CS230: Lecture 10

Class wrap-up

Andrew Ng, Kian Katanforoosh

Outline

I. Continuous Career Development

II. Choosing an AI Career

III. What's next after CS230?

IV. Closing remarks

The rapid progress in AI affects the half-life of skills.

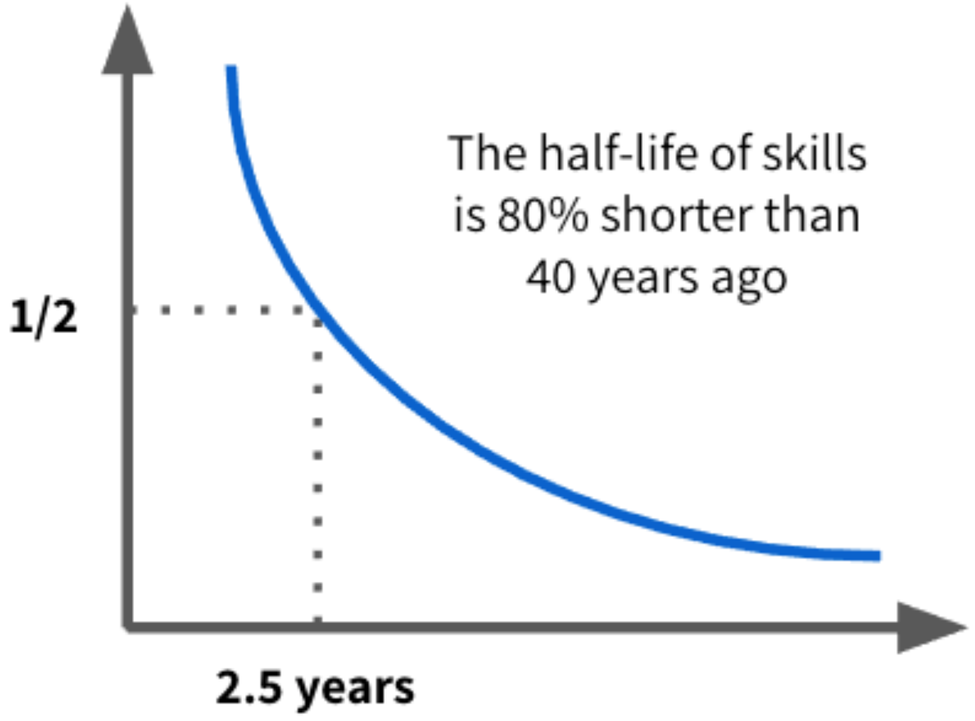

2021-22 — **Large Language Models & Latent Diffusion Models**

Massive leaps forward in large language models and multimodal models that move between image and language.

2020 — **Data-Centric AI**

Building AI systems with a focus on ensuring that the data clearly conveys what the AI must learn.

2019 — **ML Ops Tools**

Emergence of a set of tools and practices aiming to deploy and maintain ML models in production reliably and efficiently

2018 — **Responsible AI**

Discussions on AI fairness, safety, reliability, transparency, and privacy become more prominent in the community.

2017 — **Transformer Architecture**

Attention is all you need and the rise of the transformer architecture improved the performance of state-of-the-art systems in various tasks.

2016 — **Generative Adversarial Networks**

Rapid adoption of GANS after the seminal paper from Goodfellow et al.

2015 — **Modern Deep Learning Frameworks**

Google's TensorFlow and Facebook's PyTorch taking the lead over frameworks like Lua, Torch, Caffe, Lasagna, and Chainer.

How to keep up with the latest and greatest?

Precision Upskilling Cycle

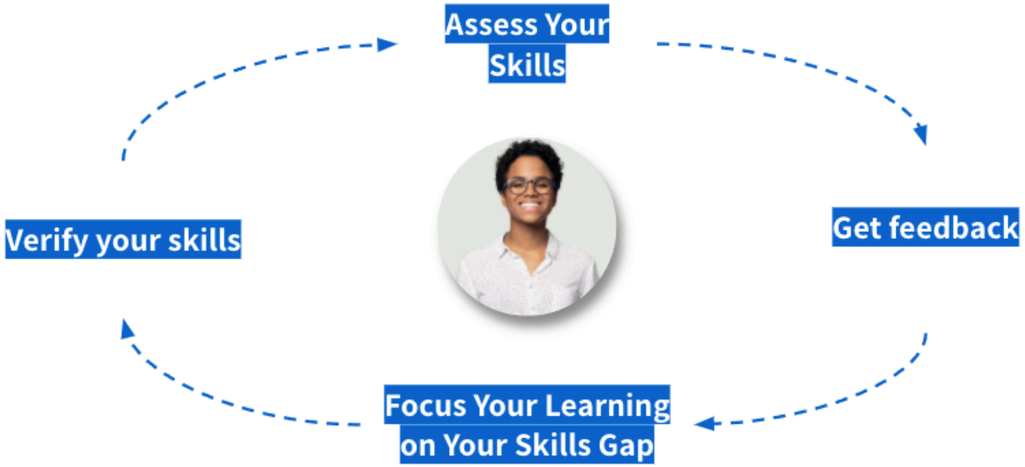

Assessments available at
https://skills.workera.ai/for-individuals

Tune your “antenna”

1. **Follow top research and applied scientists on Twitter** to stay apprised of new ML publications or blogs. Beware that you'll need to sort through the noise. For example, top institutions tend to get more attention than others – it is not always justified. If you're looking for people to follow, see the lists we follow on Twitter (@AndrewYNg, @kiankatan).

