Generative AI Foundations:

Harnessing the Power of AI-Driven Creation and Innovation



Courseware #7121



CCI Learning

Critical Career Skills Program

Generative Al Foundations

June 2025 © CCI Learning Solutions Inc

CCI Learning®

IGNITING POSSIBILITIES

Generative AI Foundations

This course book is one of a series prepared by CCI Learning Solutions Inc. for use by learners and instructors in courses on generative artificial intelligence (AI) foundations. CCI designed these materials to assist learners and instructors in making the learning process both effective and enjoyable.

This course book is copyrighted, and all rights are reserved by CCI Learning Solutions Inc. No part of this publication may be reproduced, transmitted, stored in a retrieval system, modified, or translated into any language or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual or otherwise without written permission of CCI Learning Solutions, Canada: 1-800-668-1669.

The information in this course book is distributed on an "as is" basis, without warranty. While every precaution has been taken in the preparation of this courseware, neither the author nor CCI Learning Solutions Inc. shall have any liability to any person or entity with respect to any liability, loss, or damage caused or alleged to be caused directly or indirectly by the instructions contained in this courseware or by the computer software and hardware products described therein.

Funded by the Government of Canada

Canada

Courseware Team:

Author: Jacob Wright Roberts, Mignonne van der Walt Editor: Heather Daniel, Lisa Hennessey, Sandi Langeman Copyright © 2025 CCI Learning Solutions Inc. ISBN: 978-1-55332-705-5 All rights reserved. Printed in Canada. CCI Courseware#: 7121

Any brand name or product mentioned in this publication is a trademark or registered trademark of their respective companies and are used for identification purposes only.

Unit 1: Methods

Unit Objectives

Welcome to Unit 1: Methods! In this unit, you will explore the concepts, processes, and tools that drive generative AI, distinguishing it from other AI types like predictive and analytical AI. You'll also examine the training requirements, customization options, and limitations of generative AI models, equipping you with the knowledge to select and utilize the right tools for specific tasks. Upon successful completion of this unit, you should be able to understand the following:

- □ Types of AI
- □ Processes
- □ Input and Output
- Tools and Systems
- □ Limitations

Lesson 1: Types of AI

Lesson Objectives

In this lesson, you will explore types of AI. Artificial intelligence comes in various forms, each intended to accomplish specific goals. Some forms of AI are meant to generate new content. Some are built to predict future outcomes. Some are programmed to analyze data. If we grasp the strengths and weaknesses of different forms of AI, we can decide when to use a certain tool for a certain task. Upon successful completion of this lesson, you should be able to understand the following:

- □ Artificial Intelligence
- □ Generative AI
- Predictive Al
- Discriminative Al
- Analytical AI
- Statistical AI
- □ AI vs. Search Engines

Artificial Intelligence



Artificial intelligence (AI) refers to the development of computer systems that can perform tasks that typically require human intelligence. Such tasks include visual perception, speech recognition, decision-making, and language translation.

Machine learning describes how some AI systems are designed to learn from data, adapt to new inputs, and improve over time. There are various types of AI, including narrow AI, which is designed for specific tasks, and general AI, which aims to perform

any intellectual task that a human can do.

Al has potential applications in various industries, from healthcare and finance to entertainment and education. The growth of Al technologies has led to significant advancements in automation, personalization, and data analysis, transforming how we interact with the world.

Artificial IntelligenceThe simulation of human intelligence processes by computer systems.Visual PerceptionThe ability to interpret and understand graphical information from the surrounding environment, such as images or video.Speech RecognitionThe process of converting spoken language into text by recognizing and interpreting human language.Decision- MakingThe ability to make choices or recommendations based on data.Al SystemA computer-based system that uses algorithms and data to perform tasks typically requiring human intelligence.Narrow AlA type of artificial intelligence designed to perform a specific task or a narrow range of tasks.AGI (Artificial General Intelligence)Atype of artificial intelligence that aims to perform any intellectual task that a human can do, still largely theoretical compared to current narrow Al.Data AnalysisThe process of tailoring services or content to individual preferences, sometimes using Al.Data AnalysisThe process of examining and interpreting information to extract useful facts and insights, often assisted by Al tools.Machine LearningA subset of Al that enables systems to learn from data and improve their performance over time without being explicitly programmed.NLP (Natural Language Processing)A field of artificial intelligence focused on enabling computers to understand, interpret, and generate human language.AlgorithmA set of rules or instructions given to an Al system to help it learn from data and make decisions.			
Perceptionsurrounding environment, such as images or video.Speech RecognitionThe process of converting spoken language into text by recognizing and interpreting human language.Decision- MakingThe ability to make choices or recommendations based on data.Al SystemA computer-based system that uses algorithms and data to perform tasks typically requiring human intelligence.Narrow AlA type of artificial intelligence designed to perform a specific task or a narrow range of tasks.AGI (Artificial General Intelligence)A type of artificial intelligence that aims to perform any intellectual task that a human can do, still largely theoretical compared to current narrow Al.AutomationThe use of technology to perform tasks without human intervention, often powered by Al.PersonalizationThe process of tailoring services or content to individual preferences, sometimes using Al.Data AnalysisThe process of examining and interpreting information to extract useful facts and insights, often assisted by Al tools.Machine LearningA subset of Al that enables systems to learn from data and improve their performance over time without being explicitly programmed.NLP (Natural Language Processing)A field of artificial intelligence focused on enabling computers to understand, interpret, and generate human language.AlgorithmA set of rules or instructions given to an Al system to help it learn from data		The simulation of human intelligence processes by computer systems.	
Recognitioninterpreting human language.Decision- MakingThe ability to make choices or recommendations based on data.AI SystemA computer-based system that uses algorithms and data to perform tasks typically requiring human intelligence.Narrow AIA type of artificial intelligence designed to perform a specific task or a narrow range of tasks.ACI (Artificial Intelligence)A type of artificial intelligence that aims to perform any intellectual task that a human can do, still largely theoretical compared to current narrow AI.AutomationThe use of technology to perform tasks without human intervention, often powered by AI.PersonalizationThe process of tailoring services or content to individual preferences, sometimes using AI.Data AnalysisThe process of examining and interpreting information to extract useful facts and insights, often assisted by AI tools.Machine LearningA subset of AI that enables systems to learn from data and improve their performance over time without being explicitly programmed.NLP (Natural Language Processing)A field of artificial intelligence focused on enabling computers to understand, interpret, and generate human language.AlgorithmA set of rules or instructions given to an AI system to help it learn from data			
MakingAl SystemAl SystemA computer-based system that uses algorithms and data to perform tasks typically requiring human intelligence.Narrow AIA type of artificial intelligence designed to perform a specific task or a narrow range of tasks.ACI (Artificial General Intelligence)A type of artificial intelligence that aims to perform any intellectual task that a human can do, still largely theoretical compared to current narrow AI.AutomationThe use of technology to perform tasks without human intervention, often powered by AI.PersonalizationThe process of tailoring services or content to individual preferences, sometimes using AI.Data AnalysisMachine LearningA subset of AI that enables systems to learn from data and improve their performance over time without being explicitly programmed.NLP (Natural Language Processing)AlgorithmA set of rules or instructions given to an AI system to help it learn from data	-		
typically requiring human intelligence.Narrow AIA type of artificial intelligence designed to perform a specific task or a narrow range of tasks.ACI (Artificial General Intelligence)A type of artificial intelligence that aims to perform any intellectual task that a human can do, still largely theoretical compared to current narrow AI.AutomationThe use of technology to perform tasks without human intervention, often powered by AI.PersonalizationThe process of tailoring services or content to individual preferences, sometimes using AI.Data AnalysisThe process of examining and interpreting information to extract useful facts and insights, often assisted by AI tools.Machine LearningA subset of AI that enables systems to learn from data and improve their performance over time without being explicitly programmed.NLP (Natural Language Processing)A field of artificial intelligence focused on enabling computers to understand, interpret, and generate human language.AlgorithmA set of rules or instructions given to an AI system to help it learn from data		The ability to make choices or recommendations based on data.	
AGI (Artificial General Intelligence)A type of artificial intelligence that aims to perform any intellectual task that a human can do, still largely theoretical compared to current narrow AI.AutomationThe use of technology to perform tasks without human intervention, often powered by AI.PersonalizationThe process of tailoring services or content to individual preferences, sometimes using AI.Data AnalysisThe process of examining and interpreting information to extract useful facts and insights, often assisted by AI tools.Machine LearningA subset of AI that enables systems to learn from data and improve their performance over time without being explicitly programmed.NLP (Natural Language Processing)A field of artificial intelligence focused on enabling computers to understand, interpret, and generate human language.AlgorithmA set of rules or instructions given to an AI system to help it learn from data	Al System		
General Intelligence)human can do, still largely theoretical compared to current narrow AI.AutomationThe use of technology to perform tasks without human intervention, often powered by AI.PersonalizationThe process of tailoring services or content to individual preferences, sometimes using AI.Data AnalysisThe process of examining and interpreting information to extract useful facts and insights, often assisted by AI tools.Machine LearningA subset of AI that enables systems to learn from data and improve their performance over time without being explicitly programmed.NLP (Natural Language Processing)A field of artificial intelligence focused on enabling computers to understand, interpret, and generate human language.AlgorithmA set of rules or instructions given to an AI system to help it learn from data	Narrow Al		
powered by Al.PersonalizationThe process of tailoring services or content to individual preferences, sometimes using Al.Data AnalysisThe process of examining and interpreting information to extract useful facts and insights, often assisted by Al tools.Machine LearningA subset of Al that enables systems to learn from data and improve their performance over time without being explicitly programmed.NLP (Natural Language Processing)A field of artificial intelligence focused on enabling computers to understand, interpret, and generate human language.AlgorithmA set of rules or instructions given to an Al system to help it learn from data	General		
sometimes using AI.Data AnalysisThe process of examining and interpreting information to extract useful facts and insights, often assisted by AI tools.Machine LearningA subset of AI that enables systems to learn from data and improve their performance over time without being explicitly programmed.NLP (Natural Language Processing)A field of artificial intelligence focused on enabling computers to understand, interpret, and generate human language.AlgorithmA set of rules or instructions given to an AI system to help it learn from data	Automation		
and insights, often assisted by AI tools.Machine LearningA subset of AI that enables systems to learn from data and improve their performance over time without being explicitly programmed.NLP (Natural Language Processing)A field of artificial intelligence focused on enabling computers to understand, interpret, and generate human language.AlgorithmA set of rules or instructions given to an AI system to help it learn from data	Personalization		
Learningperformance over time without being explicitly programmed.NLP (Natural Language Processing)A field of artificial intelligence focused on enabling computers to understand, interpret, and generate human language.AlgorithmA set of rules or instructions given to an AI system to help it learn from data	Data Analysis		
Language Processing)interpret, and generate human language.AlgorithmA set of rules or instructions given to an AI system to help it learn from data			
	Language		
	Algorithm		

