

AI Fundamentals and AI Advanced: Teacher Resources

Note: this is not an exhaustive list of resources. The content of this list is current as of the publication date of July 2026.

Contents

Amazon Web Services (AWS) Resources:	1
Examples of K Nearest Neighbors (KNN):.....	1
Case Studies—Responsible AI in Practice:	2
Case Studies—Failure to Apply Responsible AI Principles:	3
Additional/General List of AI Resource Links:.....	4

Amazon Web Services (AWS) Resources:

- [Free Artificial Intelligence Services on AWS](#)
 - [AWS Educate](#) (*Free | Browser-based | No AWS account required for students*) Amazon's purpose-built learning platform for students and educators.
 - [AWS Academy](#) (*Free for accredited institutions | Instructor-led curriculum*) If your school or district has a relationship with a higher education partner, AWS Academy provides instructor-led, college-credit-aligned AI and cloud curriculum.
 - [Amazon SageMaker Studio Lab](#) (*Free tier available | Browser-based*) For more advanced students or elective-level courses, SageMaker Studio Lab provides a free, no-credit-card-required environment for machine learning experimentation.

Examples of K Nearest Neighbors (KNN):

- **Recommendation Systems:** E-commerce platforms and media streaming services (like Netflix historically) use KNN to suggest products or movies to users. It finds users or items with similar preferences to recommend content that the user might enjoy.
- **Healthcare Diagnosis:** In clinical decision support systems, KNN helps predict diseases (like diabetes or heart conditions) by comparing a new patient's data with the historical data of similar patients and their diagnoses.
- **Credit Scoring and Fraud Detection:** Financial institutions employ KNN to assess the risk of granting a loan or to detect fraudulent transactions. It works by comparing new cases to past data of loan defaults or known fraudulent activities to determine creditworthiness or identify anomalies.

- **Image Recognition and Search Engines:** Early-stage image recognition systems and visual search features (such as on Pinterest) utilize KNN to retrieve similar images based on feature embeddings.
- **Ad Targeting and Customer Segmentation:** KNN helps companies like Target and Airbnb group customers based on their behavior or demographics. This allows for personalized marketing campaigns and targeted promotions to specific user clusters.
- **Geospatial Mapping:** Location-based services, including Google Maps and Uber, use the KNN algorithm to efficiently find the nearest drivers, restaurants, or points of interest in a given area.
- **Pattern Recognition:** The algorithm has been applied in identifying patterns, such as in text classification (marking emails as spam or not spam) and digit classification (interpreting handwritten numbers on forms).
- **Missing Data Imputation:** KNN can be used in data preprocessing to estimate and fill in missing values in a dataset by looking at the values of its nearest neighbors.

Case Studies—Responsible AI in Practice:

- **Finance & Insurance: The Depository Trust & Clearing Corporation (DTCC)**
 - Use Case: DTCC leverages generative AI, including Amazon Q Developer, for software development and legacy code modernization.
 - Responsible Approach: To mitigate risks related to hallucinations and intellectual property, DTCC focuses on responsible AI (RAI) frameworks that ensure code security and compliance, working closely with AWS and Anthropic to monitor output veracity.
- **Healthcare: Clinical Decision Support**
 - Use Case: Medical students and healthcare providers use large language models (LLM)-generated feedback during patient interactions to improve clinical decision-making.
 - Responsible Approach: These systems are designed with high-quality datasets to reduce diagnostic bias and adhere to strict data privacy regulations (HIPAA), ensuring AI supports rather than replaces human medical judgment.
- **Healthcare: Mental Health Triage AI**
 - Use Case: A healthcare chatbot designed to assist with mental health inquiries.
 - Responsible Approach: The bot uses "deflection" to handle topics outside its scope, guiding users to human support when necessary, and provides neutral, safe responses rather than diagnostic advice to minimize risk.
- **Public Sector: Data Reply and Red Teaming**
 - Use Case: Data Reply uses AWS to build conversational agents, incorporating red teaming to identify potential vulnerabilities.
 - Responsible Approach: They use Amazon Bedrock to build in-depth security tests that simulate adversarial attacks, ensuring the model's outputs are safe, accurate, and do not produce inappropriate content.
- **Corporate and Operational: PricewaterhouseCoopers (PwC) and Automated Reasoning**

