Careers in Artificial Intelligence

Guidance for New and Emerging Tech Talent

Interested in AI?

Are you an aspiring technologist in New York City? Do you have an interest in a career in data science, software engineering, or artificial intelligence? If so, this briefing is for you.

Background

On April 16, 2019, TTP held its fourth annual academic summit, this time focusing on careers in Artificial Intelligence. The goal of the summit was to help NYC-based colleges and tech training bootcamps ensure their curriculum helps current students prepare for roles that are either primarily AI roles or require AI skills. We have distilled the main points that came out of that summit into this document for your use.

What is Artificial Intelligence?

Artificial Intelligence (AI) is a system's ability to correctly interpret external data, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation.[1]

Machine Learning (ML) is the most common modern application of A.I. whereby algorithms use historical data to recognize patterns and create mathematical models in order to make predictions or decisions. This includes:

- **Supervised learning,** where both inputs and outputs are provided and the goal is for the program to "learn" general rules that map inputs to outputs.
- **Unsupervised learning,** where no labels are given to data and the goal is for the program to identify structure and patterns.
- **Reinforcement learning,** where the program operates in an interactive environment and "learns" based on rewards and punishments received for each action.

Deep Learning (DL) is a category of Machine Learning algorithms that use artificial neural networks to transform data at multiple layers of abstraction. It is used commonly in modern applications such as computer vision, speech recognition, and natural language processing.

[1] Kaplan and Haenlein.

How are companies in NYC using AI?

NYC-based tech companies are increasingly using machine learning and other AI techniques to conduct core business functions. This varies widely from company to company, but examples include:

- **A consumer-goods online marketplace** that uses AI to understand customers' preferences and personalize purchase recommendations
- **A news aggregation and analytics company** that uses natural language processing to categorize the articles they publish by keyword and topic
- **A social media company** that uses machine learning to identify and weed out harmful content

Companies note that AI has existed for decades, but that its use in commercial settings has become more widespread as computation power has improved and the amount of data available to inform model-building has increased.

In addition, NYC tech professionals anticipate a "snowball effect" over the next five years with the use of AI becoming more widespread as more data enables the production of better and more powerful models.

What roles require or use AI/ML?

AI is a quickly growing part of many tech roles in NYC. According to Labor Insight Jobs (Burning Glass Technologies), over the last five years, job listings that mention of Artificial Intelligence or Machine Learning grew by more than ten times in the last five years – from just over 2,000 listings in 2014 to over 21,000 in 2019.[2]

There are some roles that are mostly or exclusively machine learning or other AI-specific jobs. This includes over 700 ML-specific roles and approximately 130 Natural Language Processing-specific roles.

However, many jobs that require the use of AI are in other fields. This includes various types of software engineering jobs, such as software and app development roles, with over 3,400 jobs advertised that require some use of AI or ML. This is closely followed by data-related roles, including data science, data engineering, and data analyst roles, with over 2,600 roles that require some use or knowledge of AI or ML.

What this means is that, especially for software engineers and data scientists, it is increasingly important to build the foundational skills you may need to work with machine learning and related disciplines, even if you are not pursuing a specific AI/ML career.

How should I prepare for these jobs?

Whether you specifically want to pursue a role in AI/ML or just want to understand this well enough to do your job well, the best thing you can do now is to build a solid foundation as a software engineer or data scientist. If you've mastered the fundamentals of your field, you can pick up AI or ML-specific skills as you advance in your career.

That said, if you are specifically interested in Machine Learning or AI roles:

[2] 2,060 listings from 12/1/13 – 11/30/14; 21,774 listings from 12/1/18 – 11/30/19

- **Decide what kind of work you would like to do day-to-day, and begin working towards that role.** There are many different kinds of roles that involve the use of AI and ML. Exact job titles and responsibilities vary across companies and will continue to change, so you should read current job descriptions to understand the landscape and decide which roles you most want to pursue.

  That said, broadly speaking, role types include:

  Building models
  - **Machine Learning Scientists,** who develop the mathematical models that drive machine learning
  - **Infrastructure engineers,** who build scalable systems
  - **Data Scientists,** who build models and take them to the point of deployment; they also do analysis and contribute to business decision making
  - **Machine Learning engineers,** who solve product problems (for example, figuring out the best content to return when you search for the word "coffee"); Machine Learning engineers spend most of their day cleaning data

  Maintaining systems
  - **Software Engineers (back-end or data engineers),** who help to maintain models and ensure they run smoothly in production
  - **Dev Ops engineers,** who are responsible for the maintenance of ML models already in production

- **Begin taking Machine Learning or related courses while you are in college, or choose a Data Science or related bootcamp that includes this in its curriculum.** If you know that you're interested in this work, or just want to explore it, it is helpful to gain early exposure to the foundations that will underpin your work.

- **Get close to the work and build your skills with the ML team.** When you are ready for a job or internship, find a company that is using ML as a big part of its work. Figure out which team(s) within the company do the most ML work and try to get a role that is either on that team or does work with it. As an entry-level team member you will likely not begin as an ML specialist, but coming in as a software engineer or data analyst/scientist who does work with the ML team will allow you to begin building the skills and experience you will need to prepare and compete for those roles as you advance.

  Good roles for gaining experience at the entry level include:

  - **Data analyst** roles at a company that uses ML/AI widely. As a data analyst, you will help with day-to-day decision making for how to build or scale features across the business. This will allow you to gain exposure and experience required for more advanced roles.
  - Similarly, **software engineering** roles at a company that uses ML/AI widely. Focus on getting a role that puts you in close contact with a team that is working on ML/AI problems, and work to get on projects that allows you to build this experience.