2. **Top ML & DS conferences such as NeurIPS, CVPR, ICLR, ICML, or KDD publish their list of award winning papers**. These are peer-reviewed and have excited multiple people from the community. You'd get less real-time info than on Twitter, but also less noise. It's also a more academic view. I like to print papers and skim through them (5 min each max) focusing on abstract, methodology, and figures. If I'm interested, I'll spend more time on the paper.

3. **Reddit** is my favorite place to **discuss ML & DS research papers** or industry innovations. That's where you can go deeper into your understanding of a new approach by asking questions and scrolling through answers. The point of views are more critical. I recommend the following subreddits: r/MachineLearning, r/datascience, r/deeplearning, r/datasets, r/artificial, r/LanguageTechnology, or r/reinforcementlearning.

Outline

I. Continuous Career Development

II. Choosing an AI Career

III. What's next after CS230?

IV. Closing remarks

Backend Engineer	Frontend Engineer	Deployment Engineer	DevOps Engineer
Site Reliability Engineer	QA Engineer	Mobile Engineer	. . .

Roles in software teams are better understood than roles in AI teams.
AI as a job category is still nascent!

Kian Katanforoosh

Thus, companies write fuzzy AI job descriptions with confusing titles.

Parents: So, you're applying to a machine learning job?

Me: Trying to…

Machine Learning Engineer
Snorkel AI
Palo Alto, CA
via Lever

Machine Learning Scientist
Amazon
Sunnyvale, CA
via LinkedIn

Machine Learning Researcher
Samsung Research America
Mountain View, CA (+1 other)
via Greenhouse

Data Scientist
Waymo
Mountain View, CA
via Glassdoor

Data Analyst
SnowFlake
San Mateo, CA
via Snowflake Careers

Software Engineer
Electronic Arts Inc.
Redwood City, CA
via Glassdoor

Data Science Engineering, Machine Learning
Salesforce
Palo Alto, CA
via Glassdoor

Data Science Engineer, Recent Graduate
eBay
San Jose, CA
via EBay - EBay Inc.

Machine Learning Research Scientist
Capital One
San Francisco, CA
via Capital One Careers

Senior Machine Learning Analyst
Google
Sunnyvale, CA
via Google Careers

Research Data Analyst
Stanford University
Stanford, CA (+1 other)
via Inside Higher Ed Careers

Machine Learning Research Engineer
Embark Trucks, Inc.
San Francisco, CA
via ZipRecruiter

Software Engineer, Machine Learning, Research
Waymo
Mountain View, CA
via Waymo

Data Scientist / Machine Learning Specialist
Genuent Global, LLC
London, CA (+1 other)
via Dice

Senior Product Data Scientist
Bill.com
Palo Alto, CA
via Lever

Product Analyst, Data Science
Google
Mountain View, CA
via Google Careers

Senior Business Data Analyst
Intuit Corp
Mountain View, CA
via Intuit Jobs

Full Stack Data Engineer
Flexion
San Ramon, CA
via WFH Jobs

Full Stack Data Scientist, Supply DS
Wish
San Francisco, CA
via SmartRecruiters Jobs

AI/ML Full Stack Engineer (Quality Engineering)
Apple
Cupertino, CA
via WDTN Jobs

Software Engineer / Data Scientist, Fleet Analytics
Tesla
Palo Alto, CA
via LinkedIn

AI Deep Learning Software Engineer
Intel
Santa Clara, CA
via TOPBOTS

Machine Learning Modeler, Cash App
Square
San Francisco, CA
via SmartRecruiters Jobs

Companies are confused as much as applicants.

Kian Katanforoosh

It makes it difficult for job seekers to apply to the job that’s right for them.

How can applicants find the right role and stand out?

Identify your career pathway and benchmark your skills regularly.