- Use Case: PwC uses automated reasoning on Amazon Bedrock to evaluate and validate generative AI outputs.
- Responsible Approach: This method allows PwC to mathematically verify that AI applications comply with predetermined safety, privacy, and regulatory guidelines, reducing reliance on manual review for low-risk tasks.
- **Technology and Security: Trellix**
 - Use Case: Trellix uses Claude on Amazon Bedrock to analyze security alerts and enhance threat detection.
 - Responsible Approach: By implementing AI, they maintain 95 percent accuracy in security investigations while reducing manual effort, ensuring that automated actions remain controllable and secure.
- **Internal Amazon: Retail and Advertising**
 - Use Case: Amazon.com uses generative AI for tasks like summarizing customer reviews and generating advertising content.
 - Responsible Approach: These systems use content moderation, personally identifiable information (PII) redaction, and human-in-the-loop (HITL) checks to ensure that AI-generated content is appropriate and does not violate copyright or safety policies.
 - Key AWS Tools Used by Customers for Responsible AI
 - Amazon Bedrock Guardrails: Used to filter undesirable content, detect PII, and ensure safety across foundation models.
 - AWS AI Service Cards: Provide transparency on intended use cases, limitations, and performance metrics for services like Rekognition and Transcribe.
 - Amazon SageMaker Clarify: Used to detect bias in machine learning models.
 - AWS Well-Architected Responsible AI Lens: A framework that helps customers evaluate their AI systems against best practices.

Case Studies—Failure to Apply Responsible AI Principles:

- **Financial and Operational Failures:**
 - Zillow Offers (2021): The company's AI-powered home-buying algorithm overestimated property values, leading to a \$304 million operating loss, a 25 percent workforce layoff, and the closure of the division.
 - Air Canada Chatbot (2024): A chatbot provided incorrect information about the airline's bereavement policy. A court ruled that the airline was responsible for its agent's misinformation, establishing a precedent for AI liability.
 - Google's \$100 Billion Mistake (2023): Google's Bard AI shared incorrect information in a promotional demo, causing a massive drop in Alphabet's stock price.
- **Bias, Discrimination, and Real-World Harm:**
 - Wrongful Arrests (2020-Present): Robert Williams, a Black man in Detroit, was falsely arrested due to a flawed facial recognition algorithm.
 - Hiring Discrimination (2018): Amazon scrapped an AI recruiting tool that taught itself to discriminate against female candidates because it was trained on a decade of male-dominated resumes.

- Criminal Justice Bias (2016): ProPublica exposed the COMPAS algorithm, which was found to be biased against Black defendants in predicting recidivism rates.
- **Safety, Misinformation, and Ethical Lapses:**
 - Google AI Overviews (2024): The search tool suggested users eat rocks and use glue to make cheese stick to pizza, highlighting a lack of safety guardrails.
 - Deepfake Scam (2024): A French woman was swindled out of \$850,000 by a fraudster using advanced AI to impersonate Brad Pitt.
 - Dealer Chatbot Scam (2023): A Twitter user manipulated a Chevrolet dealership's AI chatbot into selling a 2024 Chevy Tahoe for \$1.
 - Google's Internal Ethics Conflict (2020): Google terminated researcher Dr. Timnit Gebru after she refused to withdraw a paper questioning the safety of large language models (LLM), leading to widespread criticism of their AI ethics commitment.
- **Environmental and Resource Neglect:**
 - Data Center Water Consumption (2025): The unchecked growth of AI data centers, such as those used by xAI in Memphis, has caused significant strain on local water resources, with some studies estimating a 100-megawatt center requires daily water consumption equivalent to 6,500 homes.

