


# **The Technical Recruiter's Guide to Artificial Intelligence (AI) and Machine Learning (ML)**

A Quick Guide to Learning Key Technologies in Artificial Intelligence and Machine Learning, Understand and Master AI/ML, MLOps, Connect and Call Your Candidates with Confidence.



**BY OBI OGBANUFE, PHD**

---

**The Technical Recruiter's Guide to Artificial Intelligence (AI) and Machine Learning (ML). Copyright © 2023 Obi Ogbanufe, PhD**

All rights reserved. No part of this book may be used or reproduced by any means, graphic, electronic, or mechanical, including photocopying, recording, taping, or by any information storage retrieval system without the written permission of the publisher.

All trademarks are acknowledged as belonging to their respective companies.

# Audience

In writing this book, I have assumed the following about you.

- You're a technical recruiter or want to become one
- You want to learn more about Artificial Intelligence and Machine Learning technologies
- You're interested in learning how to screen AI/ML candidates
- You want to build your confidence in technologies

# Contacting the Author

The author has done her best to provide accurate and up-to-date information in this eBook, but she cannot guarantee that the information is correct or will fit your situation. The eBook is provided with the understanding that the publisher and the author are not engaged in rendering professional or financial services.

If you have any questions regarding this eBook, need clarification, or simply want to offer your thoughts, please contact me through LinkedIn or send a note to [obi@technicalrecruitingbook.com](mailto:obi@technicalrecruitingbook.com)

This version of the eBook is offered as a FREE eBook resource.  
More books and online courses can be found at [www.simpletechbook.com](http://www.simpletechbook.com)

**The Technical Recruiter's Guide to Artificial Intelligence and Machine Learning Technologies Copyright 2023 by Obi Ogbanufe**

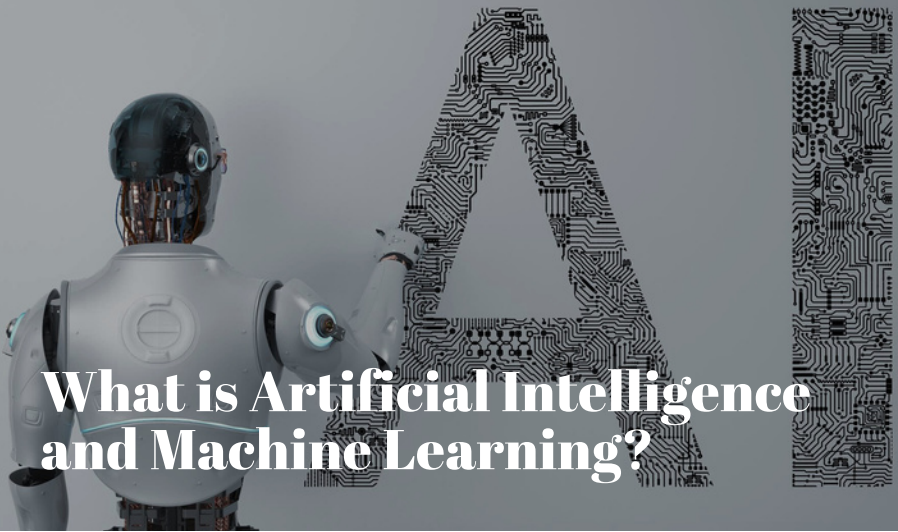
Thank you for your interest in this ebook. I hope you enjoy reading it and putting your knowledge to use.

