

COMP2200/COMP6200 Lecture 1b – Data Science Jobs

Greg Baker

23rd February 2026



Last Semester's Motivational Slide

Job ad for UnslothAI 20th February 2025

← Post

 **Daniel Han** ✓
@danielhanchen



We made 5 challenges and if you score 47 points we'll offer you **\$500K/year** + equity to join us at  @UnslothAI!

No experience or PhD needed.

\$400K - \$500K/yr: Founding Engineer (47 points)

\$250K - \$300K/yr: ML Engineer (32 points)

Challenges:

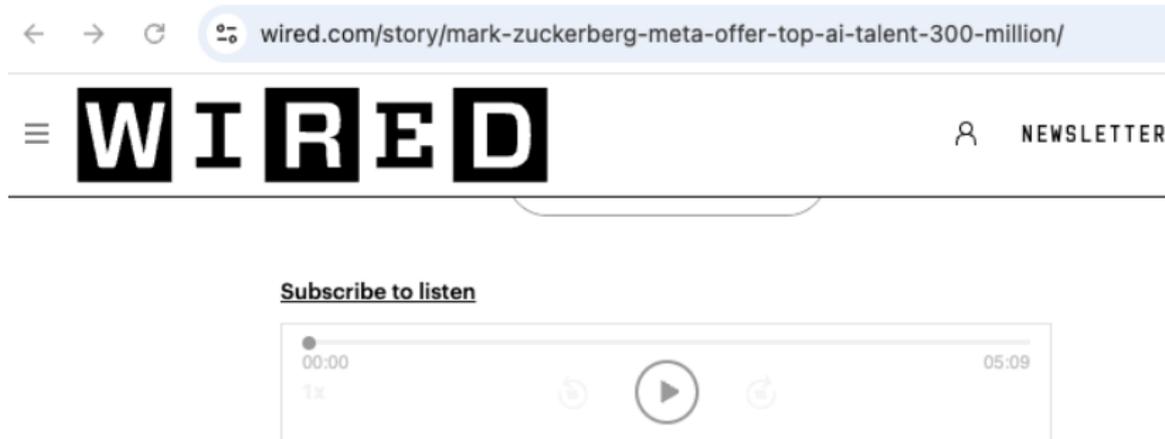
1. Convert nf4 / BnB 4bit to Triton
2. Make FSDP2 work with QLoRA
3. Remove graph breaks in torch.compile
4. Help solve Unsloth issues!
5. Memory Efficient Backprop

If you have any questions about the challenges, please feel free to ask!
We're looking for people to help push Unsloth forward - so come join us to democratize AI further!

Our past work includes:

1. 1.58bit DeepSeek R1 GGUFs: x.com/UnslothAI/stat...
2. GRPO with Llama 3.1 8B in a Colab: x.com/UnslothAI/stat...
3. Gemma bug fixes: x.com/danielhanchen/...
4. Gradient accumulation bug fixes: x.com/danielhanchen/...

This Semester's Motivational Slide



The screenshot shows a web browser displaying a Wired article. The address bar contains the URL: [wired.com/story/mark-zuckerberg-meta-offer-top-ai-talent-300-million/](https://www.wired.com/story/mark-zuckerberg-meta-offer-top-ai-talent-300-million/). The Wired logo is prominently displayed in the center. To the right, there is a user profile icon and a 'NEWSLETTER' button. Below the logo, there is a section titled 'Subscribe to listen' which contains a video player. The video player shows a progress bar from 00:00 to 05:09, with a play button in the center and skip back and forward buttons on either side.

AS MARK ZUCKERBERG staffs up Meta's new superintelligence lab, he's offered top tier research talent pay packages of up to \$300 million over four years, with more than \$100 million in total compensation for the first year, WIRED has learned.

Job titles you might end up with

Data analyst \$ Visualise data, tell stories. Often entry level

Data scientist \$\$ Manage experiments, analyse results, make and train predictive models

Data engineer \$\$\$ Automate training of models, manage data pipelines, write extract-transform-load scripts

ML engineer \$\$\$\$ Train new language and vision models

Data Analyst — Entry level

`https://www.seek.com.au/data-analyst-jobs`

Skills needed

- A pulse
- Excel
- SQL (usually)
- A business intelligence tool, e.g. Power BI, Tableau, Qlik (banking), QuickSight (startups)
 - Power BI is free on Windows
 - Tableau is free for students; note academic ambassador program
- A portfolio of interesting visualisations
- (Recently) Using ChatGPT to create visualisations

Data scientist — Entry-to-mid level

<https://www.seek.com.au/data-scientist-jobs>

Skills needed

- Python or R (Python is more common)
- Statistics and experimental design
- Machine learning libraries (e.g. scikit-learn[‡], TensorFlow, PyTorch)
- SQL and data wrangling (Pandas[‡], dplyr)
- Data visualisation (Matplotlib[‡], Seaborn, ggplot2)
- Domain knowledge – understanding the problem space (e.g. finance, health, marketing)

[‡] Topics we'll cover in COMP2200/COMP6200

Data Engineer — Senior

<https://www.seek.com.au/data-engineer-jobs>

Skills needed

- SQL (strong proficiency in query optimisation and data warehousing)
- Python or Scala (for data pipelines; PySpark is common)
- Cloud platforms (AWS, GCP, Azure)
 - Common tools: AWS Glue, BigQuery, Snowflake, Databricks
- Data pipeline orchestration (e.g. Apache Airflow, Prefect)
- Extract-Transform-Load (ETL) and data wrangling
- Containerisation and deployment (Docker, Kubernetes)
- CI/CD for data workflows (Git, Terraform, dbt)

ML Engineer — Specialised

<https://www.seek.com.au/machine-learning-engineer-jobs>

Skills needed

- Python (strong proficiency) and sometimes C++ for performance-critical code
- Machine learning frameworks: TensorFlow, PyTorch, JAX
- Deep learning models (transformers, CNNs, RNNs)
- Model deployment and optimisation (ONNX, TensorRT)
- MLOps (monitoring and maintaining ML models in production)
- Cloud services (AWS SageMaker, GCP Vertex AI, Azure ML)
- Data pipeline integration (e.g., Apache Kafka, Airflow)
- Experiment tracking (MLflow, Weights & Biases)
- High-performance computing (GPUs, TPUs, distributed training)

Elephant in the room



Will there be any junior roles left, or will AI take them all?

Elephant in the room



Will there be any junior roles left, or will AI take them all?

- Responsibility for reporting will still be around

Elephant in the room



Will there be any junior roles left, or will AI take them all?

- Responsibility for reporting will still be around
- Expect lots of jobs around “test the effect of this prompt change on the last 10,000 events”