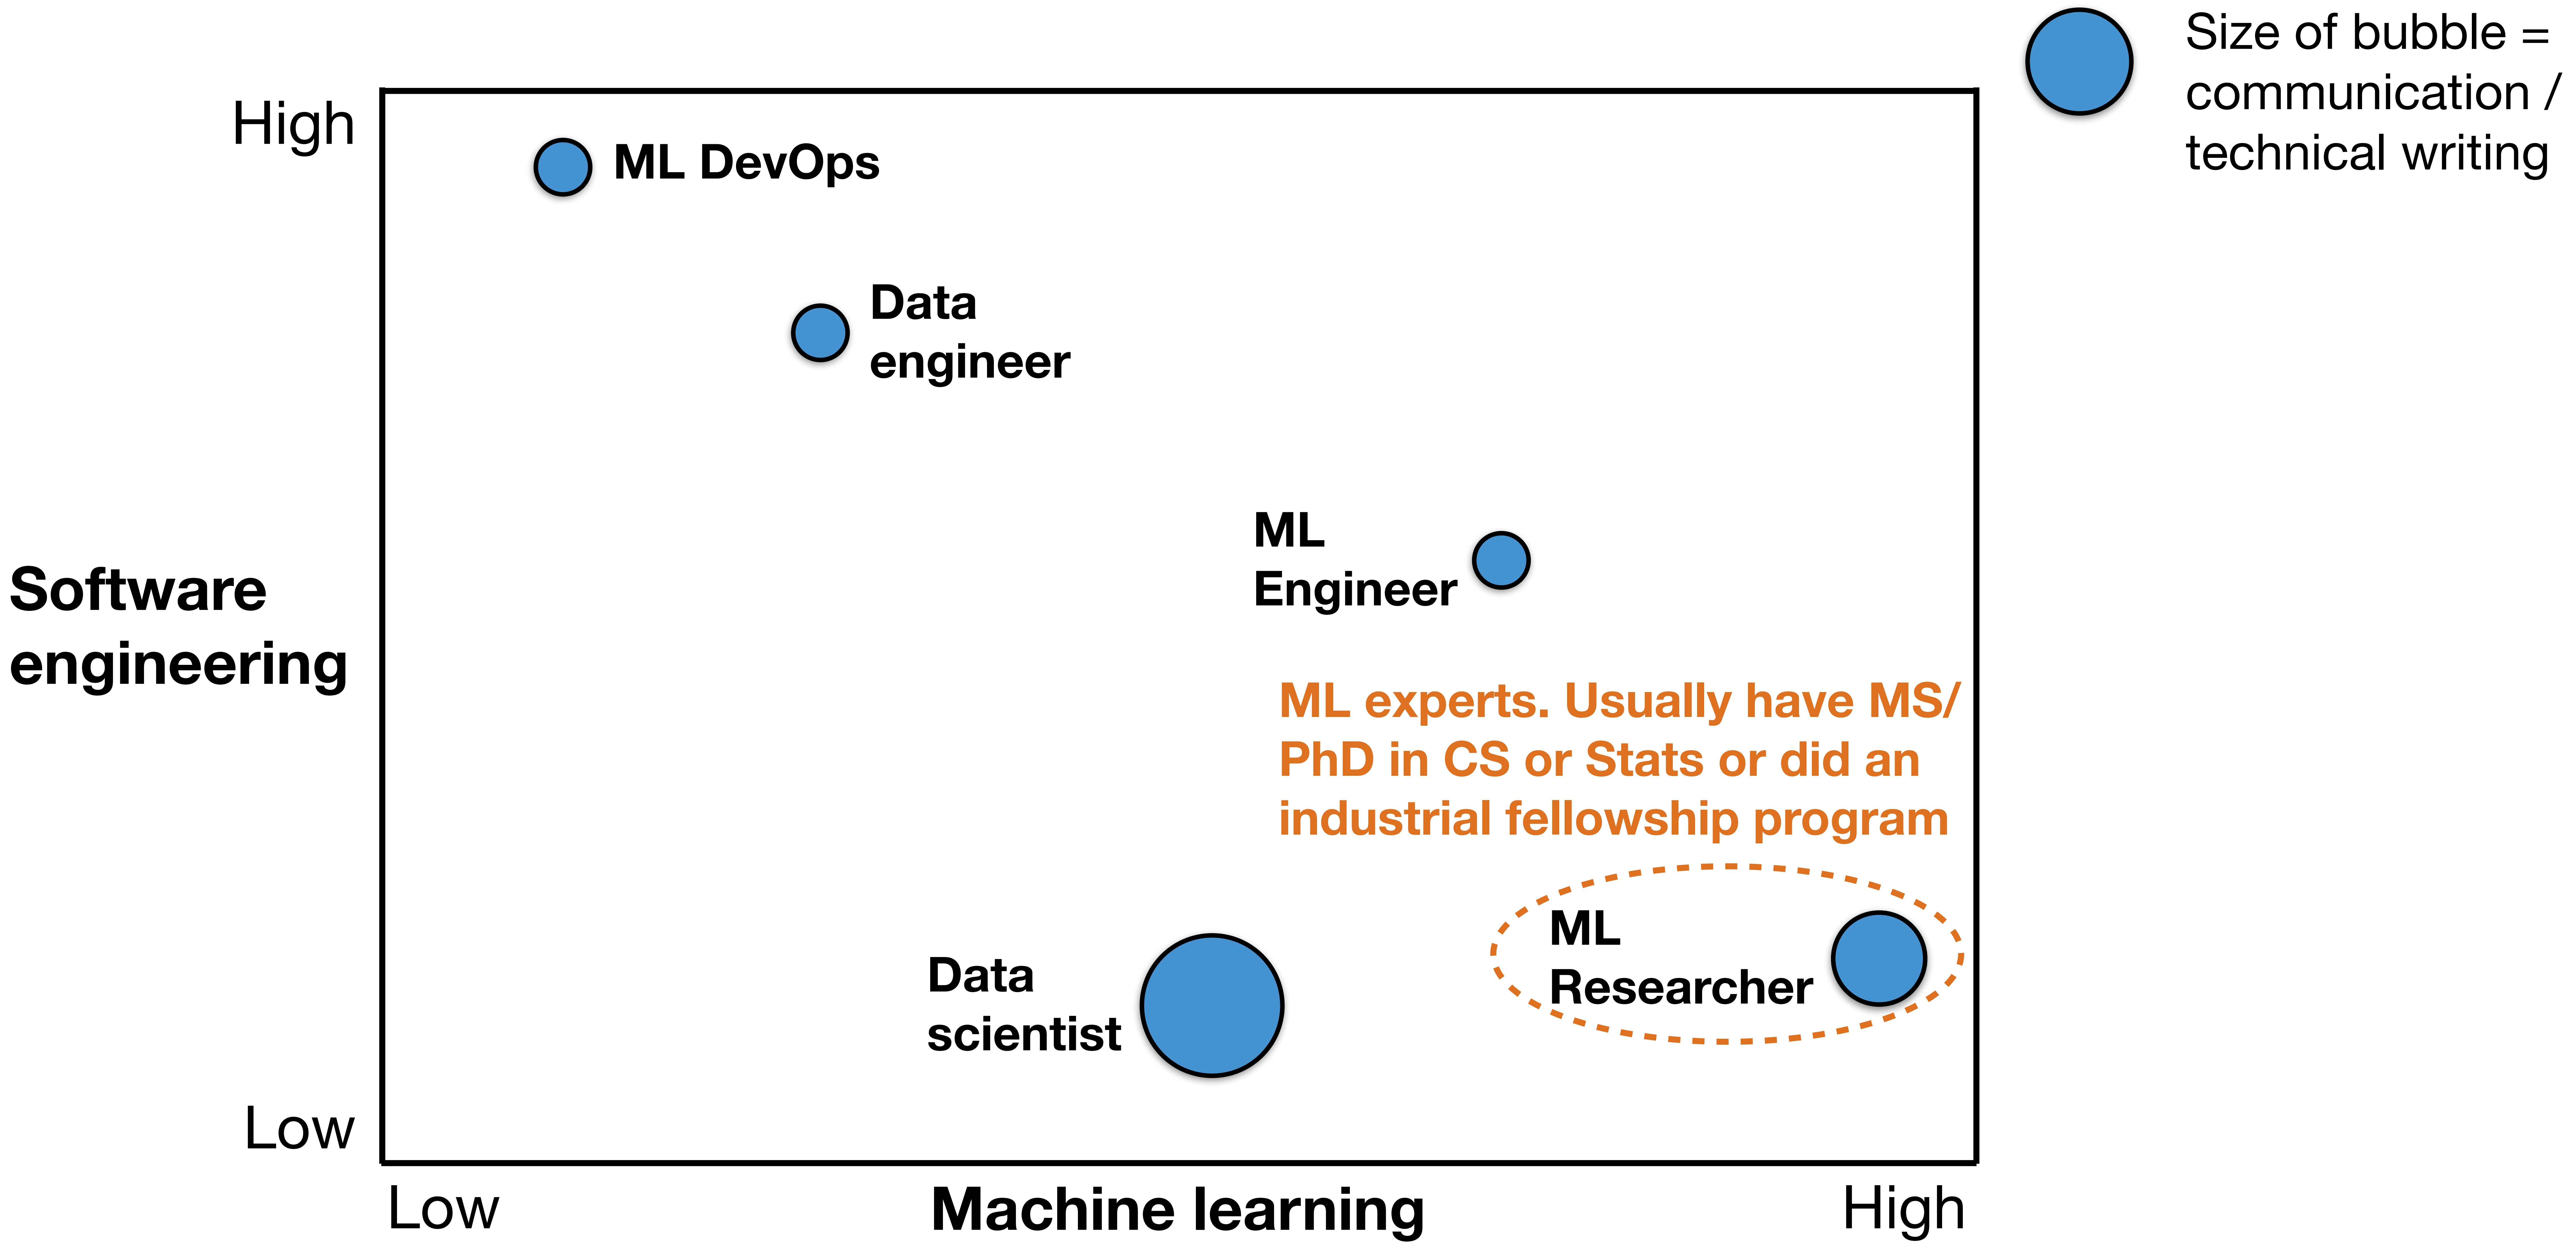

What skills are needed for the roles?