Luciana, a data scientist, uses AI to analyze large datasets and provide insights that help her company optimize its marketing strategies. The AI adapts to new trends and predicts customer behavior, giving her team a competitive edge.

Learn Task 🐺

Access the Learner Workbook to complete the Learn Task for this skill.

Generative Al

Generative AI refers to a subset of artificial intelligence focused on creating content, such as images, text, music, or videos. These AI systems are generally trained on existing, human-generated examples. As the AI examines and analyzes the human-generated content, it begins to recognize or learn the underlying patterns. Then, when prompted, generative AI can produce content that follows the same patterns and thus closely resembles human-generated content. For example, generative AI that has been trained on an archive of articles from the New York Times might be able to produce text that matches the tone and style of a New York Times article. Generative AI has potential applications in industries ranging from the creative arts to business. However, this technology has limitations, as its outputs may be unreliable, inconsistent, or unoriginal.



Generative AI	A subset of artificial intelligence focused on creating content, such as images,
	text, music, or videos.

Graphic designer Tatiana uses generative AI to design the first draft of a logo for a client. She inputs the client's specifications, and the AI successfully produces a logo that meets most of the client's requests. However, Tatiana is concerned when she notices that the AI design closely resembles a logo from an established brand, posing potential legal issues.

Learn Task 🥳

Access the Learner Workbook to complete the Learn Task for this skill.

Predictive AI

Predictive AI is a type of artificial intelligence designed to forecast future events or trends based on



historical data. By analyzing patterns in data, predictive AI can make informed guesses about future outcomes, which is particularly useful in fields like finance, healthcare, and marketing. Predictive models use various algorithms, including regression analysis, to estimate the likelihood of future events. In contrast to generative AI, which creates new data, predictive AI focuses on interpreting existing data to provide predictions.

These predictions could help organizations or individuals make proactive decisions, improve efficiency, and reduce risk. Predictive AI is often used for stock market analysis and personalized marketing.

Predictive AI	A type of artificial intelligence that analyzes existing data to forecast future events or trends.
Regression Analysis	A statistical process for estimating the relationships among variables, often used in predictive modeling.

Predictive	The process of creating, testing, and validating a model to best predict the
Modeling	probability of an outcome.

Amahle, a marketing manager, uses predictive AI to analyze customer behavior and forecast the success of a new product launch. She also relies on her own judgement and checks the AI's work for errors.

Learn Task 🥳

Access the Learner Workbook to complete the Learn Task for this skill.

Discriminative Al



Discriminative AI models are designed to classify and distinguish between different categories or classes within a dataset. These models are commonly used in tasks such as image recognition, where they identify and categorize objects within an image. Discriminative models work by learning from labeled data, identifying the features that distinguish one category from another. The goal is to maximize accuracy in predicting which category a new piece of data belongs to. Unlike generative AI, which creates new

content, discriminative AI focuses on defining the differences between categories of data. As a result, discriminative AI can be used for spam detection, sentiment analysis, and medical diagnosis. Discriminative AI can make mistakes. As a result, it is important to confirm the conclusions of discriminative AI, particularly in sensitive or high-risk scenarios.

Discriminative Al	A type of artificial intelligence that focuses on distinguishing between different classes or categories within a dataset.	
Class	A distinct group or category that data points are organized into based on shared characteristics or features.	
Classification	he task of predicting the category of a given data point.	
Image Recognition	The process of identifying and classifying objects, people, or other elements within an image using AI.	
Spam Detection	The use of AI to identify and filter out unwanted or malicious emails from legitimate messages.	
Sentiment Analysis	The process of using AI to determine the emotional tone or opinion expressed in a piece of text.	

Akira, a software engineer, uses discriminative AI to develop a spam filter for his company's email system. The filter flags suspected spam and sends it to a separate spam folder. When Akira finds that many legitimate messages are sent to the spam folder, he attempts to fine-tune the program to improve its accuracy.

Learn Task 🐺

Access the Learner Workbook to complete the Learn Task for this skill.