Regards,  
Obi Ogbanufe, PhD

# Contents

## The Technical Recruiter's Guide to Artificial Intelligence (AI) and Machine Learning (ML)

Audience	03
Contacting the Author	03
Contents	04
What is Artificial Intelligence and Machine Learning?	05
The Machine Learning Skills	07
Machine Learning Frameworks	07
Machine Learning Ops (MLOps)	08
Machine Learning Programming	08
ML Algorithms	11
The AI/ML Technologist Roles	12
Job Roles of AI/ML Technologists	12
The ML Scientist/Engineer	13
AI/ML Job Description and Screening Questions	15
Job Description: ML Engineer/Scientist	15
Conclusion	18
References	19



# What is Artificial Intelligence and Machine Learning?

A simple description of Artificial Intelligence (AI) is that it is a field of study where machines (computers) learn to perform tasks from experience and self-learning. When it receives new information, it adjusts its learning and can perform human-like activities. AI has been around since the 1950s, but its use and development were largely stagnant because of a lack of expertise, data, and funding. AI came back into focus in the 1990s, and its use and applications have continued to grow.

Only two years ago, Deloitte and Forbes projected that more than 80% of public sector organizations would be using artificial intelligence in their decision-making and that 76% of private companies are prioritizing artificial intelligence in their technology budgets, respectively. Well, that projection may very well have been overtaken when the world was introduced to artificial intelligence in November 2022. The world became acutely aware of the power of artificial intelligence with OpenAI, the organization that created the AI-powered ChatBot called ChatGPT (Chat Generative Pre-Trained Transformer).

Artificial intelligence has other subsets embedded in it. Machine learning is a subset of Artificial Intelligence, and Deep learning is a subset of machine learning. Machine learning is the most common way that organizations use artificial intelligence. Machine learning is a data analytics method that automates the building of data models through learning with little help from developers.

Deep Learning, which is a subset of Machine Learning, uses algorithms that allow it to train itself to perform different tasks. These three form a concentric hierarchy, as shown in Figure 1.




Figure 1: Concentric Nature of AI/ML and Deep Learning

AI/ML has become integrated into most of the applications and tools we use today. Think of YouTube, Facebook, LinkedIn, Instagram, and Netflix; many recommendations for connections, friends, products, and videos are based on AI/ML. Think of the recommendations for friends or product recommendations, photo cropping suggestions, product ordering suggestions, and facial recognition, to name a few.

AI/ML can be used to solve many different business problems like, fraud detection, facial recognition, personalized ads, data extraction and processing, customer recommendation, speech recognition, and spam filtering, to name a few.



# The Machine Learning Skills

Below are short descriptions of the most used skill requirements in the artificial intelligence and machine learning development community.

## Machine Learning Frameworks

A framework is like a development template or package that includes many of the tools that developers use for different tasks. A machine learning framework is a software library that includes tools used for building machine learning and artificial intelligence applications (e.g., computer vision, natural language processing). Several of them are open source.




Figure 2: Machine learning frameworks

TensorFlow, PyTorch, Keras, mlpack, scikit-learn, Theano, and Flux are examples of machine learning software libraries or frameworks. See Figure 2. Their primary function is to provide developers with the source code to develop and train a machine learning model. Machine learning frameworks provide the tools to get started with APIs and resources that simplify the development, training, testing, and deploying of machine learning models in different environments (e.g., cloud, on-premise).

# Machine Learning Ops (MLOps)

If you're familiar with DevOps (development and operations), then machine learning operations, a.k.a MLOps, is a bit similar. MLOps just has different processes. MLOps is the practice of performing DevOps for the machine learning domain. MLOps is an iterative lifecycle and includes business requirements gathering, data collection, data pre-processing, model building, model assessment and evaluation, model deployment, and model monitoring. See Figure 3. Although MLOps is a newer term, the processes have been around since 1996 when DaimlerChrysler, SPSS, and NCR established Cross Industry Standard Process for Data Mining (CRISP-DM).