Additional/General List of AI Resource Links:

- [Adobe Firefly](#) – creative AI solution from Adobe
- [AI](#) – blog posts on AI from CompTIA
- [AI4K12.org](#) (The Artificial Intelligence [AI] for K-12 initiative [AI4K12] is jointly sponsored by AAAI and CSTA.)
 - [List of Resources](#), AI4K12.org
- [AI for Oceans](#) – hour of code course from Code.org
- [Artificial Intelligence](#) – hub of AI resources from MIT Media Lab
- Artificial Intelligence Literacy Framework, U.S. Department of Labor
 - [Attachment I \(Accessible PDF\).pdf](#)
 - [Attachment II \(Accessible PDF\).pdf](#)
- [Be My Eyes Accessibility with GPT-4o](#) - YouTube video highlighting GPT-4o from Open AI
- [ChatGPT](#) - AI assistant created by OpenAI
 - [ChatGPT Prompt Training Tips Guide](#), Innovation Training
- [Claude AI](#) - AI assistant created by Anthropic
- [Code.org](#) — AI and Machine Learning Modules Code.org offers structured, standards-aligned AI curriculum modules for high school students.
- [Cybersecurity and Infrastructure Security Agency \(CISA\) AI Resources](#) - As the nation's cyber defense agency, CISA's mission to secure federal software systems and critical infrastructure is critical to maintaining global AI dominance.
- [descript](#) – AI editing for every kind of video

- [Explore the Impact of AI on Jobs, Culture, and the Future](#) – a free, six-lesson-plan unit from Marginal Revolution University (MRU) (a nonprofit housed at the Mercatus Center at George Mason University), in partnership with the Institute for Global Affairs
- Google
 - [Google Labs](#) – the home for AI experiments at Google
 - [Experiments with Google](#) – Google hub with a collection of experiments
 - [Jupyter Notebooks via Google Colab](#) For students ready to explore Python and data science; runs entirely in the browser and integrates directly with Google Drive and can connect to AWS services for more advanced projects.
 - [Quick, Draw!](#) –a neural network that includes the world’s largest doodling data set, shared publicly to further machine learning research
 - [Teachable Machine](#) - web-based tool for creating machine learning models
- [gencraft](#) – AI art image and video generator
- [Hello History AI](#) - AI technology that creates interactive experiences that allow people to engage and learn from historical figures
- [HeyGen](#) – AI that turns text to video
- [IEEE Standards Association \(SA\): AI](#) – listing of AI resources from the Institute of Electrical and Electronics Engineers (IEEE) Standards Association (SA)
- [MagicSchool AI](#) – AI-powered tools for the classroom
- [Microsoft Bing Image Creator](#) – AI that generates images based on words
- [Microsoft Copilot](#) – AI companion from Microsoft
- [Machine Learning for Kids](#) - free tool introduces machine learning by providing hands-on experiences for training machine learning systems and building things with them
- [NIST AI Resource Center \(AIRC\)](#), National Institute of Standards and Technology (NIST), U.S. Department of Commerce; the AIRC offers access to relevant technical documents and resources (including software tools and guidance) to assist in the testing, evaluation, verification, and validation (TEVV) of AI.
- [Perplexity AI](#) – AI assistant created by Perplexity AI
- [Presidential Artificial Intelligence Challenge](#) The Presidential AI Challenge is a national challenge where K-12 youth, educators, mentors, and community teams come together to solve real-world problems in their communities using AI-powered solutions with an opportunity to showcase their solutions at a national level.
- [Scratch + AI Extensions](#) Scratch is a tool for introducing computational thinking. With AI extensions from MIT Media Lab, students can integrate machine learning models into their Scratch projects.
- [Slides AI](#) – AI tool for creating slideshow presentations from text
- [Synthesia](#) – AI platform for creating videos using synthetic human-like avatars and natural-sounding voices