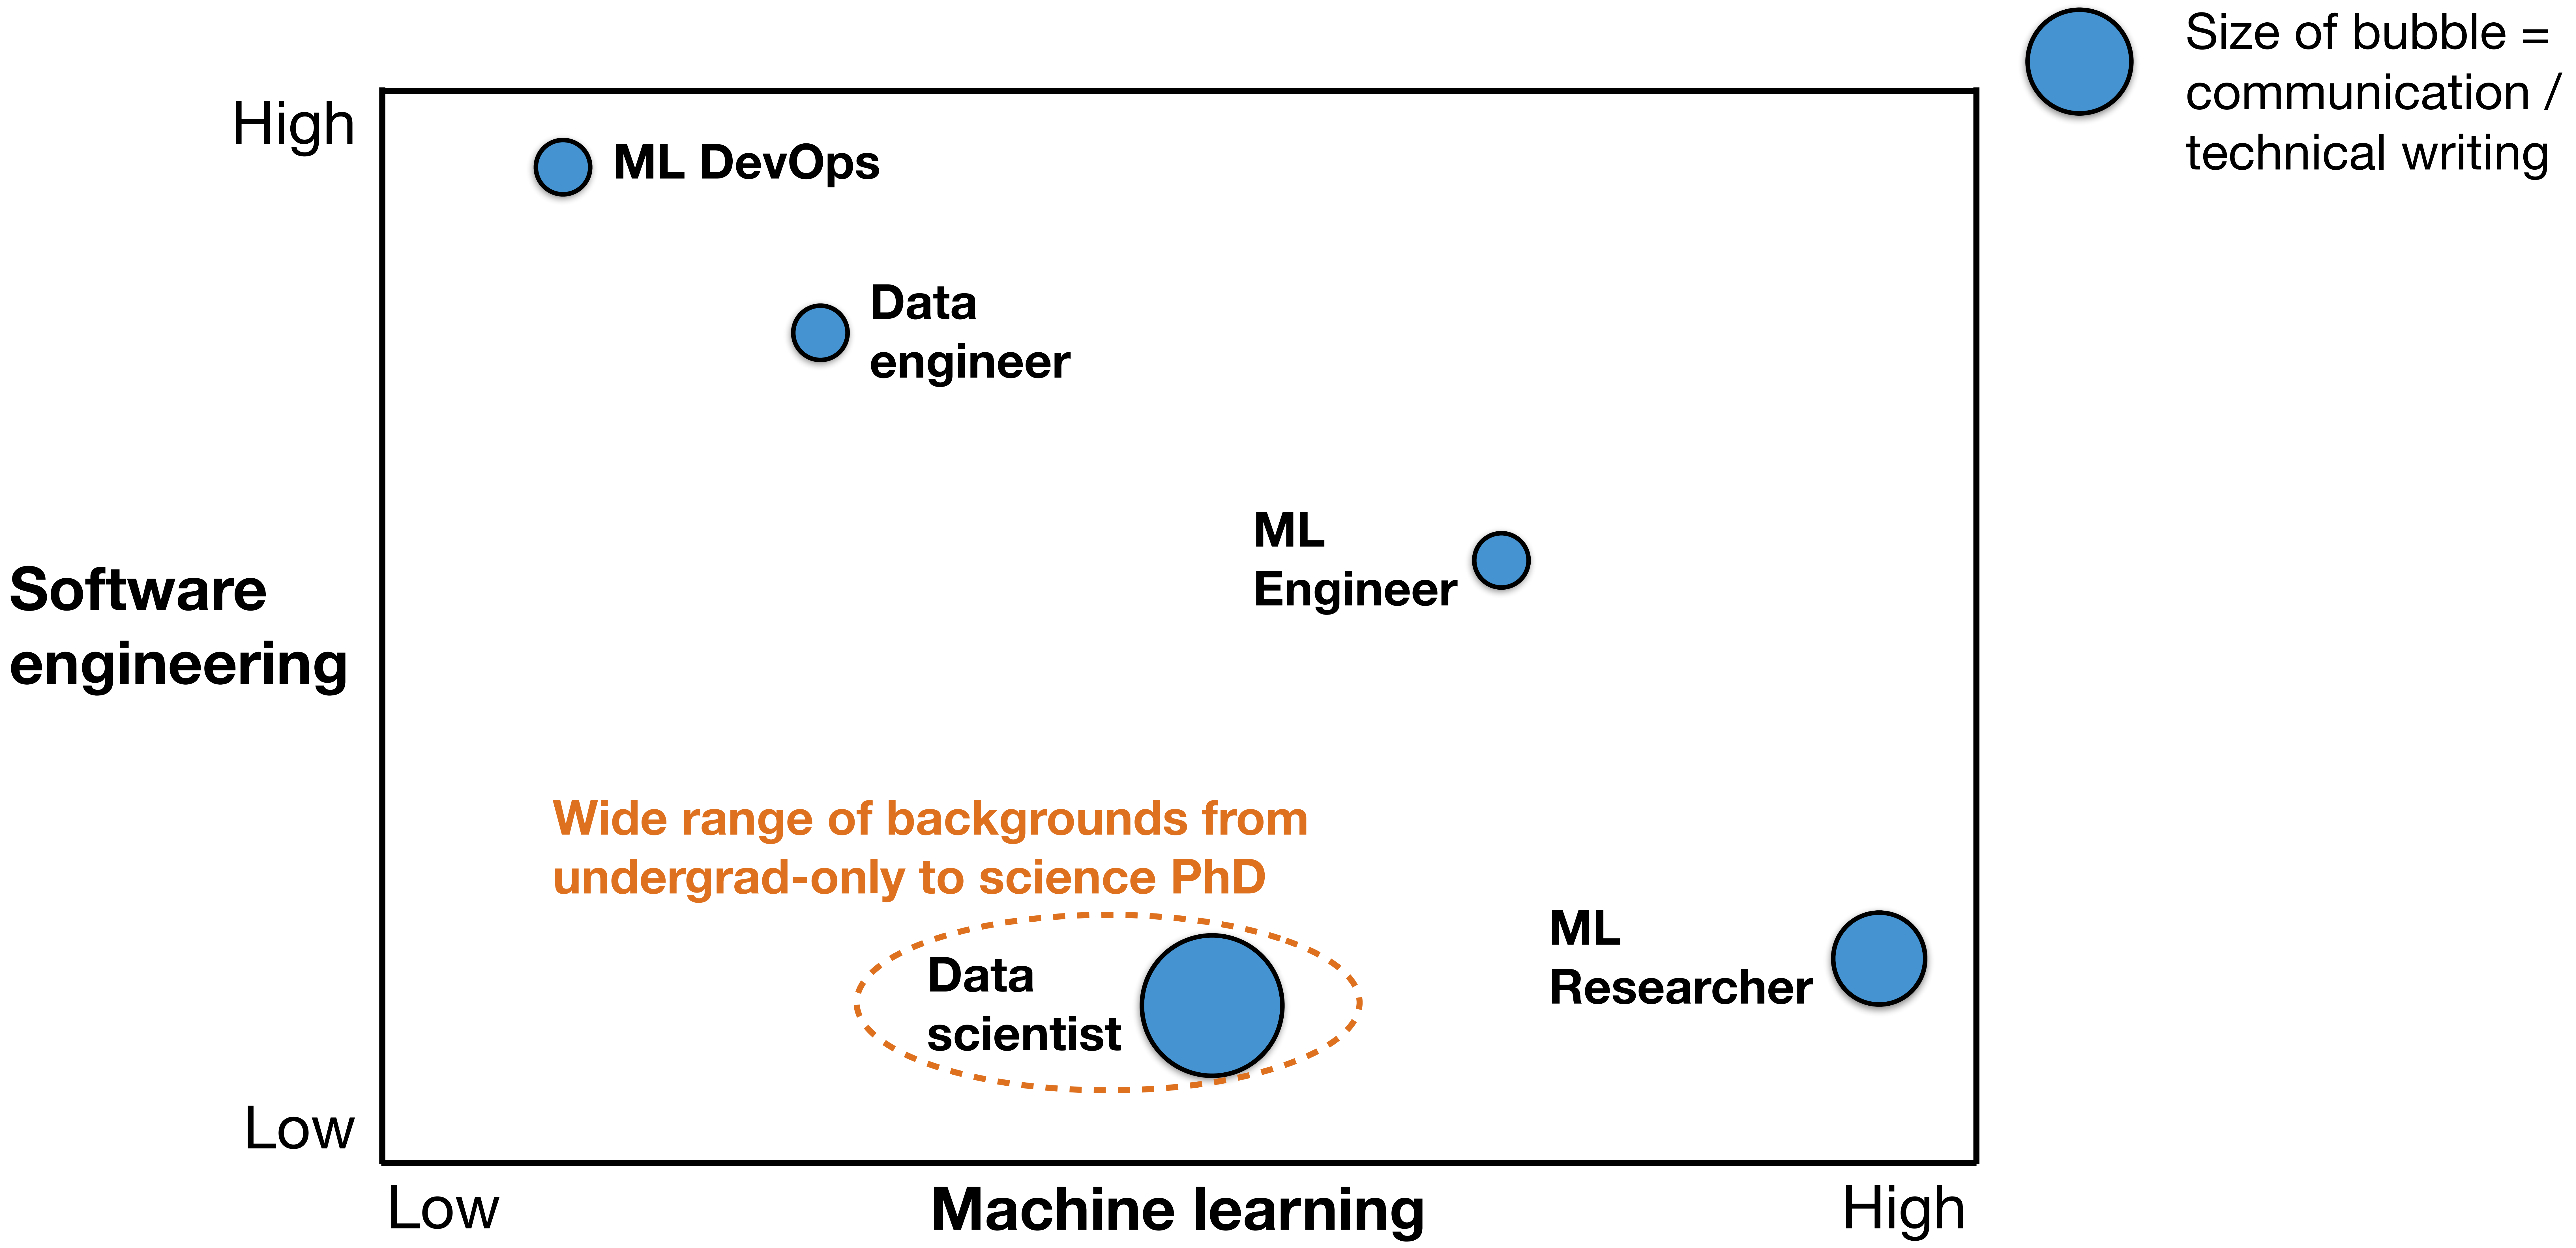

Outline

- The AI talent gap
- ML-Related roles
- **ML teams**
- The hiring process
- Exam for this course

ML team structures - lessons learned

No consensus yet on the right way to structure a ML team

ML team structures - lessons learned

- Most believe in a **mix of SWE and ML skillsets** on the team
- Many think **everyone on the team needs some level of SWE skills**
- Different views on **ML researchers**
  - Some think they are **hard to integrate with SWE** teams
  - Others think **deep ML expertise is necessary to move fast**
- Different views on **data engineering**
  - **Sits with ML team** in some orgs
  - Others think it should be a separate team (**“data warehousing”**)
  - Some think it’s also important to have **dedicated data labeling** (e.g., building labeling tools & managing outsourced labelers)