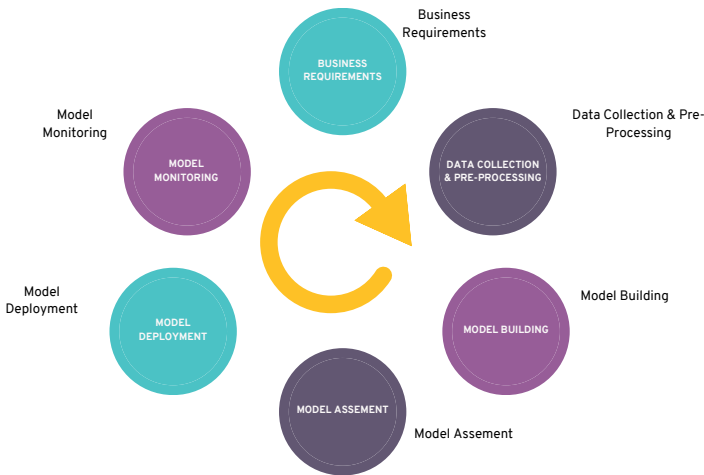


Figure 3: MLOps Lifecycle Processes

The business requirement phase is when the technologist figures out the objectives of the project and develops a project plan for bringing it to life. Data collection and pre-processing is the phase for collecting the data, understanding the data, and ensuring data quality through cleaning, formatting, and transforming the data. Extraction, transformation, and loading (ETL) may come in handy in this phase. The model building phase is when a few modeling techniques or algorithms are selected and used. The evaluation phase is when the model is assessed for performance. The deployment phase is when the model is deployed to a production environment (i.e., used by real people and processes for which it was designed). The monitoring phase is the iterative part of the project, where the model is monitored for performance degradation. Since MLOps is cyclical ( Figure 4), the process continues and never stops.



# Machine Learning Programming

Depending on the specific AI/ML job role, programming is still a prominent feature for building, testing, and deploying an AI/ML product. You might find the following programming languages in an AI/ML job description: Python, R, and Java. You will also find shell scripting languages such as Bash, Zsh, and Tcsh. See Figure 4. The scripting languages are mostly command-line based and can also be used for file manipulation, task automation, and sometimes to build machine learning models. They are not as powerful as Python or R.




Figure 4: Programming languages

## Machine Learning Models

What does an ML model even mean? In AI/ML, a model is a set of instructions or a tool that has been trained on a dataset to recognize specific patterns. For example, if I want to build a model to recognize fraud patterns, I'll use a dataset that contains records of fraudulent and non-fraudulent transactions. I'll then apply a classification algorithm (instruction) to the dataset. The algorithm spits out initial results that classify a group of transactions as fraudulent and others as non-fraudulent.

I can continue tweaking the instruction or I could apply other classification algorithms to see which algorithm provides better results and performance. I continue until I get to a stage in the training process where the majority of the truly fraudulent transactions are labeled as fraudulent (true positive). In this case, I have built a fraud classification model.

There are three approaches to machine learning: supervised, unsupervised, and reinforced. I will focus on the initial two in this audiobook because they are the major ones. Supervised learning is the approach where an algorithm is trained to recognize a pattern from a dataset that was previously labeled (e.g., fraudulent, non-fraudulent). Classification and regression methods are the two types of supervised ML. Unsupervised learning is an approach where the algorithm is not trained but recognizes hidden patterns or clusters in the data on its own. Here, there is no intervention from a person who trains or supervises the algorithm.

decades in marketing analytics to create market segmentation or user profiles. Association has also been used in marketing for creating the types of recommendations we've come to know as 'customers who bought this also bought that' or 'you may also like this item.' See Figure 5.

