



ARTIFICIAL INTELLIGENCE CAREER GUIDE

Student Affairs at Carnegie Mellon University Silicon Valley



TABLE OF CONTENTS

CARNEGIE MELLON AND ARTIFICIAL INTELLIGENCE.....	3
INTRODUCTION.....	4
COURSEWORK, SKILLS, AND COMPETENCIES.....	6
RESEARCH EXPERIENCE.....	10
ROLES IN ARTIFICIAL INTELLIGENCE.....	13
INTERVIEW PREPARATION.....	15
NEXT STEPS.....	16



ARTIFICIAL INTELLIGENCE CAREER GUIDE

Acknowledgment section

This AI career guide is updated every 6 months. If you have any suggestions on material that should be included or edited, please reach out to: career-services@sv.cmu.edu. We appreciate the hard work of Chaitra Dinesh Pai, Nursultan Dyussebayev, Arpit Gupta, Sweta Hari Kumar, and Joseph Konan who researched specific fields of interest and compiled helpful tips and resources for students to enable them to make the most of their CMU-SV experience and navigate the internship and job search process in Silicon Valley. Their valuable contributions made these career guides possible. Thank you to staff, faculty, and alumni who contributed their time and effort to make this guide a helpful tool for CMU-SV students.

Carnegie Mellon University and Artificial Intelligence

Carnegie Mellon University (CMU) has been a leader in the field of Artificial Intelligence (AI) since the “founding fathers” of AI, Herbert A. Simon and Allen Newell, helped to establish the university's first computer science degree program in 1961. Since the emergence of AI, CMU has remained a cutting-edge leader in AI research and education. As the #1 school in the field according to [U.S. News and World Report](#), CMU consistently pushes the bounds of what is considered state-of-the-art with its world-leading faculty, prestigious research labs, and diverse network of interdisciplinary researchers, alumni, and industry sponsors. The Carnegie Mellon University Silicon Valley campus is located in the heart of the San Francisco Bay Area, the center of AI industry and opportunities, making the campus an excellent location for students pursuing AI.

AI and machine learning is a branch of computer science that typically involves a computer or application trying to replicate some level of human intelligence to complete a predetermined task. Machine learning is a subfield of AI. This career guide is intended to help students understand job opportunities and career paths within the field of AI and prepare them for a job search in various AI specialties. We interviewed both students and alumni who work in AI to make this career guide even more useful to you. The information within this guide is intended to serve as a starting point for students interested in pursuing a career in AI and prompt further research into strategic planning for their career. Carnegie Mellon University has dedicated a website here, ai.cmu.edu, to showcasing how the university is currently working in AI and highlights key research, education efforts, and AI news and events.



Image: Carnegie Mellon Artificial Intelligence Beats Chinese Poker Players by \$792,327

INTRODUCTION

Careers in AI

Artificial intelligence is one of the fastest growing fields and has the potential for a tremendous impact on our society. As of June 2024, there were more than 60,000 jobs posted on LinkedIn related to AI in the United States alone and with automation and innovation increasing, the field is likely to continue to grow. Companies benefit tremendously from AI, as it helps them realize higher profits by automating tasks, improving efficiency, predicting trends, and improving experiences for products or services. The applications for AI are vast as well, ranging from predicting ad revenue, face or voice recognition on iPhones, language translation on YouTube, medical diagnostics, and more.

Students who specialize in AI will encounter a myriad of opportunities in their careers and endless paths they can follow based on their interest and skill set. Today, there are few fields where AI professionals are not found. Fields such as healthcare, finance, marketing, retail, manufacturing, education, transportation, cybersecurity, and more all have roles involving AI. Titles for roles within AI are broad and will depend on the field, duties, and specialization of the role; samples include that of Machine Learning Engineer, Deep Learning Engineer, and Machine Learning Software Engineer. In addition, AI professionals often specialize in one or more areas of computer vision, natural language processing, signal processing, deep learning, financial modeling, robotics, and/or theory, making AI as a field broad and promising. Despite the broadness of careers in AI, there are clear ways students can prepare in advance.