Managing ML teams is challenging

- It’s hard to tell in advance how hard or easy something is

Managing ML teams is challenging

- It’s hard to tell in advance how hard or easy something is

Accuracy improvement in first week

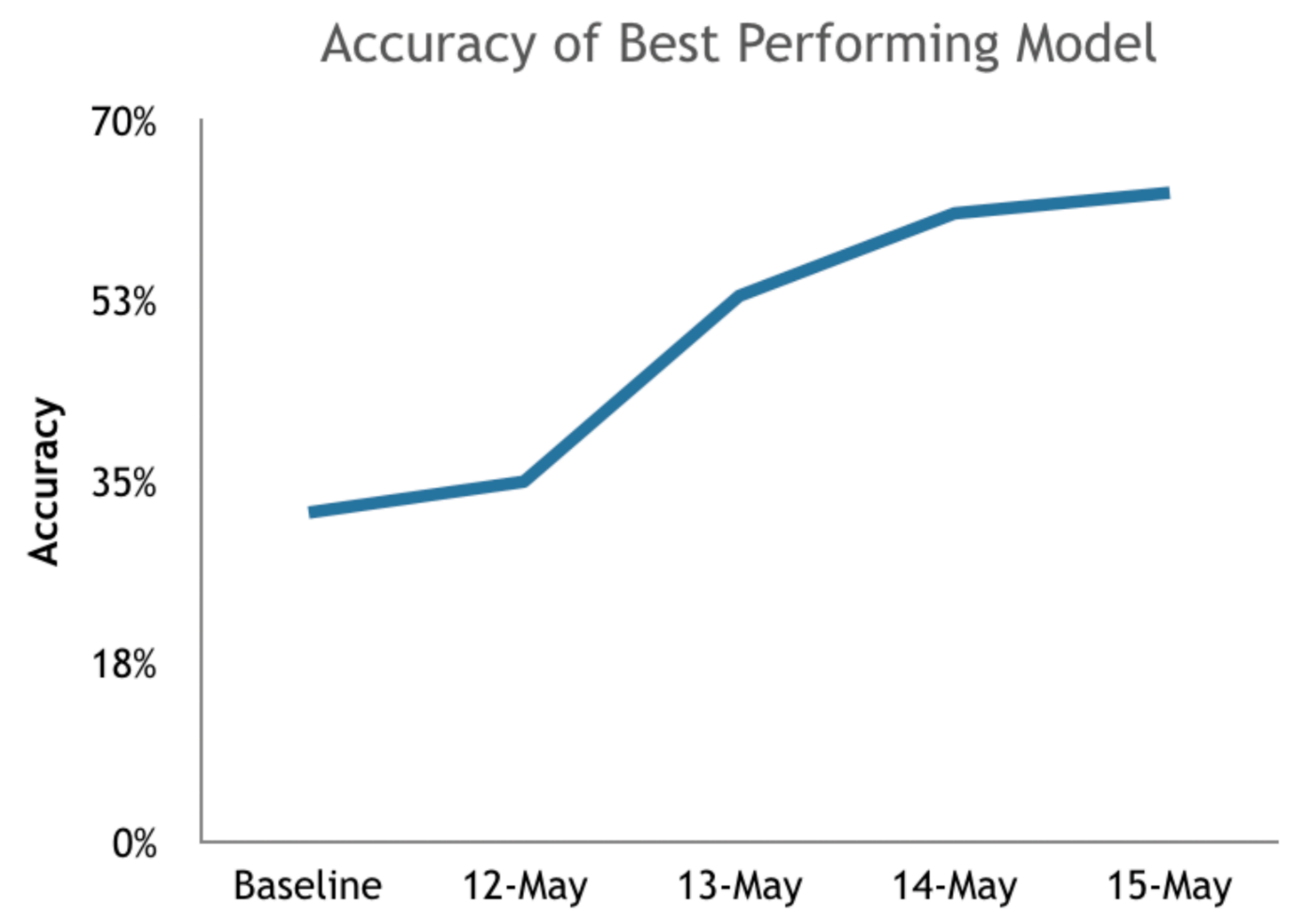

Accuracy improvement in three months

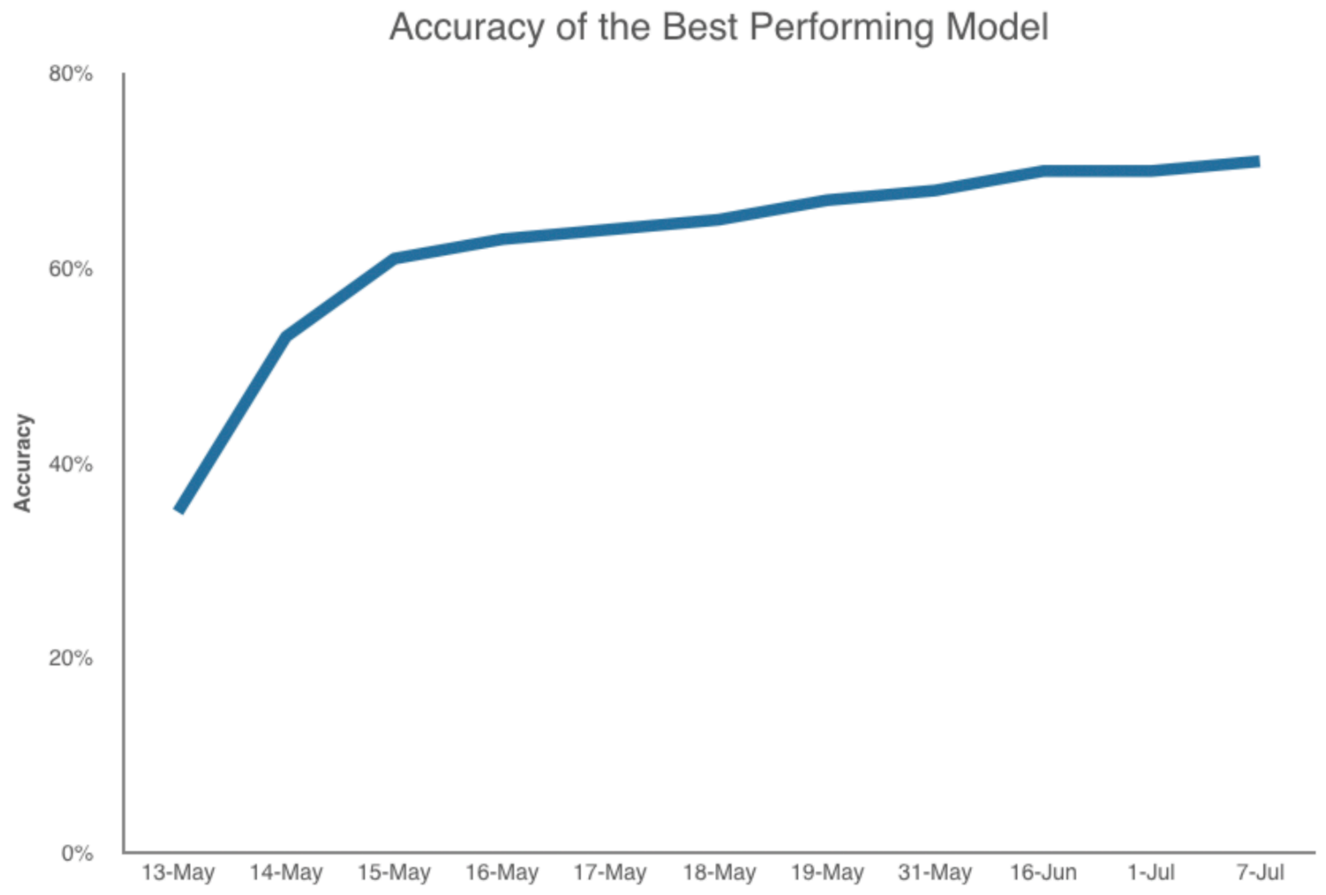

Effort

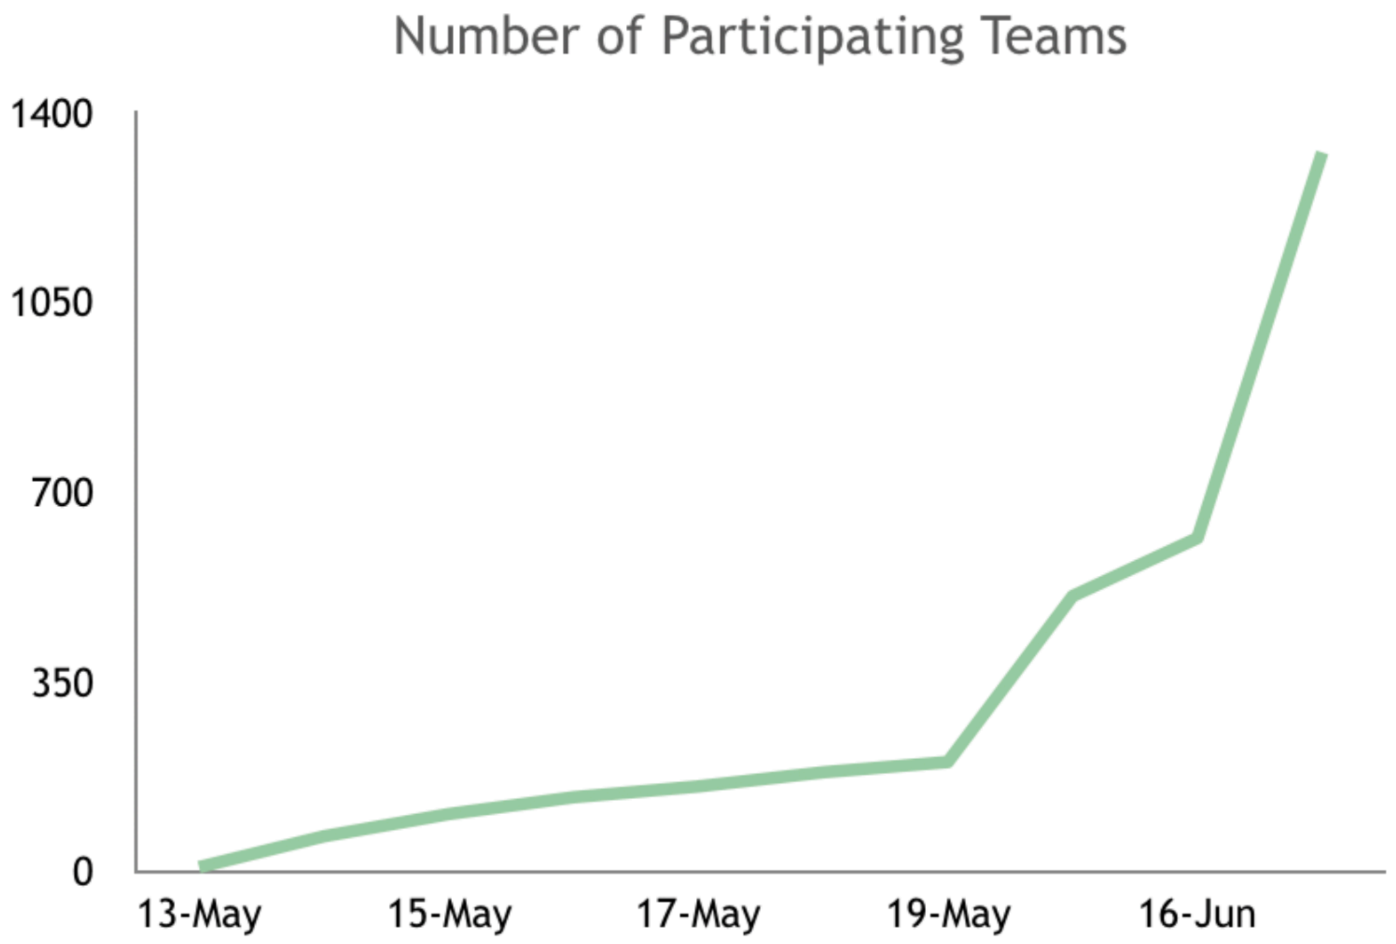

https://medium.com/@l2k/why-are-machine-learning-projects-so-hard-to-manage-8e9b9cf49641

Managing ML teams is challenging

- It's hard to tell in advance how hard or easy something is
- ML progress is nonlinear
  - Very common for projects to stall for weeks or longer
  - In early stages, difficult to plan project because unclear what will work
  - As a result, estimating project timelines is extremely difficult
  - I.e., production ML is still somewhere between "research" and "engineering"

Managing ML teams is challenging

- It’s hard to tell in advance how hard or easy something is
- ML progress is nonlinear
- There are cultural gaps between research and engineering
    - Different values, backgrounds, goals, norms
    - In toxic cultures, the two sides often don’t value one another

Managing ML teams is challenging

- It’s hard to tell in advance how hard or easy something is
- ML progress is nonlinear
- There are cultural gaps between research and engineering
- Leaders often don’t understand it

Outline

- The AI talent gap
- ML-Related roles
- ML team structures
- **The hiring process**
- Exam for this course

Where to look for ML / data science jobs?