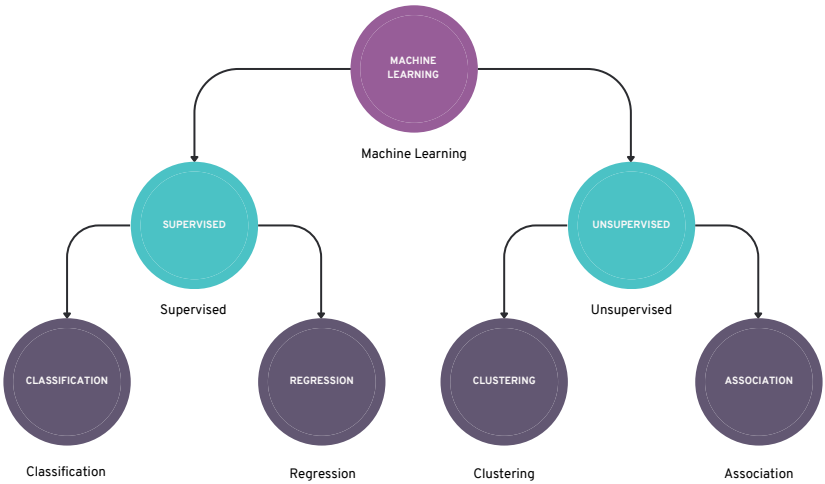



Figure 5: Machine Learning Models



# ML Algorithms

An algorithm is generally defined as a step-by-step process for carrying out a task or solving a problem. Therefore, your strawberry cake recipe can be described as an algorithm. There are specific algorithms that belong to each machine learning model. For example, the algorithms belonging to the classification model are neural networks, random forests, and decision trees. Some algorithms belonging to regression are simple regression and multiple regression. An example of a clustering algorithm is K-means. Just like the outcomes of your strawberry cake and meatloaf recipes are different, the outcomes from ML algorithms are also different.




Figure 6: ML Algorithms



# The AI/ML Technologist Roles

In recruiting AI/ML technologists, you'll find students who just graduated from school, beginners with one or two years' experience, and sometimes, the very experienced. Because of the wide variation in the processes, algorithms, and environments used for developing and implementing AI/ML solutions, you will find that the framework, tools, and methodologies in the resumes and job descriptions are different. However, the basics and foundations are similar. I describe four major AI/ML roles below.

## Job Roles of AI/ML Technologists

After reviewing and analyzing more than 400 AI/ML job descriptions, I can say (at least for now, since technologies and their job roles are dynamic) that there are three main categories of technologists in the AI/ML space. The job roles are listed as Engineers/Scientists, MLOps, and Ethics.

The Engineer/Scientist builds and develops the system, the MLOps technologist deploys and monitors the systems, and the AI/ML Ethics technologist audits the system. It is a continuous lifecycle process. Figure 7 represents what I call the AI/ML RoleCycle. In Table 1, I present the role titles and their primary tasks.

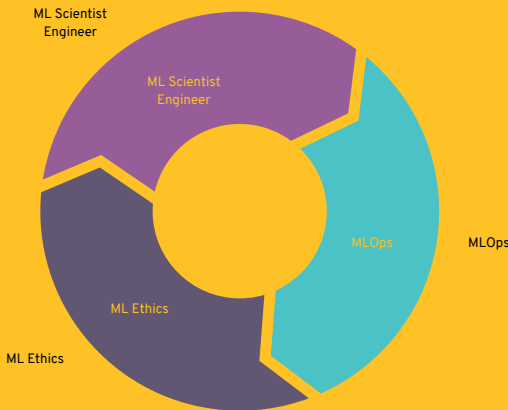


Figure 7: The AI/ML Role Cycle

Table 1: Role types, tasks, and targets.

Job Title	Main Tasks	Target
AI/ML Scientist/Engineer	They develop solutions applied to specific domains (e.g., healthcare, automotive, national security, transportation, banking, retail).	Build
AI/ML MLOps	They automate and deploy the solution.	Deploy
AI/ML Ethics	They develop AI policy and governance and audit AI solutions for fairness.	Audit
Others	Combination of all the above	All

## The ML Scientist/Engineer

About 80% of the AI/ML job roles have the “engineer” and or “scientist” designation. At first, they seem different, but on a closer look, you may find that they perform very similar tasks and have similar responsibilities and requirements. This person understands the business need or problem and works to apply AI/ML algorithms to address the need. It’s the application of ML to a specific industry or problem. You’ll see references to some experience in specific verticals – automotive, healthcare, and retail.