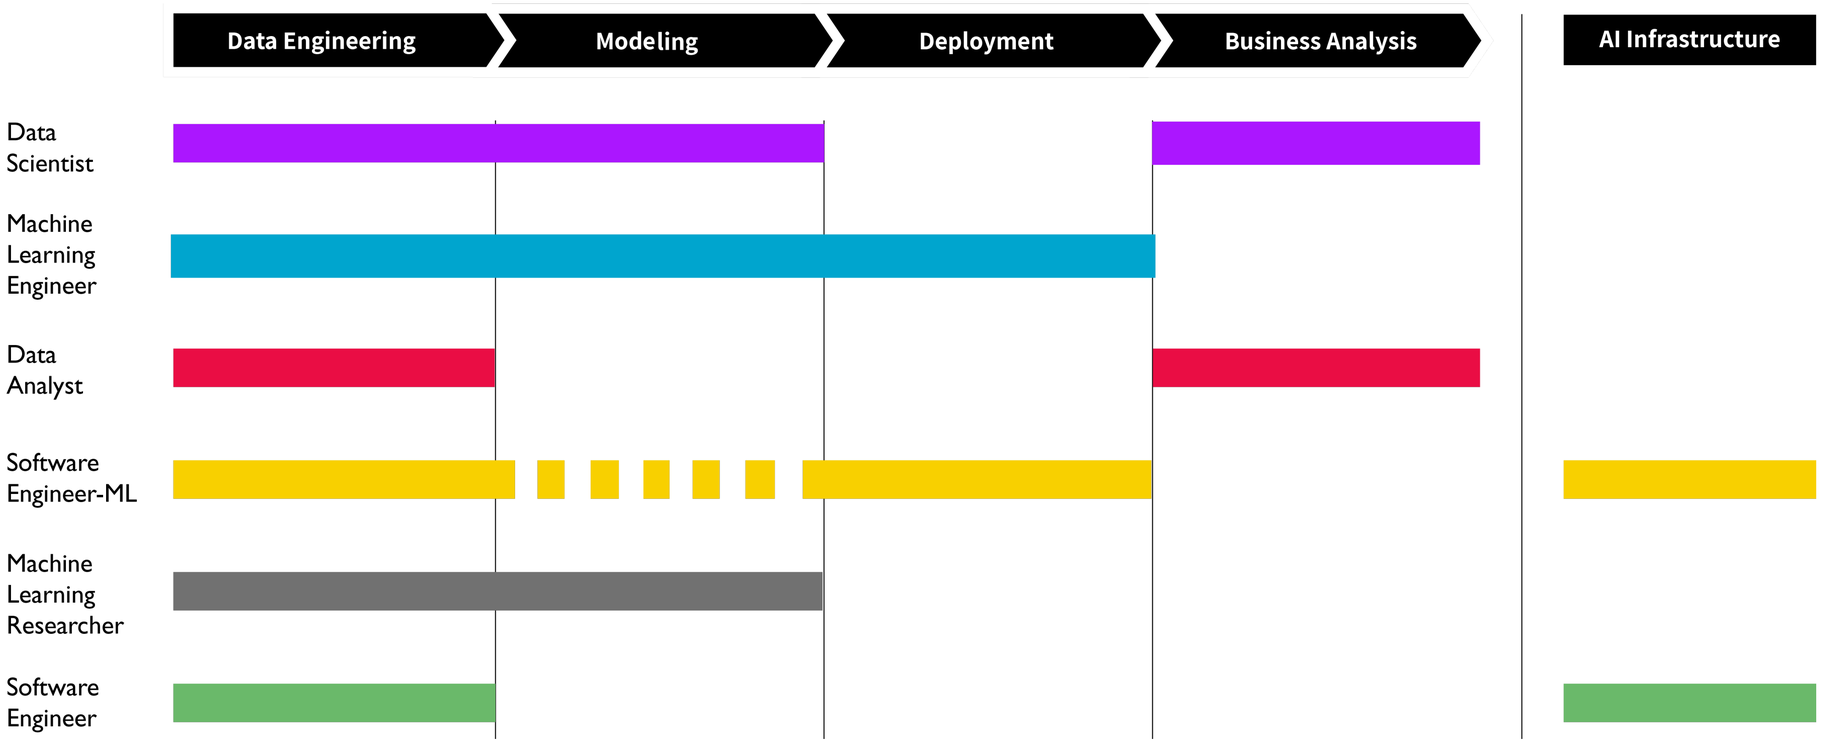

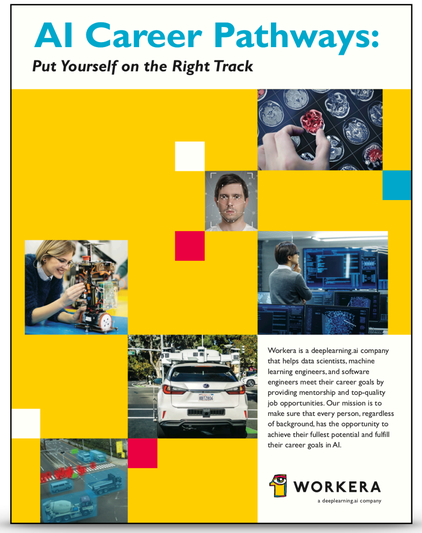

AI career pathways report (2020)

Kian Katanforoosh

AI is more than modeling.