Analytical Al

Analytical AI refers to AI systems designed to process, analyze, and interpret complex datasets,

providing insights that aid in decision-making. These systems are commonly used in fields such as finance, healthcare, and business intelligence. They can help to identify patterns, trends, and correlations within large volumes of data. While generative AI creates new content and discriminative AI categorizes data, analytical AI is focused on the extraction of meaningful information and conclusions from data. It often involves sophisticated algorithms that can handle unstructured data and



provide actionable insights. Analytical AI is crucial for understanding underlying dynamics in data and making informed predictions or strategic decisions based on that understanding.

Analytical Al	A type of artificial intelligence focused on processing and interpreting complex data to extract meaningful insights and support decision-making.
Unstructured Data	Information that is not organized in a pre-defined manner, making it difficult to analyze using traditional data processing methods.
Actionable	Clear, specific, and able to be directly used to make decisions or take steps.

Prescott uses analytical AI to interpret sales data and identify trends that could inform the company's future strategies. As a business analyst, he hopes to use the AI to uncover hidden patterns that he might not be able to detect manually.

		×17.	
	Task		
22100	I ASK	- (+ P) -	A

Access the Learner Workbook to complete the Learn Task for this skill.

Statistical AI



Statistical AI is an approach to artificial intelligence that relies on statistical methods to make predictions, classifications, or decisions based on data. It is grounded in probability theory and is commonly used in fields where understanding uncertainty is crucial, such as finance, healthcare, and social sciences. Unlike generative AI, which creates new data, statistical AI focuses on using statistical models to infer relationships and predict future outcomes. These models often involve estimating the probability of certain events based on

historical data. Statistical AI is vital for tasks that require a rigorous understanding of data variability and uncertainty.

Statistical AI	A type of artificial intelligence that uses statistical models and methods to make predictions, decisions, or classifications based on data.	
Probability	A measure of the likelihood that a specific event will occur.	
Probability Theory	A branch of mathematics that deals with the analysis and interpretation of random events and the likelihood of different outcomes.	
Uncertainty	ertainty The degree of unpredictability in outcomes or events, often modeled ar quantified in AI using probability.	
Data Variability	The extent to which points in a set differ from each other or from a central value, indicating diversity or inconsistency within the set.	

Rocco, a data scientist, uses statistical AI to predict customer churn for a telecom company, helping them retain clients by addressing issues before they lead to cancellations. The AI provides probabilities for different outcomes, which Rocco uses to focus the company's efforts where they are most needed.

Learn Task 😱

Access the Learner Workbook to complete the Learn Task for this skill.

Al vs. Search Engines

Artificial Intelligence and search engines both deal with large volumes of information, but they operate in different ways. Search engines, like Google, retrieve and rank relevant documents or

webpages based on a user's query, using algorithms designed for information retrieval. Al, on the other hand, involves systems that can understand, generate, and analyze information, offering more sophisticated responses or even performing complex tasks based on the data. While search engines find and present existing information, Al can interpret that information, draw insights, and generate new content or decisions. In recent years, Al technologies have been



integrated into search engines to improve query understanding, personalization, and the relevance of search results. Despite these overlaps, Al's capabilities extend beyond the functions of traditional search engines.

Search Engine	A software system that searches a database or the internet for information based on user queries and returns relevant results.
Google	A widely used search engine that indexes and retrieves vast amounts of information from the web in response to user queries.
Query	A request for information or data submitted to a search engine or database.
Search Algorithm	A set of rules and processes used by a search engine to determine and rank the most relevant results for a given query.
Information Retrieval	The process of obtaining relevant data or documents from a database in response to a user query.
Integration	The process of combining different systems, tools, or data sources to work together.
Query Understanding	The ability of a search engine to interpret the intent and context of a user's question to deliver the most relevant results.
Relevance Ranking	The process by which a search engine orders search results based on their relevance to the user's query.

Elena, a journalist, uses a search engine to gather information for an article. She then uses an Al assistant to analyze the data and provide a summary of the most critical points. Elena uses this as a jumping off point for her research and reporting, but she always goes back to check and confirm the information.

Learn Task 🥳	Access the Learner Workbook to complete the Learn Task for this skill.
Assessments	Access the Learner Workbook to complete the Practice Exercise and Practice Ouestions.

Lesson 2: Processes

Lesson Objectives

In this lesson, you will discover the basic processes that generative artificial intelligence uses to create content. You will discover how different types of AI models, such as text models, image models, and large language models (LLMs), are trained to perform specific tasks. By understanding how these models learn from vast datasets, identify patterns, and generate outputs, you will gain insight into the complexities of AI model training and the importance of computational resources. The objective is to provide a comprehensive overview of how AI systems are developed, highlighting key AI models. Upon successful completion of this lesson, you should be able to understand the following:

- Text Models
- □ Image Models
- □ Large Language Models
- □ Diffusion
- □ Transformer
- Variational Autoencoders
- □ Generative Adversarial Networks
- □ Convolutional Neural Networks
- Model Training

Text Models

Text models are a type of AI designed to process and generate human language. These models are



trained on vast amounts of text data, enabling them to understand context, generate text, and perform language tasks like translation, summarization, and question-answering. Examples of advanced text models include OpenAI's GPT series, Google Gemini, Anthropic Claude, and Meta's LLaMA. These models learn through a process called model training, where they analyze patterns in the data, learning the structure and rules of language. To perform effectively, these models require an enormous amount of training data, which includes a wide range of opinions and points of view. The effectiveness of

these models depends on the quality and diversity of the training data, as well as the algorithms used to fine-tune them for specific tasks.

Text Model	An AI model designed to process, understand, and generate human language.	
Training Data	Data used to teach AI models, enabling them to learn patterns, structures, and relationships to perform tasks effectively.	
Multilingual	The ability of an AI model to understand and generate text in multiple languages.	
Fine-Tuning	The process of making specific adjustments to a pre-trained AI model to improve its performance on a particular task.	
OpenAl	An artificial Intelligence research organization known for developing advanced AI models, including the GPT series.	
GPT (Generative Pre-trained Transformer)	A type of AI model developed by OpenAI, designed for generating human-like text based on large-scale pre-training.	
Google Gemini	A multilingual AI model developed by Google, known for its capability to handle diverse language tasks across different languages.	
Anthropic	An AI research company focused on developing AI systems that are safe, ethical, and aligned with human values.	
Anthropic Claude	An AI model developed by Anthropic, designed with a focus on safety, ethics, and alignment with human values.	
Meta	A technology company that develops social media platforms and AI research projects, including advanced AI models like LLaMA.	
LLaMA (Large Language Model Meta AI)	A research-focused AI model developed by Meta, optimized for studying language processing and AI model behavior.	

Mariana, a writer, uses OpenAI's GPT4 to help brainstorm ideas for her new novel. She also experiments with Google Gemini to translate her book into different languages.

Learn Task 🦷

Access the Learner Workbook to complete the Learn Task for this skill.

Image Models

Image models are AI systems designed to generate or interpret visual content, such as images and

graphics. These models are trained on large datasets of images paired with textual descriptions, which are often manually tagged, enabling them to recognize patterns, objects, and styles, and create or analyze visual content accurately. Examples of leading image models include DALL-E, developed by OpenAI, which can generate highly detailed and creative images from textual descriptions, and Adobe Firefly, designed for creative professionals to generate and enhance visual content with ease. Training these models is computationally intensive, requiring significant energy and powerful GPUs to process the



massive amounts of data involved. The quality of these models depends heavily on the diversity of their training data and the sophistication of their learning algorithms.