Preparing for a Career in AI

Perhaps the first and most obvious way to prepare for a career in AI is through coursework. The most common roles in artificial intelligence involve data science, machine learning, and deep learning, so coursework within these areas are important. Of the three, companies leveraging AI often look for knowledge and experience with deep learning, so taking courses and getting experience where students are able to expand their deep learning knowledge is advised and will make getting an internship or full-time role in AI easier. In addition to courses, it is important for those entering the field to gain practical experience through things such as projects and internships. Employers especially like to see practical experience through relevant, high-impact projects that have been deployed, published, or have won Kaggle competitions. In the case of internships, projects directly related to their specific data and desired outcomes are required. Networking and community engagement also offer promising opportunities for students to further develop their AI portfolio.

Coursework

Students who are serious about getting involved in this field will want to take courses that expose them to machine learning and deep learning foundations, such as 11-785, Introduction to Deep Learning. This course gives students access to the network of other students who have completed the course -- connections directly relevant to the field and industry. The course is well-rounded in terms of concepts and helps with understanding the fundamentals of Deep Learning. In addition, students should take a course that teaches them how to work with data, such as Cloud Computing and Distributed systems, as being able to work with data is key to the development, training, and deployment of machine learning systems. Beyond that, students can take courses that allow them to specialize in various areas relevant to AI related to their interests. More information, including a map of courses offered at the CMU-SV campus, can be found in the Coursework, Skills, and Competencies section of this guide below.

INTRODUCTION

Practical Experience

Students interested in pursuing a career in AI can enhance their portfolio through projects and internships in order to gain industry experience and understand real-world applications of AI. At CMU-SV, students who have gone on to a career in AI have had internships at companies including TikTok, Equinix, Amazon, Intuit, AI startups and others in roles ranging from product manager intern, software engineer intern, machine learning engineer intern and others. In addition, competitions and hackathons are a great way to solve practical problems, improve your skills, and build a portfolio. Hackathons students & alumni have attended in the past include Brembo Hackathon, AGI House Hackathon, and others posted on Slack. Student organizations are also a great way to gain practical experience in a field as well as improve leadership and organizational skills. Contact Student Affairs to start or join an AI-related student organization.

Networking and Community Engagement

Attending AI-centered conferences and meetups is also recommended for students interested in pursuing a career in AI. Conferences attract industry leaders and can provide insights into AI trends and emerging technologies, as well as offer students networking opportunities. In addition, students can expand their knowledge in areas outside of the classroom and gain increased exposure to diverse perspectives. Top conferences in the field of AI include but are not limited to GenAI Summit San Francisco, SuperAI, and Data + AI Summit. Several of the AI conferences are located in the San Francisco Bay Area, very close to CMU-SV's campus, which gives CMU-SV students an added boost. In addition to conferences, students can attend more informal meetups, which can also further business contacts, opportunities, and knowledge. There are also professional organizations in AI, such as the Association for the Advancement of Artificial Intelligence, which offer student discounts for memberships and benefits that range from discounts to conferences, publications, and events, access to a social network of members, and more.



COURSEWORK, SKILLS, AND COMPETENCIES

Having a strong project portfolio is vital to having a strong application. At Carnegie Mellon, you have access to world-class faculty who will help you recreate state-of-the-art results, design scalable infrastructure, or implement complex algorithms. Relevant courses include:

Coursework

Group 1: Foundations

- 18661 Intro to Machine Learning for Engineers
- 11785 Introduction to Deep Learning (Recommended)

Group 2: Working with data

- 18656 Functional Programming in Practice
- 18668 Data Science for Software Engineering
- 14736 Distributed Systems: Techniques, Infrastructure, and Services
- 14848 Cloud Infrastructure and Services
- 15619 Cloud Computing

Group 3: Specialization 1

- 18797 Machine Learning for Signal Processing
- 14795 AI Applications in Information Security

Group 4: Specialization 2


- 49784 Artificial Intelligence for Product Managers
- 49779 Applied Artificial Intelligence
- 49781 Introduction to Machine Learning
- 14-691 AI Ethics: Multidisciplinary Perspectives and Industry Insights

Group 5: Research Experience

- 18980 M.S. Graduate Project (repeatable)
- NA Course Project with Advisor
- NA AI Internship
- 49-790 Independent Study
- 49-796 Independent Study
- 49-791 Capstone

Group 6: Other

- 18645 How to Write Fast Code
- 18751 Applied Stochastic Processes



COURSE PROJECTS

Below are examples of course projects. Please email career-services@sv.cmu.edu if you have any questions about your own resume or if you need to schedule a resume appointment, please do so through Handshake.