- Apply directly to companies you’re interested in
- On-campus recruiting (for those in school)
- LinkedIn recruiters (for more experienced)
- Recruiting fairs at ML conferences (NIPS / ICML / ICLR / etc.)
- This course!

What to expect in the interview process?

- Much less well-defined than software engineering interviews
- Common types of assessments:
  - Background & culture fit
  - Whiteboard coding (similar to SWE interviews)
  - Pair coding (similar to SWE interviews)
  - Pair debugging (often ML-specific code)
  - Math puzzles (e.g., involving linear algebra)
  - Take-home ML project
  - Applied ML (e.g., explain how you'd solve this problem with ML)
  - Previous ML projects (e.g., probing on what you tried, why things did / didn't work)
  - ML theory (e.g., bias-variance tradeoff, overfitting, underfitting, understanding of specific algorithms)

How to prepare for the interview?

- Prepare for a general SWE interview (e.g., “Cracking the Coding Interview”)
- Prepare to talk in detail about your past ML projects (remember details, prepare to talk about tradeoffs and decisions you made)
- Review how basic ML algorithms work (linear / logistic regression, nearest neighbor, decision trees, k-means, MLPs, ConvNets, recurrent nets, etc)
- Review ML theory
- Think about the problems the company you’re interviewing with may face and what ML techniques may apply to them

Outline

- The AI talent gap
- ML-Related roles
- ML team structures
- The hiring process
- **Exam for this course**

Overview of the exam

- Designed to help you prepare for ML engineering interviews
- Take on your own time

Outline

- **Problem setup** (e.g., for a particular problem, what data would you look for, what model would you choose, and what metric would you optimize?)
- **Algorithm knowledge** (e.g., how does an LSTM work?)
- **Understanding of commonly conventions** (e.g., what problem with RNNs does an LSTM solve?)
- **ML theory** (e.g., understanding bias and variance)
- **ML debugging** (e.g., what's likely to cause the bug here?)

Example question 1

- Why does the Residual Block in the ResNet architecture help with the vanishing gradient problem?

Example question 2

- Suppose you see the following learning curve when training on a single batch. Which of the following could be the cause? (Select all that apply)
  - Shuffled labels
  - Learning rate too low
  - Learning rate too high
  - Numerical instability
  - Too big a model