Image Model	An AI representation designed to generate or interpret visual content, such as images and graphics.	
Labeled Data	Images that have been tagged with descriptions or classifications, used to train AI models.	
Visual Content	Any matter that is viewed, such as images, graphics, or videos.	
Text-Image Pairs	Data sets where images are matched with corresponding textual descriptions, used in training image models.	
GPUs (Graphics Processing Units)	Specialized hardware used to accelerate the training of AI models, especially in tasks involving large datasets and complex computations.	
DALL-E	An AI model developed by OpenAI that generates images from textual descriptions, allowing users to create detailed and imaginative visuals based on prompts.	
Adobe Firefly	An AI-powered tool by Adobe designed to enhance creative workflows by generating images, effects, and other visual content from text-based inputs.	

Kwesi, a graphic designer, uses DALL-E to create artwork for his clients. He uses the program to generate images based on specific descriptions of his clients' visions.

Learn Task 🥳

Access the Learner Workbook to complete the Learn Task for this skill.

Large Language Models

Large language models (LLMs) are a class of AI models that have been trained on extensive datasets to understand and generate human language at a high level of complexity. LLMs can perform tasks such as summarization, translation, and even creative writing. LLMs are built using deep learning techniques and contain billions of parameters which are the adjustable elements of the model that influence how it processes information.



To perform effectively, LLMs require a vast amount of training data that encompasses a wide range of opinions and perspectives. The enormous scale of these datasets can enable LLMs to generate text that is contextually relevant and coherent across a wide range of topics. Training these models also consumes significant computational resources. The large amounts of energy and powerful GPUs needed make the process both resource-intensive and expensive.

LLM (Large Language Model)	An AI system trained on extensive text data to perform advanced language tasks.
Parameters	The adjustable elements within an AI model that influence how it processes and interprets data.
Contextual Relevance	The ability of an AI model to generate or interpret text that is appropriate to the given context.
Summarization	The process of condensing a large amount of information into a shorter, more digestible form.

Halyna, a researcher, uses a large language model to help her analyze academic papers and generate summaries. She carefully reviews the Al's conclusions for accuracy.

Learn Task 🙀

Access the Learner Workbook to complete the Learn Task for this skill.

Diffusion

Diffusion in the context of AI refers to a generative model that produces images by gradually refining random noise into a coherent image. The process involves starting with a noisy image and

iteratively applying a diffusion process that removes the noise, guided by a trained neural network. Diffusion models, like DALL-E, use this approach to generate highquality images from text descriptions, where the model learns to reverse the noise generation process and create detailed visuals.

In other words, diffusion in AI is like starting with a messy, blurry picture and gradually cleaning it up until you get a clear and detailed image. Imagine you have a photo that's been covered with static, like an old TV screen. The AI



slowly removes this static bit by bit, guided by what it has learned from other pictures, until the final image appears.

Diffusion	A generative process where an image is gradually refined from random noise into a coherent and detailed visual by iteratively removing the noise, guided by a trained neural network.
Iteratively	Repeatedly applying a process or set of steps multiple times, often with the goal of refining or improving the outcome with each repetition.
Neural Network	Computational models inspired by the human brain, used in AI to recognize patterns and make decisions.

Noise Generation	The process of adding random data (noise) to an image, which is later refined by the diffusion model to create a coherent image.
Refinement Algorithm	Processes used in diffusion models to iteratively reduce noise and produce a clear, high-quality image.
Dataset	A collection of information used to train AI models, often consisting of images, text, or other types of information.

Marta uses a diffusion model to generate concept art for her upcoming exhibition. She inputs a brief description into the tool and watches as the model gradually transforms random noise into an image.

Learn Task 🦷

Access the Learner Workbook to complete the Learn Task for this skill.

Transformer

A transformer is a type of neural network architecture designed for handling sequential data, such as text, by processing entire sequences in parallel rather than step-by-step. Transformers are particularly effective in tasks like language modeling and translation because they use mechanisms called attention layers to weigh the importance of different parts of the input data. This architecture underpins many advanced models, including GPT, enabling them to perform complex tasks like text generation, tokenization, and understanding context.



A transformer is a type of AI that reads and understands a whole sentence or piece of text all at once, instead of going through it word by word. It figures out which parts of the text are most important and focuses on these parts to understand the meaning. This helps transformers translate languages and write text because they can quickly grasp the context and generate accurate and meaningful sentences.

Transformer	A type of neural network architecture that processes sequential data, like text, by considering entire sequences in parallel and using attention mechanisms to understand context and relationships within the data.
Sequential Data	A type of information that is ordered or arranged in a specific sequence, such as time-series information or text, where the order of elements is important for understanding the context.
In Parallel	Refers to performing multiple tasks or processes simultaneously rather than one after the other.
Tokenization	The process of breaking down text into smaller units, like words or subwords, that can be processed by a transformer model.
Attention Layers	The components of a transformer that allow the model to focus on specific parts of the input data when making predictions.
Hyperparameters	The settings that define the architecture and behavior of a neural network, such as learning rate and number of layers.
Learning Rate	A hyperparameter that controls how much the model's guesses are adjusted with each step during training.

Jorge, a software developer, uses a transformer-based model to automate the translation of technical documents into multiple languages.

Learn Task 🕷

Access the Learner Workbook to complete the Learn Task for this skill.

Variational Autoencoders

Variational autoencoders (VAEs) are a type of generative model that learns to encode input data into a compressed representation and then decode it back into the original or similar data. Unlike

traditional autoencoders, VAEs introduce variability in the encoding process, allowing them to generate new data that is similar to the input dataset. VAEs are used in applications like image generation, where they can create variations of images by sampling from the learned latent space. In other words, variational autoencoders take something, like a picture, and shrink it down to a basic version of itself. Then, they use that basic version to recreate the picture, but with a twist. They can add small changes to make new, similar pictures. VAEs are



great for making new images that appear similar to the ones they learned from, with slight changes.

Variational Autoencoder	A type of generative model that learns to encode input data into a compressed representation and then decodes it to generate new, similar data by introducing variability in the encoding process.
Latent Space	A compressed representation of data learned by a model, from which new variations can be generated.

Graphic designer Lloyd uses a VAE model to generate variations of his logo designs, exploring different styles and colors. The model helps him quickly create a variety of logos, giving his client more options to choose from.

Learn Task 🐺

Access the Learner Workbook to complete the Learn Task for this skill.

Generative Adversarial Networks

Generative adversarial networks (GANs) consist of two neural networks, a generator and a discriminator, that work together in a competitive process to create realistic data. The generator creates data samples, while the discriminator evaluates them against real data, providing feedback



to improve the generator's outputs. This adversarial process can result in the generation of realistic images, audio, or other types of data, with GANs being widely used in applications like image synthesis, video generation, and deepfakes.

In other words, generative adversarial networks are like a game between two players. One creates something new, like a picture, and the other tries to figure out if it's real or fake. The first player, called the generator, keeps getting better at making realistic pictures because the second

player, the discriminator, gives feedback on how convincing they are. Over time, this back-and-forth

competition helps the generator make images, sounds, or videos that appear so real, it's hard to tell they were made by a computer.

GAN (Generative Adversarial Network)	A type of AI model that consists of two neural networks, a generator and a discriminator, which compete against each other to create highly realistic data, such as images or videos.
Generator	The part of a GAN that creates new data samples, attempting to mimic real data from the dataset.
Discriminator	The component of a GAN that evaluates the generated data and helps the generator improve by distinguishing between real and fake data.

Filmmaker Ren uses a GAN to generate a deepfake of a celebrity. Ren posts the content on Instagram. Content moderators flag the post as misinformation, and it is eventually removed. Ren may face legal consequences for his misleading use of the celebrity's image and voice.