Self-Attention Networks for Connectionist Temporal Classification

Developed end-to-end speech recognition encoder-decoder RNN. Generated transcripts, surpassing human-level performance w/ Levenshtein distance of XX.X. PyTorch, Python.

Deep Residual Learning for Image Classification and Verification

Developed deep residual learning CNN architecture for face identification with XX.X% accuracy and 0.XX AUC, leveraging center loss and triplet loss. PyTorch, Python.

Sparse Multilayer Perceptron for Phoneme Recognition

Created high-efficiency low-parameter MLP network to identify phoneme state labels using raw mel spectrogram frames. XX.X% accuracy. PyTorch Python

Deep Learning Automatic Differentiation Libraries

Implement low-level logic of an end-to-end deep learning framework, together with automatic differentiation and optimization or MLPs, CNNs, and RNNs. PyTorch, Python.

Eigenfaces for Face Recognition and Bounding

Implemented eigenfaces as an unsupervised dimensionality reduction algorithm, leveraging results by reshaping eigenvectors for principal component analysis. Matlab.

Source Separation w/ Single Channel Independent Component Analysis

Implemented my own ICA based on freeing fourth moments method, extracting independent components to simultaneously speaking sources from a recording. Matlab.

Signal Separation using Non Negative Matrix Factorization

Separated mixed signal into component signals, extracting speech from recordings with background sound and music by minimizing KL divergence between spectrograms. Matlab.

Sparse Recovery of Streaming Signals Using Iterative Hard Thresholding

Implemented iterative hard thresholding with projected gradient descent algorithm, recovering sparse image transform from 3 compressed measurements. Python, Matlab.

Big Data Twitter Analytics Web Service Design and Implementation

Developed data warehouse complete with front-end query response system, backend database, ETL tool, to process terabytes of tweets on AWS. Java, Tomcat, Python, Hbase.

Blockchain Ledger Design and Implementation

Developed blockchain mining application with ledger that will award coins if a block is successfully solved given a difficult work and add a new block into the network. Go, Java.

Distributed File System Design and Implementation

Implement a simple distributed file system using my own RMI library, allowing files to be hosted remotely on one or more storage servers complete with API. Go, Java.

Distributed Key-value Store Design and Implementation Using Raft

Implement Raft algorithm as a replicated state machine protocol, achieving fault tolerance by storing copies of data on multiple replica servers. Go, Java.

Bigram ETL Tool Design and Implementation Using Apache Spark

Counts number and frequency of bigrams, as well as the number of bigrams required to add up to 10% of all bigrams within corpus. Apache Spark, Java.

BigTable, a Scalable NoSQL Database Design and Implementation

Implemented a distributed storage system for managing structured data, designed to scale to a very large size: petabytes of data across thousands of commodity servers. Java.

Multi-path TCP Analytic Tool of Instant Virtual Networks in Mininet

Analyzed Three-Tier and Fat Tree network topologies as well as east-west congestion at the transport level using MPTCP approach by simulating network traffic. Mininet, Python.

Remote Method Invocation API Tool for Distributed Microsystems

Implemented remote method invocation (RMI) library for calls over network connections, permitting objects located on one Java virtual machine to call methods on another. Java.

```
[1])===?1&&.stoponfalse){t=1;break}n=1,de
?o=u.length:r&&(s=t,c(r))}return this},remove
nction(){return u=[],this},disable:function()
re:function(){return p.fireWith(this,argument
ending",r={state:function(){return n},always:
romise)?e.promise().done(n.resolve).fail(n.re
dd(function(){n=s},t[1^e][2].disable,t[2][2].
=0,n=h.call(arguments),r=n.length,i=1!==r||e&
(r),l=Array(r);r>t;t++)n[t]&&b.isFunction(n[t
/><table></table><a href='/a'>a</a><input typ
yTagName("input")[0],r.style.cssText="top:1px
test(r.getAttribute("style")),hrefNormalized:
```