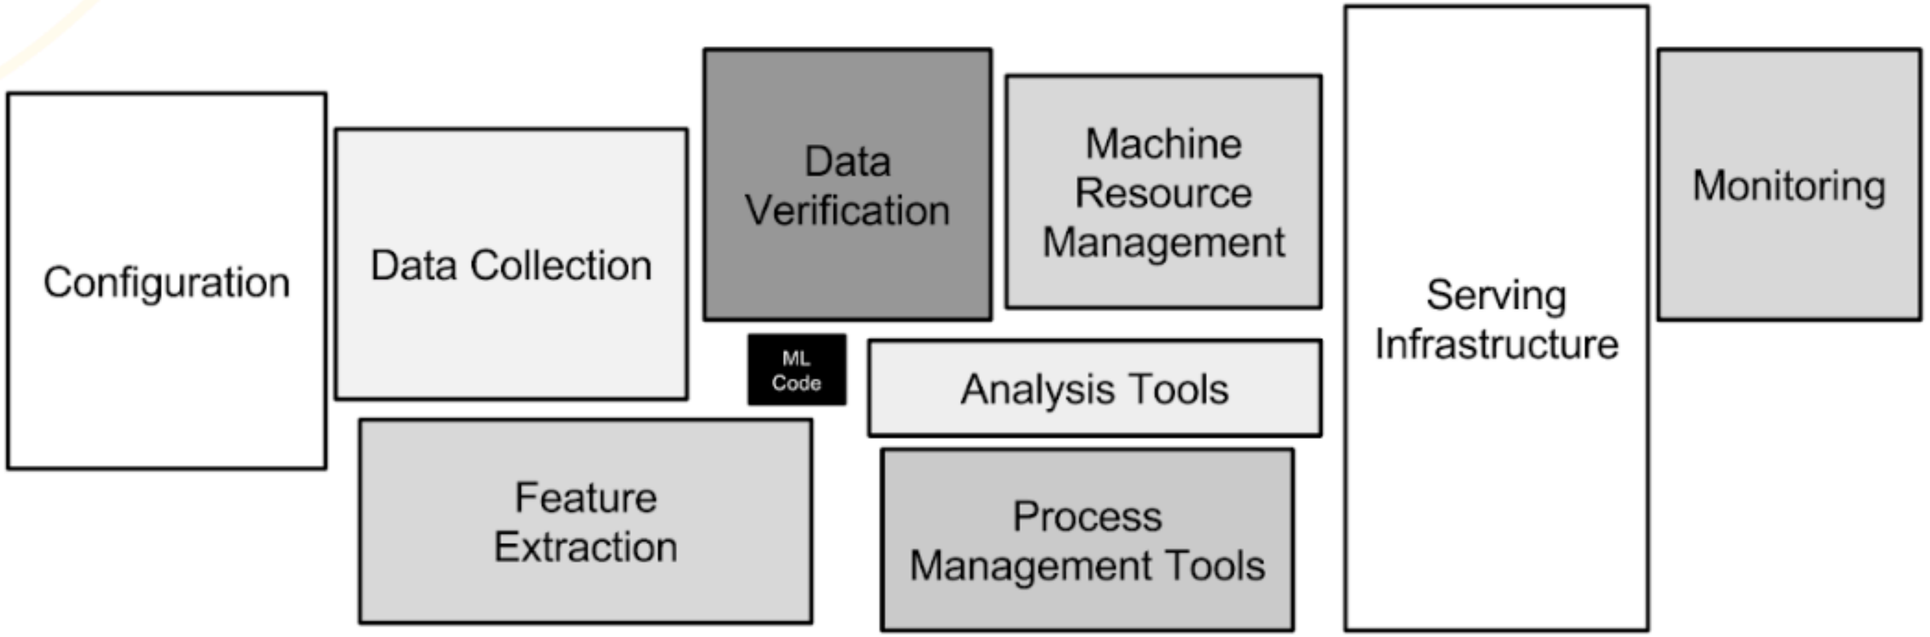

Figure 1: Only a small fraction of real-world ML systems is composed of the ML code, as shown by the small black box in the middle. The required surrounding infrastructure is vast and complex.

[Source: Sculley et. al, Hidden Technical Debt in Machine Learning Systems]