<u>Learn T</u>ask 🥳

Access the Learner Workbook to complete the Learn Task for this skill.

Convolutional Neural Networks

Convolutional Neural Networks (CNNs) are a specialized class of deep learning models primarily used for processing and analyzing visual data, such as images and videos. Unlike traditional neural networks, which treat each input pixel independently, CNNs apply convolutional layers to detect patterns like edges, textures, and shapes at different levels of abstraction. As the data passes through successive layers, the CNN learns increasingly complex features. Eventually, the CNN can



classify objects, recognize faces, or detect anomalies with high accuracy.

In other words, CNNs are like smart visual detectors that scan images in pieces, recognize patterns, and combine these pieces to understand and classify the whole picture. Instead of examining each part of an image separately, they learn to perceive the big picture by recognizing important features like edges and textures, just like how our brains process visual information.

CNN (Convolutional Neural Network)	A type of deep learning model designed to analyze visual data by automatically detecting and learning features through layers.
Convolutional Layer	A component in a neural network that applies filters to input data to capture important visual features like edges and textures.
Image Classification	The task of categorizing images into predefined classes based on their visual content using models like CNNs.

Yoshi uses a CNN to build an app that can identify different types of flowers from photos. The CNN detects the flowers in order to classify them.

Learn Task 🧣

Access the Learner Workbook to complete the Learn Task for this skill.

Model Training

Model training is the process by which AI systems learn from data to perform specific tasks, such as generating text, recognizing images, or making predictions. This process involves feeding the AI large amounts of data and adjusting its internal parameters to improve its performance on the task

at hand. During training, the model learns to identify patterns, relationships, and structures within the data, which it uses to make accurate predictions or generate content. The training process can be supervised, where the model learns from labeled data, or unsupervised, where it discovers patterns in data without explicit guidance. Training а model is requiring computationally intensive, significant energy and powerful GPUs, especially for large-scale models like LLMs or image models. Fine-tuning is often applied after initial training to optimize the



model for specific applications, enhancing its accuracy and relevance.

Model Training	The process of teaching an AI system to perform specific tasks by learning from data.
Supervised Learning	A training method where the AI model learns from labeled data with known outcomes.
Unsupervised Learning	A training method where the AI model identifies patterns in data without explicit labels.
Computational Resources	The hardware and software tools required to train AI models, including processors, memory, and storage.
Energy Consumption	The amount of energy required to power the computational resources used during model training.

Machine learning engineer Mei Ling spends weeks training a new AI model on a large dataset of medical images. Her objective is to teach the AI to detect early signs of disease. Mei Ling is proud when the model achieves high accuracy. However, she is concerned about the potential consequences of misdiagnosis and continues to work to improve the model's accuracy results.



Lesson 3: Input and Output

Lesson Objectives

In this lesson, you will encounter the fundamental concepts of inputs and outputs in AI. You will learn how different types of data are processed to produce various forms of content. You will explore the types of inputs (such as text, audio, video, and images) and outputs (including generative text, video, images, and audio), and learn how AI tools can be customized to meet specific needs. Additionally, you will explore the process of selecting the right AI tool, considering factors such as functionality, ease of use, cost, data privacy, and output control to ensure the tool aligns with your objectives. Upon successful completion of this lesson, you should be able to understand the following:

- Understand Input and Output
- Input Types
- Output Types
- □ Customization
- □ AI Tools
- Select a Tool

Understand Input and Output

In the context of AI, inputs refer to the data or information that is fed into a model for processing,



which can include text, audio, video, images, or a combination of these. Outputs are the results generated by the AI after processing the inputs. The relationship between inputs and outputs is fundamental to AI operations, as the quality and type of input directly affect the nature and quality of the output. Different AI models are optimized for specific types of input and output, making it crucial to match the right model with the right task. Understanding how inputs are transformed into outputs helps users better choose, customize, and use AI tools for tasks like content creation, data analysis, and automation.

Input	The data or information provided to an AI model for processing, such as text, images, or audio.
Output	The result generated by an AI model after processing the input, such as generated text, video, images, or audio.

Rania, a content creator, inputs a script into an AI tool to generate a promotional video. She finetunes the visuals and adds personalized elements to match her brand's style.

Learn Task 🥷

Access the Learner Workbook to complete the Learn Task for this skill.

Input Types

Input types refer to the different forms of data that can be fed into an AI model for processing. Common input types include text, which might consist of written words or commands; audio, which can include spoken language, music, or sound effects; video, involving moving images and

sound; and images, which are static visual representations. The choice of input type depends on the task at hand and the capabilities of the AI model being used. For example, text inputs are commonly used in language models for generating text-based responses, while image inputs are used in models designed for visual tasks like object recognition or image generation. Understanding the variety of input types helps in selecting the appropriate AI tool and ensuring that the model processes the data correctly to produce the desired output.



Text	Written words or commands provided as input to an AI model.
Audio	Sound recordings, including speech or music, used as input for processing by AI models.
Video	Moving visual content, often accompanied by sound, used as input for AI models designed to handle dynamic media.
Image	Static visual representations used as input for AI models focused on tasks like recognition or generation.

Koko, a podcaster, inputs a recorded audio file into an AI tool designed to enhance sound quality. The tool outputs a new audio file with less background noise.

```
Learn Task 🐺
```

Access the Learner Workbook to complete the Learn Task for this skill.

Output Types

Output types refer to the various forms of content that an AI model can generate after processing



input data. Common output types include generative text, where the model produces written content based on the input it receives; generative video, which involves creating video content from scripts or visual descriptions; generative image, where the model creates images from textual or visual input; and generative audio, which involves producing sound or music based on input parameters. Each output type is tailored to specific tasks and applications, such as generating articles, creating videos, producing artwork, or composing music. The versatility of Al

models in generating different output types makes them valuable tools across multiple industries, from content creation to entertainment and marketing.

Generative Text	Written content produced by an AI model based on input data.
Generative Video	Video content created by an AI model from scripts, storyboards, or visual descriptions.
Generative Image	Visual content generated by an AI model from textual or visual input.
Generative Audio	Sound or music produced by an AI model based on input data.

Filmmaker Sven inputs a storyboard into an AI tool to generate an animated video.

Learn Task 🐺

Access the Learner Workbook to complete the Learn Task for this skill.

Customization

Customization in AI refers to the ability to tailor AI models to perform specific, individualized tasks

according to the user's needs. Generative AI models can be customized through various means, including fine-tuning existing models, adjusting parameters, or developing self-contained applications that perform targeted functions. Examples of customized AI tools include Custom GPT, where users can create tailored AI chatbots, and Google Gems, designed for specific language tasks. Customization allows users to optimize AI models for specific use cases, ensuring that the outputs meet their particular requirements. This flexibility makes AI tools more



relevant and effective in various contexts, from business automation to creative projects.

Customization	The process of tailoring AI models to perform specific, individualized tasks.
Custom GPT	A tailored AI chatbot that customizes OpenAI's GPT models to perform specific functions for users.
Google Gems	A tailored AI chatbot that customizes Google's Gemini AI for specific language- related tasks.

Olli, a project manager, customizes Google Gems to automate repetitive tasks in their team's workflow.