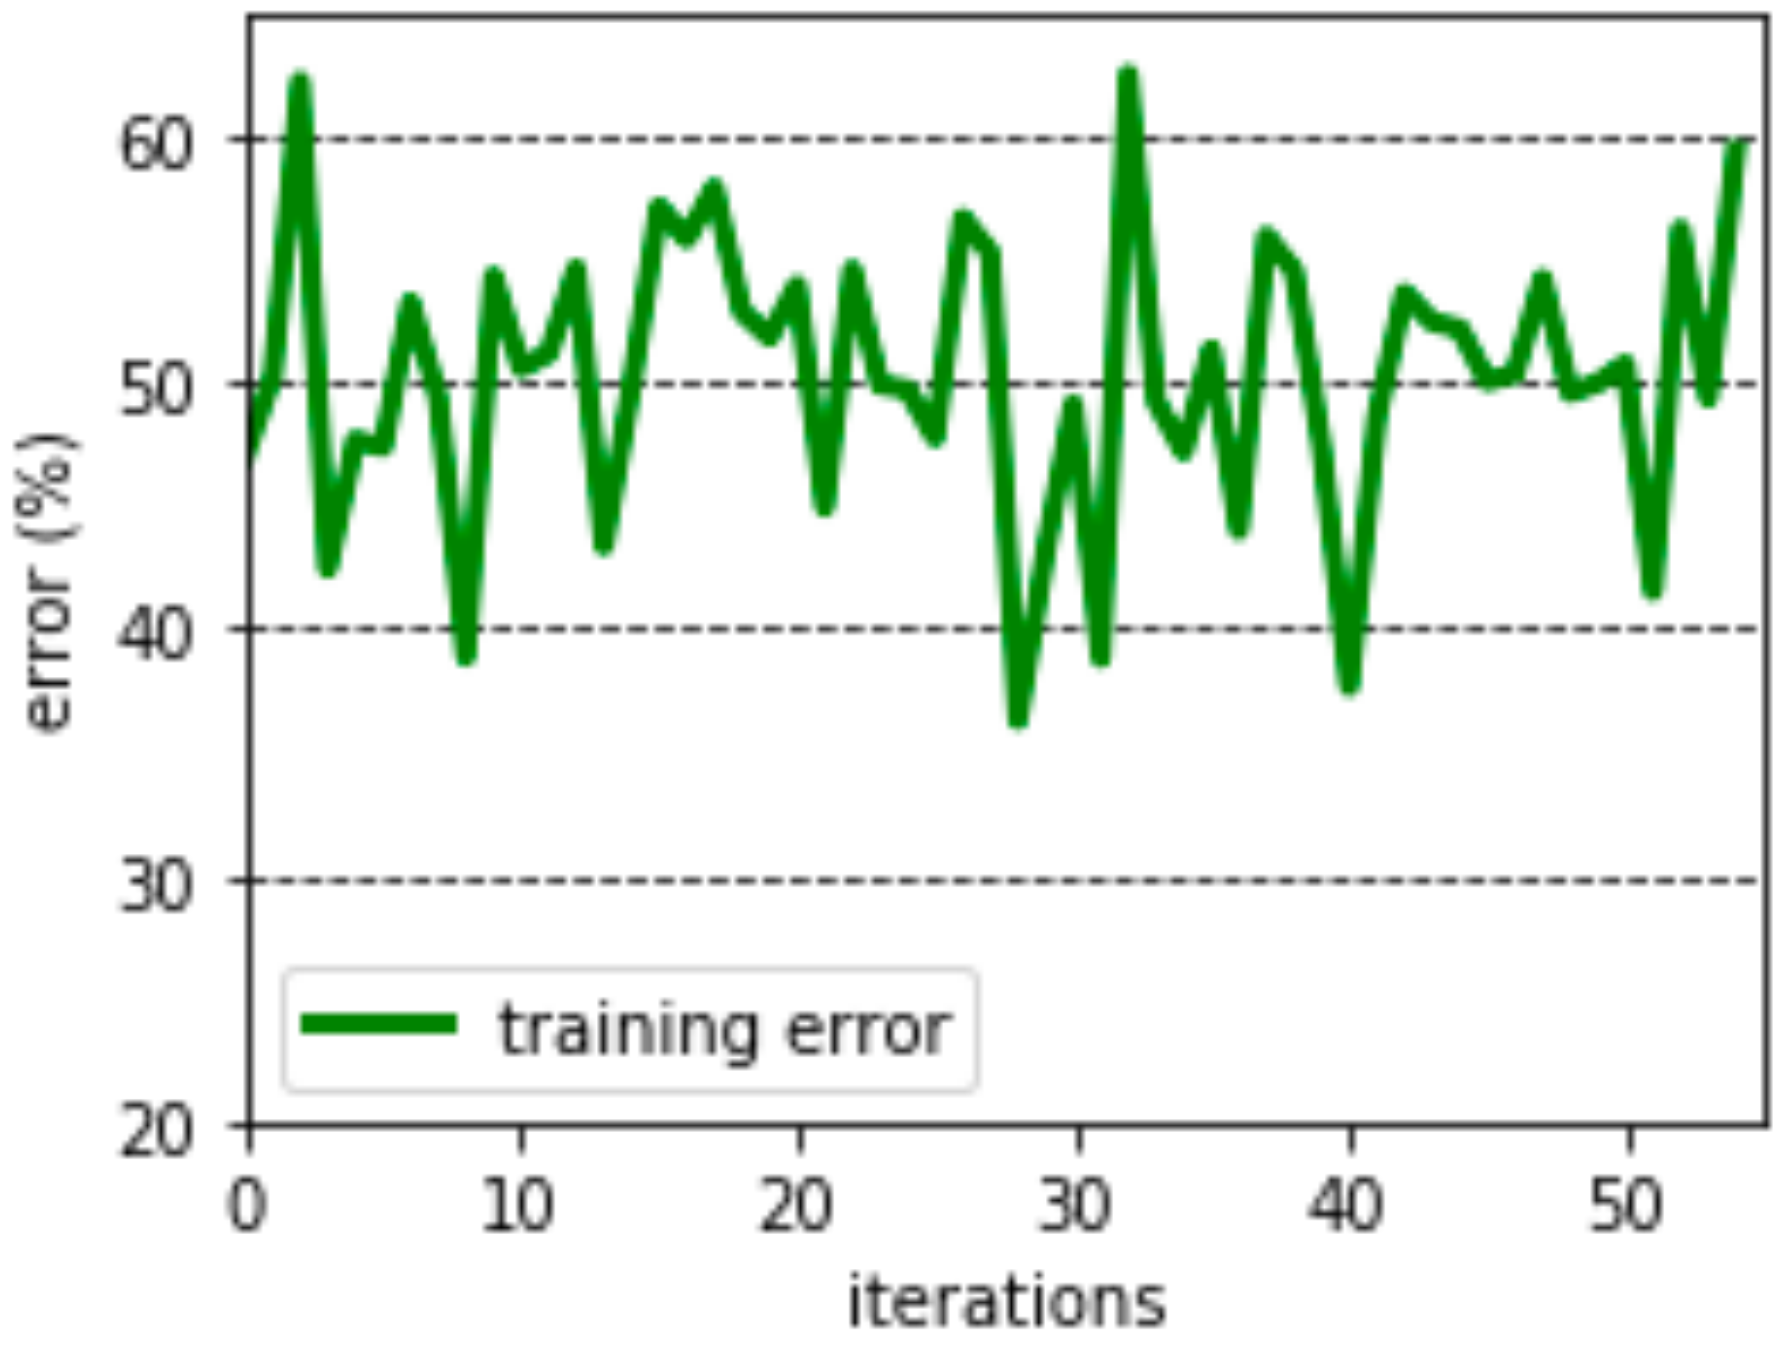

Example question 3

- **For each of the following prediction tasks, select the loss function that is best suited for it.**
- Predict sale price of a house listed for sale.
- Predict whether an image contains pornography or not.
- Predict the category of an email: personal, promotional, reminder, or spam.
- Predict whether a voice sample belongs to the owner of the phone.

(a) Mean Squared Error
(b) Categorical cross-entropy
(c) Binary cross-entropy
(d) CTC loss
(e) GAN loss

Questions?