Like Software, AI is becoming more specialized as a discipline

The π-shape Data Scientist

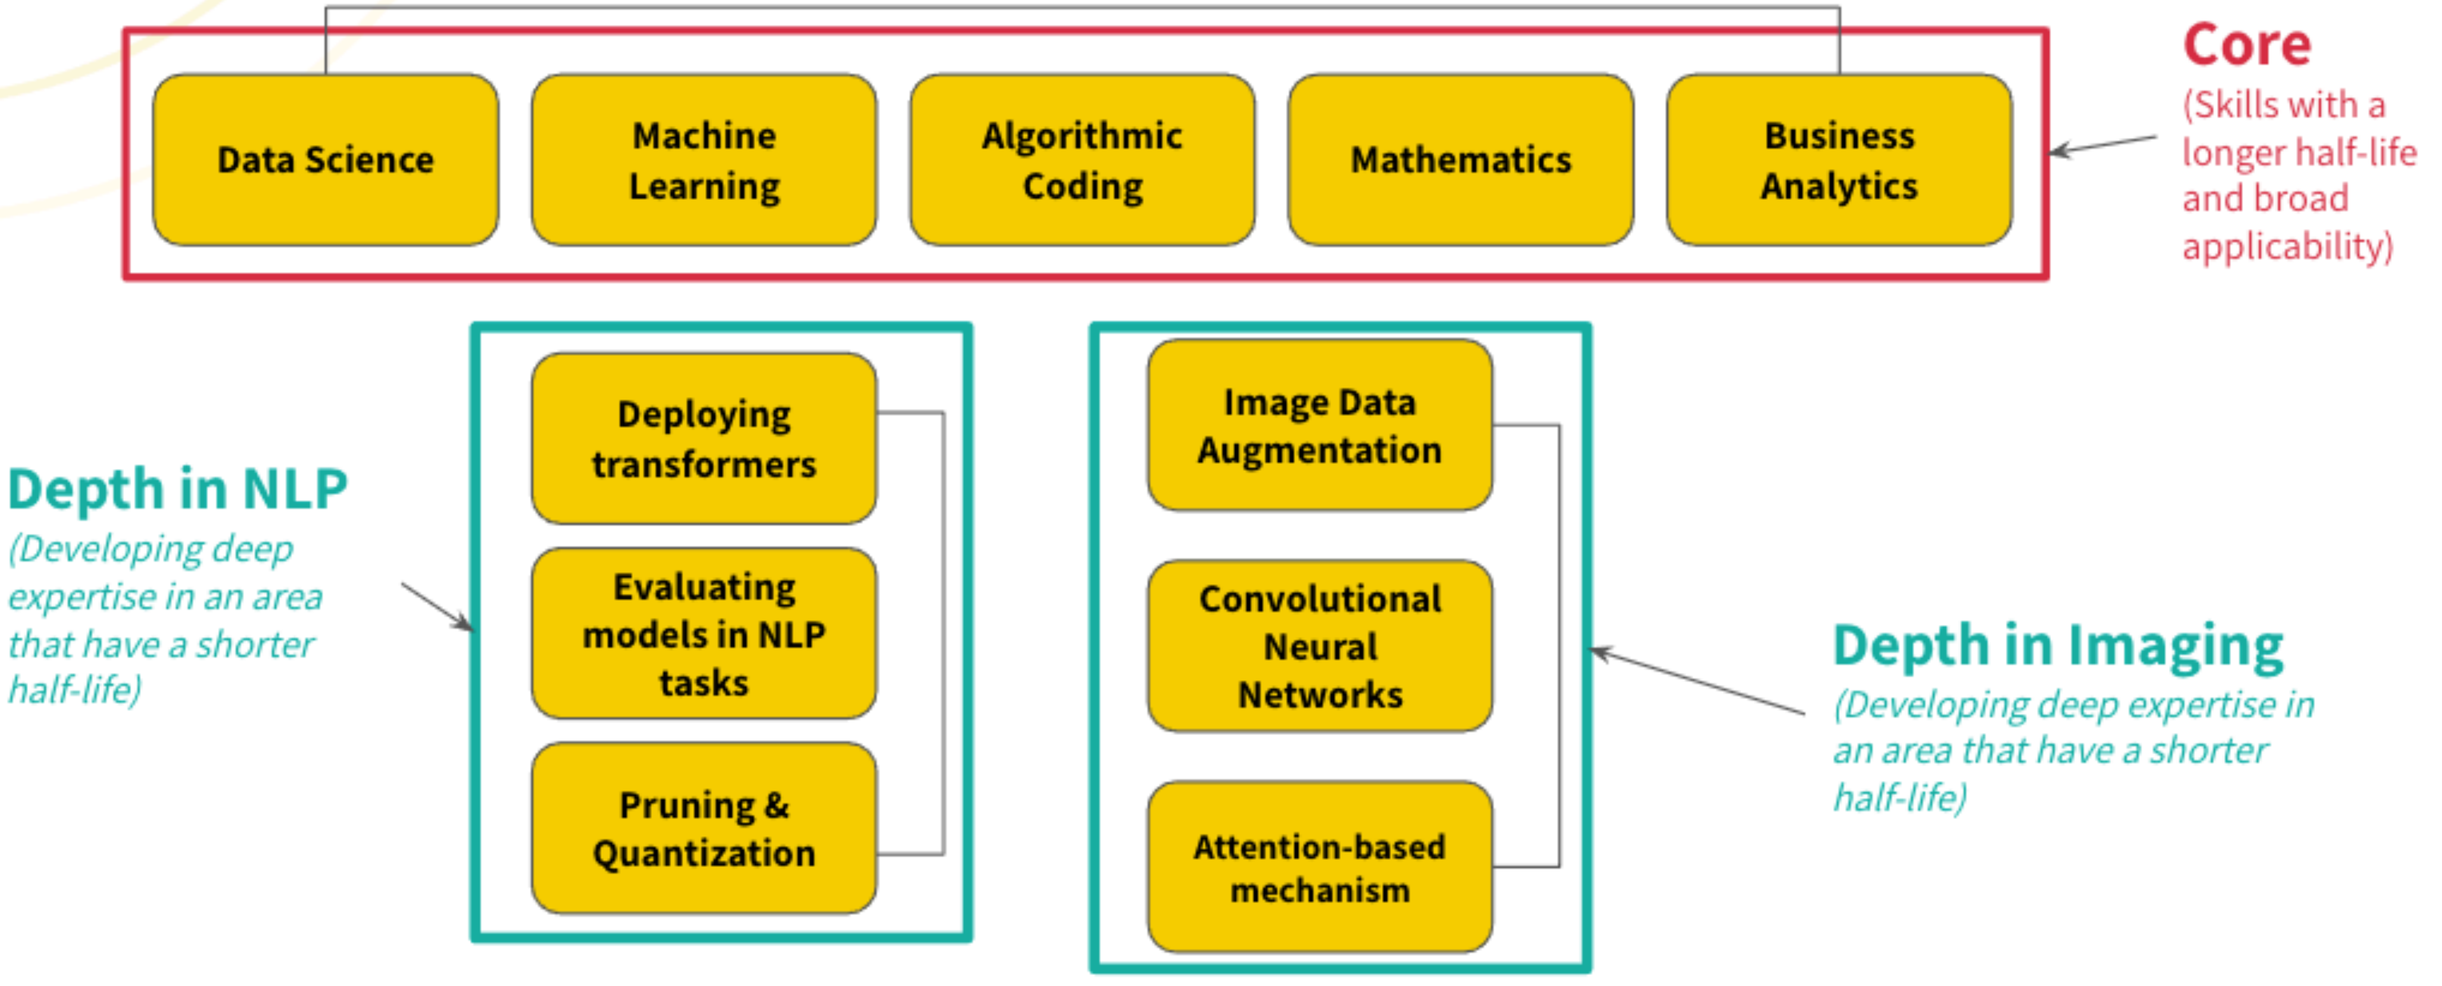

AI+X: Develop your AI competency in conjunction with your subject-matter expertise.