Learn Task 🙀 Access the Learner Workbook to complete the Learn Task for this	skill.
--	--------

AI Tools



Al tools are software applications or platforms that leverage artificial intelligence to perform specific tasks, ranging from content creation to data analysis and productivity enhancement. Examples of popular AI tools include Microsoft Copilot, which assists with tasks in Word and Excel; Google Gemini, which integrates with Google applications; Meta GPT, serving Meta's social media platforms; Adobe Express, which uses AI for graphic design and content creation;

and OpenAl's ChatGPT, a conversational AI model that interacts with users in natural language. Other tools like Canva use AI to simplify graphic design, while Claude by Anthropic focuses on ethical AI applications, and Microsoft Azure AI Studio provides a platform for building and deploying AI models. Stable Diffusion is another tool used for generating high-quality images from textual descriptions. Each tool has its unique strengths and applications, making them valuable in different fields.

Adobe Express	A tool that uses AI to assist with graphic design and content creation.
Canva	An online design tool that incorporates AI to simplify graphic design tasks.
Microsoft Azure Al Studio	A platform for building, deploying, and managing AI models.
Stable Diffusion	An AI tool used for generating high-quality images from textual descriptions.

Marketing professional Lara uses Canva's AI tools to quickly create visually appealing social media posts. She also uses OpenAI's ChatGPT to draft engaging content for her campaigns before fine-tuning the text to match her brand's voice.

Learn Task 🙀

Access the Learner Workbook to complete the Learn Task for this skill.

Select a Tool

Selecting the right AI tool involves evaluating various factors to ensure it meets the user's needs

and is suitable for the intended purpose.



One key consideration is the tool's purpose and functionality. What use is the tool intended for? Does that use match your goals? Can the tool perform the specific tasks required?

Also consider ease of use. Is the tool user-friendly? Is it accessible? Can it be used by people of varied abilities, ages, and cultural backgrounds?

Cost is another important factor. What initial investment is required to use the tool? Are there ongoing expenses associated with continued use? How does the value of the tool compare to the cost?A fourth consideration is updates and support. Is

technical support available? Are there regular updates to maintain the tool's effectiveness?

Do not overlook privacy and security. Does the tool comply with relevant regulations? Does the tool protect user data? Is the tool vulnerable to cyberattacks?

Quality is paramount. Are the outputs accurate? Are they consistent? Are they reliable? Are they original?

Consider customizability. Can the tool be tailored to meet specific needs?

Finally, the parameters available for output control are also crucial. Can you adjust the model's output? Is the tool responsive to feedback and guidance? Does the tool allow you to fine-tune its behavior to produce your desired results?

Purpose	The specific objective or intended use of a tool or system.
Functionality	The range of tasks and operations that a tool or system can perform.
Ease of Use	The degree to which a tool or system is user-friendly and intuitive to operate.
Cost	The financial expenditure required to acquire, maintain, and operate a tool or system.
Updates	The regular improvements or patches provided to a tool or system to enhance performance or security.
Support	The assistance provided by the tool or system's provider to help users resolve issues or optimize usage.
Data Privacy	The protection of personal and sensitive information from unauthorized access or disclosure.
Security	The measures taken to safeguard a tool or system from threats such as cyberattacks or data breaches.
Quality	The degree to which a tool or system produces reliable and accurate results.
Customizability	The ability of a tool or system to be tailored to specific user needs or preferences.
Output Control	The settings and parameters that allow users to influence the nature and quality of the outputs generated by a tool or system.

Theo, a small business owner, selects an AI tool to manage his company's social media. His priorities are security and user-friendliness. After comparing several tools, he chooses one that offers data privacy features and customer support.

Learn Task 🥳	Access the Learner Workbook to complete the Learn Task for this skill.
Assessments 🔰	Access the Learner Workbook to complete the Practice Exercise and Practice Questions.

Lesson 4: Tools and Systems

Lesson Objectives

In this lesson, you will explore specific AI tools and systems. Through practical examples and scenarios, you'll learn how these tools can support tasks like writing, data analysis, customer support, and creative design. Upon successful completion of this lesson, you should be able to understand the following:

- □ Conversational Models
- □ ChatGPT
- □ Microsoft Copilot
- □ Google Gemini
- 🛛 Meta Al
- □ Adobe Express
- □ Claude
- □ Microsoft Azure Al Studio
- □ Stable Diffusion
- DALL-E
- □ Adobe Firefly

Conversational Models

Conversational AI models, such as OpenAI's ChatGPT, are designed to generate human-like text based on the input they receive. They can assist with a range of tasks, which include answering questions, providing explanations, and generating creative content. These chatbots are sometimes used for customer support, content creation, or as a virtual assistant to enhance productivity.



To use a conversational AI model, you start by entering a text prompt or question into the chat interface. The AI processes the input and generates a text-based response, which you can then review or use as appropriate.

Conversational AI may be useful when you need to generate large amounts of text quickly, such as drafting emails, writing articles, or brainstorming ideas. They may also be effective for answering questions in customer support scenarios, where they can provide fast responses based on a large database of knowledge. Certain conversational AI models

are integrated into applications or productivity software to support in-app tasks.

Conversational Al	Artificial intelligence models designed to interact with users in natural language, simulating human dialogue.
Virtual Assistant	A software agent that can perform tasks or services for an individual based on user input.

Yoko uses a conversational AI model to help her brainstorm ideas for a speech at her sister's wedding. She then refines the ideas and adds her own insights and anecdotes.

Learn Task 🦷

Access the Learner Workbook to complete the Learn Task for this skill.

ChatGPT

OpenAl's ChatGPT is a conversational AI model. ChatGPT can be accessed via OpenAl's website. Developers can also integrate ChatGPT into their applications through application programming interfaces (APIs). ChatGPT partners with Microsoft to offer ChatGPT in Microsoft applications through Microsoft Copilot, and ChatGPT is also available as a bot in the workplace software Slack.



The term "Chat GPTx" is often used to describe the evolving versions of OpenAl's GPT models, where "x" serves as a placeholder for the specific version number (e.g., GPT-3, GPT-3.5, GPT-4). It represents the family of models that are designed for conversational AI and natural language processing tasks. As the models evolve, "GPTx" signifies the evolving AI capabilities offered by each new version.

API (Application Programming Interface)	A set of tools and protocols that allow different software applications to communicate with each other.
GPTx	The evolving versions of OpenAI's GPT models, where "x" serves as a placeholder for the specific version number (e.g., GPT-3, GPT-3.5, GPT-4).

Fernando asks ChatGPT4 for recommendations of healthy options for dinner. The AI suggests a recipe for quinoa with roasted eggplant, and Fernando cooks the dish that night.

Learn Task 😭 Access the Learner Workbook to complete the Learn Task for this skill.

Microsoft Copilot

Microsoft Copilot is a tool that integrates Open AI's technologies, including ChatGPT, into Microsoft applications. Copilot functions as an AI-powered productivity tool built into software like Microsoft



Word, Excel, PowerPoint, Outlook, and Teams. Copilot can assist with in-app tasks such as writing, data analysis, and document formatting. It can provide suggestions and automate repetitive tasks, which may allow Microsoft users to focus on other aspects of their work.

In Word, Copilot can assist with writing and editing by suggesting text completions, generating content based on prompts, and adjusting grammar and style. It can also summarize long documents, suggest relevant content based on the context, and help with formatting.

In Excel, Copilot can help analyze data by creating charts and summarizing trends. It can automate repetitive tasks like formula creation, data cleaning, and pivot table generation.

In PowerPoint, Copilot can assist in creating presentations by generating slide content, suggesting layouts, and adjusting design elements. It can also translate text-heavy slides into bullet points.

In Outlook, Copilot can aid with email management by suggesting responses, summarizing lengthy email threads, and scheduling meetings based on context. It can draft emails from brief prompts and help prioritize messages, making inbox management more efficient.

In Teams, Copilot can summarize meeting notes, generate action items, and provide real-time language translation. It can also assist with creating agendas and drafting follow-up emails.