You will notice that this isn’t the case for the MLOps technologist. Because the scientist/engineer interfaces with customers to understand their needs and then apply them, there’s often a requirement for strong communication skills. Communication skills help the candidate gather requirements, ask questions, build relationships, and present solutions to customers.

The AI/ML Scientist/Engineer is the applied scientist. They are the ones who apply AI/ML technologies to business problems. Therefore, you will often find that the job description for this candidate requires some experience in a vertical domain, like healthcare, life sciences, transportation, retail, etc. Some experience in the domain helps them quickly identify the business issues and associated solutions.

For example, a technologist familiar with the healthcare domain already understands the nuances in patient care and the stringent regulations in that space and would be able to design ML solutions that account for the domain and recognize opportunities in the domain. The scientist/engineer is the core developer with experience in AI/ML algorithms and core statistical methods. They perform the requirements gathering, research, data wrangling, model building, and present the results to others.

## Characteristics of the ML Scientist/Engineer's resume or job description

- Business solution: Some experience in a vertical (e.g., Transportation, Healthcare, Manufacturing). Gather business requirements, identify opportunities, work with business stakeholders
- Frameworks: Some experience with ML frameworks (e.g., Tensorflow, PyTorch, Keras)
- Programming: Python, Java, C/C++
- ML Algorithms: Some experience in ML algorithms (deep learning, neural networks, clustering, decision trees)
- Statistical Analysis: Some experience in regression analysis, hypothesis development, time series
- Database Systems and Engineering: Some experience in big data systems (Hadoop, Hive, Spark) and data warehouse systems (Redshift, Snowflake). Expertise in relational and non-relational database concepts (SQL, NoSQL, ETL)
- ML Lifecycle: Some experience in building and optimizing data and ETL pipelines. Experience in ML workflow involving data prep, modeling, deployment, and monitoring
- Cloud: Some experience with deploying ML solutions in any of the major ML Cloud platforms (AWS GCP, Azure)



# AI/ML Job Description and Screening Questions

Here are questions the recruiter may ask after reviewing a job description. In this section, the job descriptions are reviewed to understand the job requirements and enable the recruiter to create Request for Information (RFI) questions that could be directed to the hiring manager (through an account manager in the case of the agency recruiter).

The objective of the screening process is not only to screen out candidates but also to screen them in. If a candidate is not suited for the current position, they might be for a future one. If a candidate does not qualify for an interview or placement at one company, it should not disqualify them from other positions. You want to understand how they have used their skills to solve similar business issues. The aim is to ensure you can present qualified candidates in the best possible light.

## Job Description: ML Engineer/Scientist

Based on what we learned from the previous sections (The AI/ML Technologists' Skills and The AI/ML Technologist Roles), you may already be familiar with the terms in this sample job description. This position is for an ML Engineer/Scientist who can design and build ML models. To do that, this person must be familiar with the AI frameworks (e.g., TensorFlow) and be familiar with programming languages for doing development in this area (e.g., Python). The candidate must also have expertise in working with at least one of the cloud providers (e.g., Azure, AWS) and with cloud computing (e.g., Kubernetes, Docker). See Figure 13 for some of the qualifications that need verifying.

After some pleasantries between you and the candidate, below are some sample technical screening questions you might ask based on the job requirement.