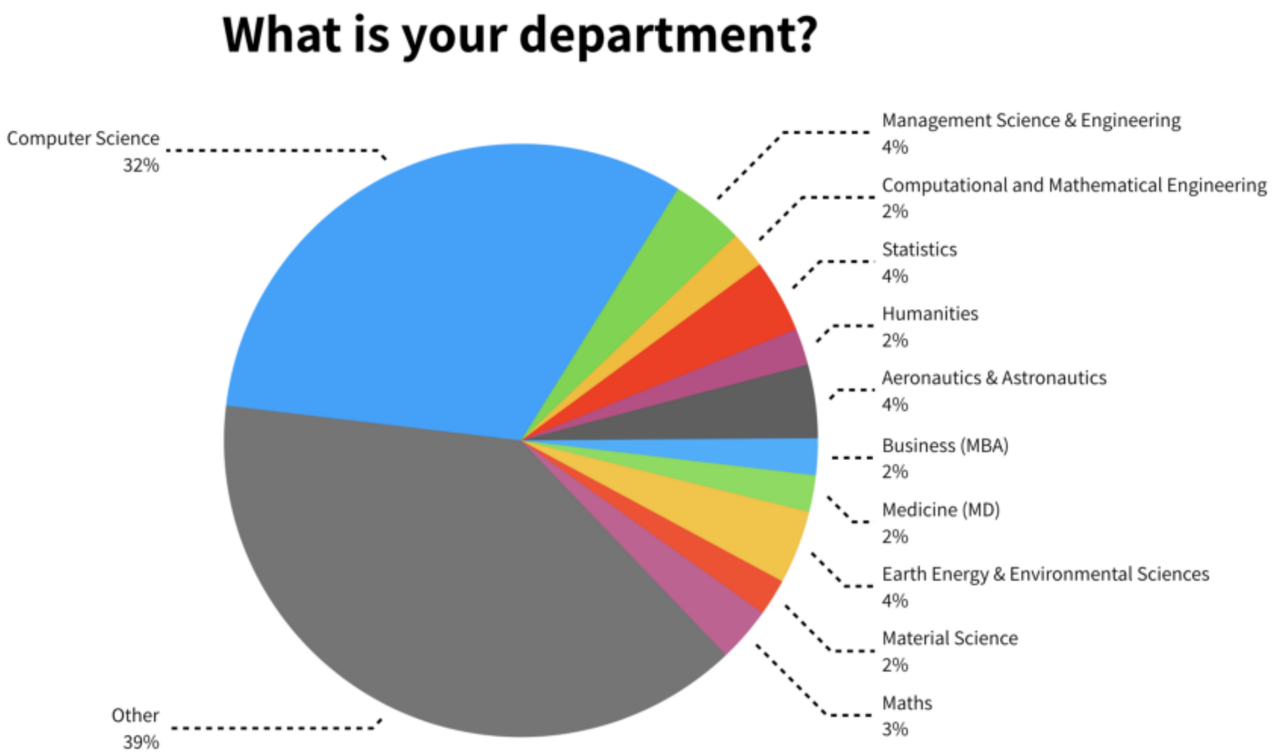

Kian Katanforoosh

CS230

Predicting Ionic Liquid Materials Properties from Chemical Structure

Tyler Quill, Shayta Roy and Yaakov Tuchman
Department of Materials Science & Engineering
Stanford University