Skills and Competencies

Skills and competencies within the field of AI vary depending on the role and specialization. AI professionals are expected to have technical skills that include a deep knowledge of programming languages, such as Python and R, as well as a strong grasp on data processing, cloud computing, and handling. Analytical and problem-solving skills are also important, as AI professionals must be able to analyze complex problems and develop novel solutions. In addition, soft skills like communication, teamwork, and leadership skills are nearly universally sought after. Below is a list of competencies and skills that are particularly important in the field of AI:

Technical Skills and Competencies

- Artificial Intelligence and Machine Learning
- Advanced Mathematics and Statistics
- Cloud Computing Platforms (e.g. AWS, GCP)
- Computer Vision
- Data Science and Management (digestion, extraction, analysis, visualization, etc.)
- Data analysis and visualization
- Machine Learning (Machine learning frameworks such as: TensorFlow, PyTorch)
- Deep Learning (concepts including: convolutional neural networks, recurrent neural networks)
- Mathematical Reasoning
- Natural Language Processing
- Programming Prowess (Python, R, Java, C++, etc.)
- Modeling and Simulation
- Software Engineering
- Testing and Validation

Soft Skills and Competencies

- Attention to Detail
- Continuous Learning
- Creativity and Innovation
- Emotional Intelligence
- Flexibility and Resilience
- Mental Visualization
- Integrity and Ethical Competence
- Interpersonal Skills
- Oral and Written Communication
- Problem Solving and Critical Thinking
- Project Management
- Self-Management
- Teamwork and Digital Collaboration

References for Skills & Competencies Listed Above –

1. <https://chcoc.gov/sites/default/files/Skills-Based%20Hiring%20Guidance%20and%20Competency%20Model%20for%20Artificial%20Intelligence%20Work.pdf>
2. <https://www.linkedin.com/pulse/must-have-competences-skills-our-new-ai-world-reform-fawzi-mnf7c/>
3. <https://www.datacamp.com/blog/essential-ai-engineer-skills>

RESEARCH EXPERIENCE

Students who are interested in careers in AI are expected to have done research with a faculty member that specializes in the domain. In particular, research assistantships are easily attainable and lead to high impact projects given you are working with faculty members at Carnegie Mellon University. For perspective, a Research Scientist in industry will have hundreds of citations. In contrast, ECE and SCS professors will have thousands of citations. High-tier internships have “a publication in top journals” as a sought after or required qualification, and you will be more successful if you publish in your interest area.

Many courses, including 11-785, 14-736, 14-848, and 18-797, require you to create a project proposal and realize that idea by getting results. These courses, in addition to others not listed, allow you to have a faculty advisor on your project. This means, you are able to choose a research idea that the faculty member is interested in working on and get research experience while doing your coursework. In addition, 18-980 is an option for students who have a research project not linked to a traditional course, allowing you to get core credit for doing research with faculty. There is no excuse for not having at least one research assistantship working with faculty that specialize in your area of interest. Further, if you are starting at CMU with no relevant experience in the field, it is readily attainable.

If you want to work at a specific company doing AI research and development, then cut to the chase. People who want to work at Amazon should have research projects improving the products or experience of Amazon’s engineers or customers. A simple google search yields multiple big-data resources that are used by Amazon and researchers:

- [SNAP Web Data: Amazon Reviews](#)
- [Topical Chat Dataset: Amazon Alexa](#)
- [Sales Rank Data: Amazon Print and Kindle Books](#)
- [Movie Review Web Data: Amazon Prime Video](#)
- [Stock Price Data: AMZN \(Amazon\)](#)

Applying to Amazon and having multiple projects on Amazon will most likely get your application noticed by Amazon. Moreover, a code you write for these specific datasets relevant to Amazon can be ported to get similar insights and results for other companies. To emphasize this point, here are some equivalent Google datasets for those listed above:

- [Google Play Store App Data](#)
- [TensorFlow Speech Recognition Challenge: Google Home](#)
- [Google Analytics Customer Revenue Prediction](#)
- [Video Comments and Metadata: YouTube](#)
- [Stock Price Data: GOOG \(Google\)](#)

And more [datasets for Amazon](#) or [datasets for Google](#) are out there, and you can simply search for your top company’s data so you can make the best resume for your application.