Productivity Tool	Software designed to help users complete tasks more efficiently.
Productivity Tool	Software designed to help users complete tasks more efficiently.

Jaime, a financial analyst, spends much of his day creating and analyzing charts. Jaime uses Copilot in Excel to automate some of this work, freeing up time in his schedule for other critical tasks.

Learn Task 😭 Access the Learner Workbook to complete the Learn Task for t	his skill.
---	------------

Google Gemini

Google Gemini is an AI tool that integrates with Google applications like Docs, Sheets, Slides, and Gmail.



In Google Docs, Google Gemini can assist with writing and editing by offering real-time language suggestions, grammar corrections, and style improvements. It can also help generate text based on prompts, summarize documents, and translate content into different languages directly within the document.

In Google Sheets, Gemini can help analyze data and generate summaries. Users can ask questions about their data, and Gemini will generate responses.

In Google Slides, Gemini can help generate text for slides, suggest visual elements, and ensure that content is concise and well-structured. It also assists with language translation, making it convenient to create multilingual presentations.

In Gmail, Gemini can suggest responses and help with translations. It can also generate full email drafts based on brief prompts.

Multilingual	Possessing the ability to understand, generate, or process text in multiple
	languages.

Luca uses Google Gemini in Gmail to draft a reply to their colleague's email inquiry about an upcoming deadline. They then review, edit, and send the email.

```
Learn Task 🥳
```

Access the Learner Workbook to complete the Learn Task for this skill.

Meta Al

Meta Al is a suite of Al tools that integrate with Meta's social media applications, such as Facebook, Instagram, Messenger, and WhatsApp. On Facebook, Instagram, and Snapchat, users can use Meta Al to assist with content creation by generating captions, hashtags, or entire posts based on prompts. The Al can suggest posting times and content types to increase engagement. Meta Al can also enhance photos and videos by applying filters and improving image quality.

For group admins on Facebook, Meta AI can identify and remove harmful content, recommend moderation actions, and suggest ways to increase engagement within the group.



On Messenger and WhatsApp, Meta AI enables smart replies, automated customer support, and chatbots. Businesses can use AI-driven chatbots to engage with customers, answer frequently asked questions, and provide real-time assistance.

Filters	Tools or features that modify or enhance the appearance of images or videos, often by applying visual effects, color changes, or overlays.
Smart Replies	Response suggestions that help users quickly reply to messages or emails.
Automated Customer Support	Systems that handle consumer inquiries without the need for human intervention, often through chatbots or automated messaging.

Portia uses Meta AI's tools to apply filters to her photographs before sharing them with her family and friends.

Learn Task 🕷

Access the Learner Workbook to complete the Learn Task for this skill.

Adobe Express

Adobe Express is an AI-powered graphic design tool that simplifies the creation of professionalquality visuals for various purposes, including social media, marketing, and branding. It offers a wide range of templates and design elements that can be customized to fit the user's needs. Adobe Express integrates AI features that help automate design tasks, making it accessible even to those with little design experience.

To use Adobe Express, start by selecting a template or creating a design from scratch using the platform's tools. You can customize the design by adding images, text, and other elements, with the AI offering suggestions and adjustments to enhance the visual appeal. Once the design is complete, you can export it in various formats for use in different platforms or media.

Adobe Express is ideal for small business owners, marketers, and social media managers who need to create eye-catching visuals quickly without requiring advanced design skills. It's also useful for individuals seeking to enhance their personal projects, such as creating invitations, posters, or presentations. Additionally, Adobe Express is beneficial for content creators who need to produce consistent and branded visuals across different platforms.



Template	A pre-designed layout that can be customized to create various types of documents or visuals.
Visual Appeal	The attractiveness of a design, which can be enhanced through the use of colors, layouts, and imagery.

Leslie, a small business owner, uses Adobe Express to create promotional graphics for his new product launch. This allows him to focus on other aspects of his business while still maintaining a visual presence.

Learn Task 😽

Access the Learner Workbook to complete the Learn Task for this skill.

Claude

Claude is an AI model developed by Anthropic, focused on creating ethical and safe AI applications that align with human values. Claude is designed to minimize biases, avoid harmful outputs, and operate within strict ethical guidelines. It is particularly useful for organizations that prioritize ethical AI use and need reliable tools that adhere to high standards of safety and fairness.

To use Claude, input your queries or tasks into the platform that supports the model, and it will generate responses or perform actions based on the ethical frameworks embedded in its design. Claude is often integrated into customer service platforms, content moderation systems, or any application where ethical considerations are paramount. Users can customize the model's behavior by defining specific ethical parameters or guidelines that align with their organizational values.

Claude is especially useful for companies that need to ensure their AI tools are free from bias and are operating within ethical standards, such as in content moderation, automated decisionmaking, or customer interactions. It's also beneficial for educational institutions and non-profits that require AI tools to adhere to strict ethical guidelines.

*claude Do your best work	Claude, convert this design into a react component using Tailwind CSS. Team Budget \$12,567 of \$25,000
Start using Claude for yourself or your team G Continue with Google OR name@yourcompany.com Continue with email	Image: state of the state
Learn more ↓	

Content Moderation	The process of monitoring and managing user-generated content to ensure it complies with standards and policies.
Human Values	The moral principles and beliefs that guide the behavior and decisions of individuals and organizations.

Ollie, a policy advisor, integrates Claude into his organization's content moderation system. He aims to maintain the ethical standards of the organization while managing a large volume of content efficiently.

Learn Task 🥳

Access the Learner Workbook to complete the Learn Task for this skill.

Microsoft Azure Al Studio

Microsoft Azure AI Studio is a comprehensive platform that enables users to build, deploy, and manage AI models and applications. It provides a suite of tools for machine learning, data analysis, and AI model development, allowing users to create customized AI solutions for various business needs. Azure AI Studio is highly scalable, making it suitable for projects ranging from small experiments to large-scale enterprise applications. To use Microsoft Azure AI Studio, you start by selecting the AI tools and services you need from the platform, such as machine learning models or data analytics tools. You can then build and train your AI models using Azure's resources, deploying them directly to your applications. The platform also offers monitoring and management features to help you optimize performance and ensure your models are functioning as expected.

Azure AI Studio Develop and deploy generative AI apps and APIs responsibly with a comprehensive platform		Construint C			
Build copilots and AI solutions faster with prebuilt and customizable models, using your data to innovate at scale.	0 Aferen	Anny to top tasks and too	seed. September 2014 Sectors 2014 Australia Arrian Array Control Control Array Control Control Array Control Control Array Control Array Control Array Control Control Array Control Array Control Control Array Control Array Array Control Control Array Control Array Array Control Control Control Array Control Array Control Control Control Control Control Control Control Control Control Array Control		And hading a south In 120% and incommendation ID Designs large language models (LDM) Gamps an UMP prompt from and marks for the animate for some production adversers (I
Build with Azure Al Studio Get started with Azure		Explore cutting-edge mod	Lexten Dennis helf par il austant Sells Mitte Lance 3 (19) instruct Sensible des pression	Say in this solve polytic in the plagment stagement Colours command - plase Colours Ladie 473 (September)	Type & I or on Andrewson's Mitted Surger Mitted With and Antoneousless,
				<	

Microsoft Azure AI Studio is ideal for businesses and developers who need to create custom AI solutions, such as automating workflows, analyzing large datasets, or developing new products. It's also useful for research teams working on AI innovations, providing the necessary tools and computational resources. Additionally, Azure AI Studio is beneficial for organizations seeking to integrate AI into their existing systems and scale as their needs grow.