CS230

Learning Power Flow Mappings for Power Grid Simulation

Lily Buechler
Department of Mechanical Engineering
Stanford University

CS230

Deep Learning for Local Ancestry Inference

Jan Sokol
Biomedical Informatics Training Program
Stanford University

Matthew Aguirre
Biomedical Informatics Training Program
Stanford University

CS230

CADDYIAN - Diver Gesture Language Classification

Veronica Peng
Dept. of Computer Science

Xi Yu
Dept. of Bioengineering

Wenxi Zhao
Dept. of Civil & Env. Engineering

CS230

Predicting Parkinson's disease behavioral state from neural and kinematic data

Marissa Lee, Johanna O'Day, Kirsten Seagers
Department of Computer Science
Stanford University

And many more..

Kian Katanforoosh

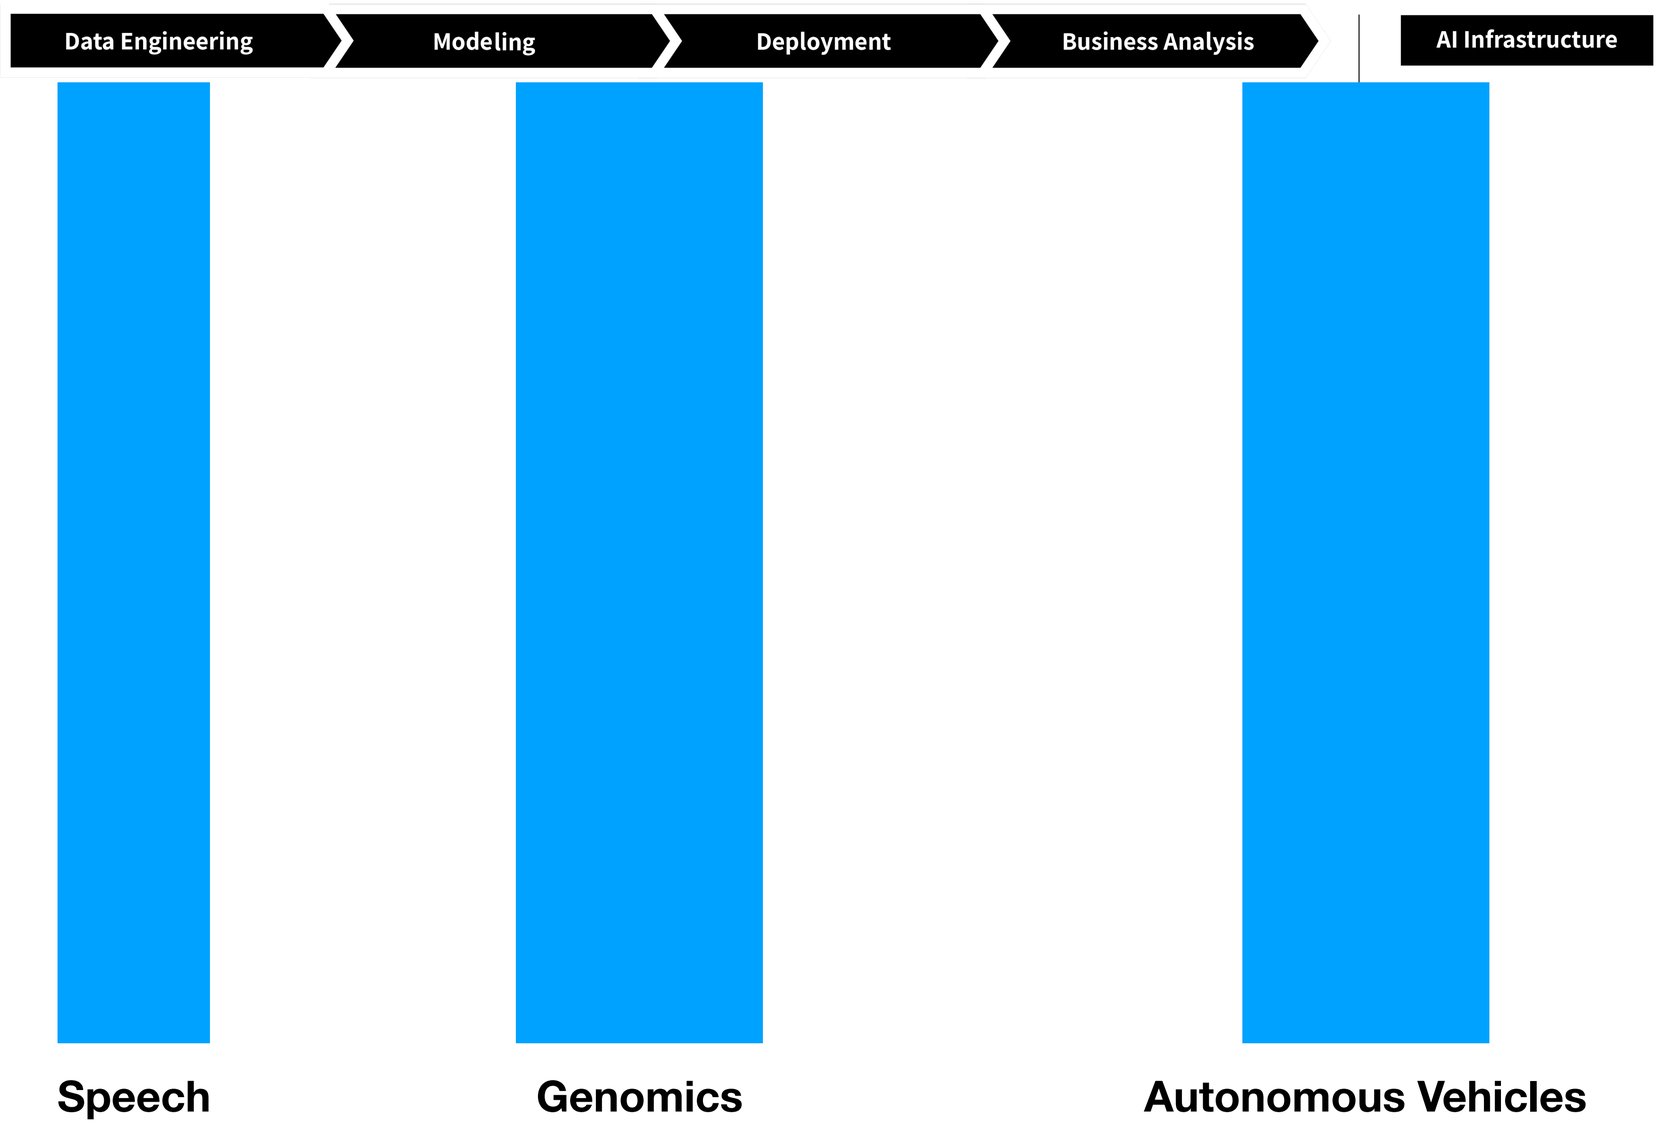

**“I’m a Data Scientist,
with expertise in Speech”**

**“I’m a MLE,
with expertise in Genomics”**

**“I’m a Data Analyst,
with expertise in Marketing”**

**“I’m a Software Engineer,
with expertise in
Manufacturing Automation”**

Kian Katanforoosh

If you’re transitioning to AI, be an AI+X.

Kian Katanforoosh

It’s also a great time to start a company!

Watch for those:

1. **Market** dynamics / Business opportunity: Funding availability?
2. Tech/**Product**/Defensibility: Proprietary data? Proprietary model?
3. **Team**: Timing?

Outline

I. Continuous Career Development

II. Choosing an AI Career

III. What's next after CS230?

IV. Closing remarks

What's next?

Classes at Stanford

Natural Language Processing

CS 124: From Languages to Information (LINGUIST 180, LINGUIST 280)