PATHWAYS

Depending on the amount of AI related coursework you take, and the extent to which you engage with faculty, it may make sense for you to consider a variety of options.

Data Science Bootcamps. These, in terms of competitiveness, are among the least competitive given you are an engineering student from Carnegie Mellon University. Bootcamps often cost money, so, like school, you are paying to build your human capital.

Data Science Fellowships. These are like Data Science Bootcamps but are more competitive and are often sponsored by employers. These opportunities are free, and, for this reason, there is a more competitive process getting in. Being a student who has more relevant coursework will give you more intuition than the people who are very smart but want a career change.

AI Residencies. These are ideal if you are shooting for a competitive AI role, but have either published in a low- to mid-tier journal or have not published at all. Even if you have published in a high-tier journal, these are still worth considering because you are able to do similarly impactful research but in an industry setting.

[Consider this list if you are interested in AI Residencies.](#)

DS and ML Roles at Mid-tier Companies. These roles are not hard to find because Data science and AI is so popular. There are many people who want to break into data science, and they tend not to first consider bootcamps or fellowships (though they should) so the applicant pool is very large. To stand out at mid-tier companies, it is not necessary that you have published, but you should have relevant project experience on the same problem they are working on. Even better, if you got results using their data or improved their product, then you will be taken more seriously by the company you are applying to. The prestige of this is between Fellowships and AI Residencies, and pays similar to the latter.



PATHWAYS

Let's explore a list of roles that use AI in their work everyday:

Pure Technical Roles:

- **Data Scientist (DS):** Unveils hidden insights from data using statistical methods and machine learning. They need strong analytical skills, programming expertise (Python, R), and data visualization know-how.
- **Machine Learning Engineer (MLE):** Builds and maintains AI models trained by Data Scientists. MLEs require expertise in software development, machine learning algorithms, and cloud computing platforms.
- **AI Engineer:** Designs, develops, and implements AI systems. They need a blend of computer science, software engineering, and machine learning skills.
- **Robotics Engineer:** Creates intelligent machines. Robotics Engineers combine mechanical engineering, electrical engineering, and computer science with AI to design robots.
- **Software Engineer (SE) with AI Focus:** Develops software applications infused with AI functionalities. SEs with an AI focus need strong programming languages (C++, Java) and an understanding of machine learning concepts.
- **Research Engineer/Scientist:** Pushes the boundaries of AI by developing new algorithms and techniques. They require a Ph.D. in a relevant field (computer science, engineering) and strong research and experimentation skills.
- **Domain-specific Scientist/Engineer utilizing AI:** Applies AI to solve problems in a specific field (e.g., healthcare, finance). They need expertise in their domain along with a strong foundation in AI techniques. For example Bioinformatic Research Scientists utilizing Machine Learning to study and derive insights from their focus areas.
- **Security Specialist:** Safeguards AI systems from cyberattacks. Security Specialists need a deep understanding of AI systems and cybersecurity best practices.

Techno-Business Roles:

- **Data Analyst:** Cleans, organizes, and analyzes data to support business decisions. Data Analysts need strong data manipulation skills (SQL) and data visualization tools (Tableau).
- **Business Analyst/BI Developer:** Bridges the gap between business needs and AI solutions. Business Analysts/BI Developers need business acumen, data analysis skills, and knowledge of BI tools.
- **Product Manager:** Oversees the development and launch of AI-powered products. Product Managers need strong communication, product strategy, and user experience design skills.
- **VC Associate:** Analyzes and leads investments in promising AI startups. VC Associates also need financial expertise and a keen eye for spotting innovative AI companies.