1. How did you start your career as a ML Engineer?
  - a. This question helps you understand the candidate's journey into this relatively newer area of technology and development. This question helps you ascertain:
    - i. Their years of experience in AI/ML space.
    - ii. Their previous job role before they switched to this area.
    - iii. Their interest in a new position.
2. The hiring manager is interested in candidates with solid experience in AI/ML; could you please describe an AI/ML project you've worked on using any of the algorithms (e.g., neural networks, decision trees, random forest, clustering)?
  - a. This question helps you ascertain their real-life experience in this domain. You might be able to ascertain:
    - i. Their level of experience relative to the job requirement.
    - ii. Their ability to communicate.
3. There is a general requirement that the person in this position is able to collaborate with customers to understand the business requirements for designing an AI solution. Could you please describe your experience in this area? Perhaps, describe when you've designed an AI solution based on a business need in an organization.
  - a. This question further helps you ascertain their ability to communicate. Recall that technologists in this position should be able to ask and answer questions and present to their stakeholders.
4. This job also needs someone with strong programming skills; could you please describe some of your previous experiences using Python, R, or Java in a data-centric environment?
  - a. This question ascertains their proficiency in programming. It is possible to do AI/ML work without heavy skills in Python and R. It is rare but possible.
  - b. You will also learn whether they are programmers who have upskilled into the AI/ML space.
5. If given the opportunity, which major cloud providers would you promote for their strength in AI/ML processes? What is your experience using any of them for AI/ML?
  - a. This question helps you understand their experience using cloud computing
  - b. You might learn about their experience in container technologies (e.g., Kubernetes, Docker).



6. I'm still learning all I can about AI/ML, and from what I've learned so far, AI/ML development goes through several stages. Which of the stages do you enjoy the most?
  - a. This question is looking for their understanding and experience in the AI/ML phases described in the "MLOps Skills" section (data gathering, data cleansing, model building, testing, deployment, and monitoring)
7. How do you usually determine the ML algorithm to apply to a business problem?
  - a. This is an open-ended question where the candidate describes a problem (e.g., classification, clustering problem) and then proceeds to tell you how they would select an algorithm to match the expected outcome. If the candidate is experienced and open to sharing, prepare for a lengthy response.
8. Knowing that there are many frameworks for building AI/ML solutions, can you describe the libraries/frameworks you've used and why you would choose one over the other (e.g., TensorFlow, PyTorch)?
9. In your experience in the AI/ML domain, how do you select and collect the data used for training, building and testing your models?
10. Overall, how would you describe yourself and your skills to a hiring manager?
  - a. From this question you might ascertain:
    - i. If the candidate can perform well in a telephone/face to face interview
    - ii. The level of confidence the candidate has in their skills and abilities



## Conclusion

Technologies are changing at an even faster rate than previously thought. Organizations are scrambling to train, recruit, and retain employees who embody the potential to handle the power of these technologies. One of those technologies is Artificial Intelligence and Machine Learning. As a technical recruiter, you make it possible for organizations to find the technologists they need. However, to excel at the technical recruiter role, you need to understand the technologies that your candidates, clients, or hiring managers use. By reading this book, I hope you have come away with a better understanding of AI/ML technologies, the principal job roles, and some questions to ask your candidates to ascertain their knowledge and experience in this area.

# References

1. Deloitte, 2020, Government Trends 2020: What are the most transformational trends in government strategy today?
2. Forbes, 2021, <https://www.forbes.com/sites/louiscolumnbus/2021/01/17/76-of-enterprises-prioritize-ai--machine-learning-in-2021-it-budgets/?sh=1cf9a009618a> Accessed May 2022.
3. Friedrich, S., Antes, G., Behr, S. et al. Is there a role for statistics in artificial intelligence?. Adv Data Anal Classif (2021). <https://doi.org/10.1007/s11634-021-00455-6>
4. Harvard, 2022, <https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/>. Accessed April 2022.
5. SAS, 2020, [https://www.sas.com/en\\_us/insights/articles/big-data/artificial-intelligence-machine-learning-deep-learning-and-beyond.html](https://www.sas.com/en_us/insights/articles/big-data/artificial-intelligence-machine-learning-deep-learning-and-beyond.html). Accessed in January 2020.
6. TELUS International, 2022, <https://www.telusinternational.com/articles/ai-fraud-detection-in-banks>. Accessed May 2022.