CS 224N: Natural Language Processing with Deep Learning (LINGUIST 284)
CS 224U: Natural Language Understanding (LINGUIST 188, LINGUIST 288)
CS 276: Information Retrieval and Web Search (LINGUIST 286)

Computer Vision

CS 131: Computer Vision: Foundations and Applications
CS 205L: Continuous Mathematical Methods with an Emphasis on Machine Learning
CS 231N: Convolutional Neural Networks for Visual Recognition
CS 348K: Visual Computing Systems

Others:

CS 273B: Deep Learning in Genomics and Biomedicine (BIODS 237, BIOMEDIN 273B, GENE 236)
CS 236: Deep Generative Models
CS 228: Probabilistic Graphical Models: Principles and Techniques
CS 337: AI-Assisted Care (MED 277)
CS 229: Machine Learning (STATS 229)
CS 129: Applied Machine Learning
CS 234: Reinforcement Learning
CS 221: Artificial Intelligence: Principles and Techniques
CS 217: Hardware Accelerators for Machine Learning
CS 329S: Machine Learning Systems Design

AI for Healthcare Bootcamp

AI for Climate Bootcamp

Outline

I. Continuous Career Development

II. Choosing an AI Career

III. What's next after CS230?

IV. Closing remarks

Closing remarks

1. *You can **feel proud**. Despite the difficult circumstances, you've built a project that you can carry with you after the class.*

2. *If you enjoyed working on your project, **keep developing it after the class**.*

3. *We encourage you to **review other students' projects** which will be shared on the CS230 website.*

4. ***Keep in touch**. You're officially a CS230 graduate, and we hope that you've built a relationship with your team members or a teaching staff member.*