References for Roles Shared Above:

- www.datacamp.com/blog/jobs-in-ai-to-pursue
- www.coursera.org/articles/artificial-intelligence-jobs
- www.springboard.com/blog/data-science/careers-in-ai/
- appacademy.io/blog/artificial-intelligence-ai-careers-to-consider
- datategy.net/2024/01/19/artificial-intelligence-jobs-in-2024/

MANAGING AI PRODUCTS AS A PRODUCT MANAGER

Why is AI Important for a Product Manager in 2024?

AI has become a transformative force in product management due to its ability to enhance decision-making, streamline operations, and drive innovation. For product managers in 2024, understanding AI is crucial because it enables them to harness data-driven insights, automate routine tasks, and deliver personalized user experiences. The proliferation of AI technologies, such as machine learning and natural language processing, has made it possible to solve complex problems, predict market trends, and create smarter products that meet evolving customer needs.

Utilizing AI vs. Managing AI Products

While utilizing AI involves integrating AI tools to improve product management processes, managing AI products as a PM requires overseeing the development and lifecycle of products built on AI technologies. Utilizing AI might include using AI-driven analytics for better decision-making or employing AI in customer support. In contrast, managing AI products involves strategic planning, defining AI model requirements, collaborating with data scientists and engineers, and ensuring the product's alignment with business objectives and market demands.

What is an AI Product Manager?

An AI Product Manager (AI PM) is responsible for guiding the development of AI-based products. They lead roadmap planning given the unique lifecycle of the model development, manage stakeholders' expectations of the probabilistic-based solutions, and drive market ramp strategy. They possess a blend of technical knowledge in AI and expertise in product strategy and market positioning. AI PMs work at the intersection of technology and business, ensuring that AI solutions not only leverage cutting-edge technology but also address customer needs and enhance business value. This role involves strategic planning, cross-functional team leadership, and continuous innovation to keep up with AI advancements.

How is AI Product Management Different from Generic Product Management?

AI product management differs from generic product management in several ways, primarily due to the complexity and technical nature of AI technologies. AI PMs need to understand machine learning models, data requirements, and ethical considerations specific to AI. Despite these differences, core aspects of product management such as user-centered design, market analysis, and cross-functional collaboration remain unchanged. AI PMs must also prioritize features, manage product roadmaps, and ensure alignment with business goals, much like their non-AI counterparts.

Leveraging AI Across Different Stages of Product Management

AI can be leveraged across various stages of product management and development, including:

- **Ideation:** Using AI to analyze market trends and generate innovative product ideas.
- **Design:** Employing AI for user experience research and creating personalized designs.
- **Development:** Integrating AI models to enhance product functionality and performance.
- **Testing:** Utilizing AI for automated testing and quality assurance.
- **Launch:** Applying AI for targeted marketing campaigns and user engagement strategies.
- **Post-Launch:** Using AI for continuous product improvement through user feedback and data analysis.

How Much AI Do I Need to Know? (Proficiency Level)

An AI Product Manager does not need to be an expert coder but should have a strong understanding of AI concepts, machine learning fundamentals, and data science principles. Proficiency in AI for a PM involves being able to communicate effectively with technical teams, understanding the limitations and possibilities of AI technologies, and making informed decisions based on AI-driven insights. A foundational knowledge of Python and familiarity with AI tools and frameworks can be beneficial.

MANAGING AI PRODUCTS AS A PRODUCT MANAGER

Skill Sets Needed to Manage AI Products

Key skill sets for managing AI products include:

- **Technical Knowledge:** Understanding AI/ML concepts, data analysis, and model development.
- **Strategic Thinking:** Ability to align AI initiatives with business goals and market needs.
- **Project Management:** Coordinating cross-functional teams and managing development timelines.
- **Ethical Awareness:** Ensuring AI products adhere to ethical standards and regulatory requirements.
- **Communication Skills:** Translating technical details into business insights for stakeholders

Key Roles and Responsibilities of AI Product Managers

AI Product Managers are responsible for:

- **Product Vision and Strategy:** Defining the vision and roadmap for AI products.
- **Collaboration:** Working with data scientists, engineers, and designers to develop AI solutions.
- **Market Analysis:** Identifying market opportunities and user needs.
- **Project Management:** Overseeing product development from ideation to launch.
- **Ethical Oversight:** Ensuring compliance with ethical standards and regulations

How to Be an Ethical AI Product Manager

Ethical AI product management involves:

- **Transparency:** Ensuring AI operations and decision-making processes are transparent to users.
- **Fairness:** Mitigating biases in AI models to ensure equitable outcomes.
- **Privacy:** Protecting user data and complying with privacy regulations.
- **Accountability:** Being responsible for AI product decisions and their impacts

AI Product Management in Startups vs. BigTech

AI product management in startups often involves more flexibility and faster iteration cycles, allowing for rapid innovation and experimentation. In contrast, BigTech companies usually have more resources, established processes, and a larger user base, which can lead to more rigorous development and deployment practices. Both environments require AI PMs to be adaptable, innovative, and strategic in leveraging AI technologies

Key Challenges Faced by AI PMs

AI PMs face several challenges, including:

- **Data Quality:** Ensuring the availability and accuracy of data for training AI models.
- **Technical Complexity:** Managing the intricacies of AI algorithms and integrating them into products.
- **Ethical Concerns:** Addressing biases, privacy issues, and regulatory compliance.
- **Market Readiness:** Gauging user acceptance and readiness for AI-driven solutions

INTERVIEW PREPARATION FOR ROLES IN AI

For software engineering roles & machine learning roles, block part of your daily work to complete LeetCode questions. Attend Coding Gym weekly sessions if you need support with this as well at CMU-SV. Ensure that you conduct technical mock interviews and also prepare for behavioral questions (examples listed below).

General AI and Industry Behavioral Interview Questions to Consider:

- AI's Impact: How do you see AI transforming the industry you're in? What are the biggest challenges and opportunities?
- Ethical Considerations: What are the most pressing ethical concerns surrounding AI development and deployment? How do you see these issues being addressed in your industry?
- Talent and Skills: What skills and knowledge do you believe are essential for AI professionals today? How is the talent landscape evolving?
- Industry Trends: What are the most exciting developments in AI that you're following? How do you see these trends impacting your work?

Role-Specific Questions

AI/Software Engineers

- What AI tools & frameworks do you use most frequently in your work?
- Can you describe a challenging project you've worked on? What was your role and what were the outcomes?
- How do you approach debugging and troubleshooting AI models? How do you ensure model accuracy and reliability?
- How do you approach data collection, cleaning, and preparation for AI projects? What role does data quality play in model performance?

Product Managers

- AI Product Development: Can you describe the process of developing an AI-powered product from concept to launch? What metrics do you use to measure success?
- User Experience: How do you ensure that AI features are user-friendly and provide real value? How do you gather and incorporate user feedback?
- Business Impact: How do you measure the impact of AI on your company's bottom line? Can you share an example of how AI has driven revenue or cost savings?
- AI Strategy: How do you align AI initiatives with overall business goals? What is your approach to prioritizing AI projects?
- Product Development and Strategy: How do you handle the integration of AI technologies into existing products? What key skills do you believe are essential for managing AI products?

Cybersecurity Professionals

- Threat Detection: Can you describe how AI can be used to detect and respond to cyber threats? What are the limitations of AI-based threat detection?
- Privacy and Security: How do you balance the need for data privacy with the use of AI for security purposes? What measures can be taken to protect sensitive information?
- What are the unique security challenges posed by AI technologies?
- How do you stay ahead of emerging threats and vulnerabilities in AI systems?
- What are some common misconceptions about AI security?
- What tools and methodologies do you use to secure AI systems?
- How do you evaluate the security of an AI model or system?

NEXT STEPS

CMU alumni shared their top tips for breaking into AI and here they are:

- During your internship, build a reputation as being a go-to person for AI
- Learn to make AI work for you
- If you start out at a company not in AI, consider working on a side project related to AI and invite yourself to AI related meetings at your company to get your foot in the door on a future AI project (share your interest in AI with your manager too!)
- Find a mentor at your company that works in AI & share your current work
- Share your projects on LinkedIn. Whether it's a course project or a passion project (the latter is especially worth highlighting), showcasing your work can make a big difference. Completing a project is important, but it's equally vital to ensure that others know about it. In many ways, storytelling is the currency of your career!
- Be present in your classes and always be curious!