Deployment	The process of making a software application or AI model operational and
	accessible to users.

Deegan, a data scientist, uses Microsoft Azure Al Studio to develop a machine learning model that predicts customer behavior. The platform's scalability lets him handle increasing data volumes as the project expands.

Learn Ta		
learn la	- Jul -	
LCUITIO		

Access the Learner Workbook to complete the Learn Task for this skill.

Stable Diffusion

Stable Diffusion is an AI tool designed for generating high-quality images from textual descriptions, allowing users to create visuals based on specific prompts or concepts. It uses advanced machine learning techniques to interpret text and produce corresponding images, often used in creative fields like graphic design, marketing, and content creation. Stable Diffusion is known for its ability to generate detailed and visually appealing images that align closely with the input descriptions.

To use Stable Diffusion, you input a textual description or prompt into the tool, specifying the type of image you want to generate. The AI processes the input and produces an image that matches the description, which you can then download or further refine. Users can adjust parameters such as style, color, and composition to achieve the desired visual outcome.

Stable Diffusion is particularly useful for artists, designers, and marketers who need to create custom visuals quickly without the need for extensive manual design work. It's also beneficial for content creators who want to generate unique images for blogs, social media, or digital marketing campaigns. Additionally, Stable Diffusion can be used in prototyping and conceptualization, helping teams visualize ideas before moving to full-scale production.



Prototyping The process of creating a preliminary model or concept to test ideas before full-scale development.

Galinda uses Stable Diffusion to generate concept art for her latest project. She fine-tunes the Al's output by adjusting the color palette and composition to make sure the final images align with her artistic style.

Learn Task 🦷 🛛 🗛	ess the Learner Workbook to complete the Learn Task for this skill.
------------------	---

DALL-E

DALL-E is an AI tool developed by OpenAI that generates images from textual descriptions, allowing users to create visuals based on prompts. It uses a diffusion model to turn words into pictures. DALL-E is sometimes used in creative industries, where it helps artists, designers, and marketers bring their ideas to life without needing advanced graphic design skills.



To use DALL-E, you start by entering a detailed description of the image you want to create. The AI processes this input and attempts to generate an image that matches the description, which you can then refine by adjusting the prompt or generating new variations. Users can experiment with different descriptions to explore the range of creative possibilities that DALL-E offers, generating images for projects or inspiration.

DALL-E could be useful for marketing campaigns, product concepts, or certain artistic projects. It can also be used for brainstorming sessions, where quickly generated visuals could help communicate ideas and concepts. Additionally, DALL-E can be used by individuals or small businesses without access to a graphic design team.

Textual
DescriptionA written input provided to an AI model to generate corresponding images or
other outputs.

Ronald, a marketing manager, uses DALL-E to brainstorm visuals for an upcoming ad campaign. He inputs a description of a futuristic cityscape, and the AI generates several pictures. Ronaldo does not use the images in the campaign because of his concerns about their originality. However, they serve as valuable inspiration.

Learn Task 😽

Access the Learner Workbook to complete the Learn Task for this skill.

Adobe Firefly

Adobe Firefly is an AI-powered creative tool that can help users generate and enhance visual content, offering features like image generation, effects, and style transfers. It is designed to integrate with Adobe's suite of design tools, allowing Adobe users to automate or accelerate parts of their workflow.

Adobe	Creativity & Design \checkmark	Adobe Firefly	Products ~			Sign in
me / Adobe	Creative Cloud / Adobe Fireft	,				
			Create w	ith Firefly	/	
			genera	tive Al.		
		Firefly m		eatures in Adobe creative apps. Learn abo Lightroom, and the Firefly web app.	ut the latest	
	C Genedit Inge Activity reception	SR Dry				
	Generative Fill		Text to Image	Generative Shape Fill	Generative Remove	
	The next generation Fill is now in the Pho app. Create richer, m images with more or before. See how in Photosh	otoshop (beta) Iore lifelike ontrol than ever	With the new Firefly Image 3 Model, you can create higher- quality images with better composition, photorealistic details, and improved mood and lighting. Create with Firefly	With new Generative Shape Fill, you can quickly fill your vector outline and explore a variety of options that match the look and feel of your own artwork. See how in Illustrator	See everything you want and nothing you don't. It's easy to remove distractions from the background with Generative Remove powered by Firefly generative AL	

To use Adobe Firefly, you can start by selecting a base image or creating one from scratch using the tool's AI-powered features. You can then apply various effects, styles, or adjustments to the image, guided by Firefly's AI suggestions. Once satisfied, you can export the final image for use in different media or use Adobe products like Photoshop or Illustrator for further refinement. Adobe Firefly might be useful in situations where a designer needs to quickly generate multiple visual concepts for a client presentation or project. It can also apply different styles or effects to existing images, allowing for experimentation. Firefly may also serve users with limited time or design skill, as it can eliminate the need to spend hours on time-consuming manual work.

Style Transfer	A technique that applies the visual style of one image to another, often used in creative applications.
Workflow	The sequence of processes through which a piece of work passes from initiation to completion, especially in creative industries.

Nikki uses Adobe Firefly to create a series of posters for a music festival. They quickly apply different styles and effects to the base images, producing a variety of designs.

Learn Task 🥳	Access the Learner Workbook to complete the Learn Task for this skill.
Assessments 🔰	Access the Learner Workbook to complete the Practice Exercise and Practice Questions.

Lesson 5: Limitations

Lesson Objectives

In this lesson, you will engage with the limitations of artificial intelligence. Al technologies come with significant limitations that users must understand to use them effectively and responsibly. This lesson explores key challenges such as the reliability of Al outputs, the technological resources required, privacy concerns, the absence of universal standards, issues with consistency, and the problem of obsolescence. By understanding these limitations, you can better navigate the complexities of Al and ensure that your Al use is ethical, secure, and aligned with your goals. Upon successful completion of this lesson, you should be able to understand the following:

- □ Reliability
- Tech Requirements
- □ Privacy
- Lack of Standards
- □ Consistency
- □ Obsolescence

Reliability

Generative AI models, while powerful, produce outputs that may not always be reliable. These



outputs can often include biases that reflect the data the model was trained on, leading to unfair or skewed results. Additionally, AI models can generate misinformation, especially if the input data is incomplete or incorrect. As a result, AI use can lead to the spread of false information. Another common issue is AI hallucinations, where the model produces outputs that are entirely fabricated and have no basis in reality. These limitations are critical to consider when using AIgenerated content, as they can impact the accuracy and trustworthiness of the information. Ensuring that AI

models are used responsibly requires ongoing monitoring and validation of the outputs.

Reliability	The degree to which an AI model consistently produces accurate and trustworthy outputs.
Bias	A tendency of AI models to reflect and perpetuate prejudices found in the training data, leading to unfair or unbalanced outputs.
Misinformation	The false or inaccurate information that may be generated by AI, often due to flawed or incomplete input data.
AI Hallucination	An instance where an artificial intelligence model generates content that is entirely fabricated and has no connection to reality.

Kodzo uses a generative AI tool to draft an article. He notices that some of the information it produced is inaccurate and reflects biases he hadn't intended to include. He must spend extra time fact-checking and revising the output to ensure the final article is fair, accurate, and bias-free.

Learn Task 👾

Access the Learner Workbook to complete the Learn Task for this skill.

Tech Requirements

Generative AI models require significant technological resources to function effectively. High processing power is necessary to handle the complex computations involved in generating



content, which typically involves using powerful GPUs (Graphics Processing Units). Additionally, generative AI often needs continuous access to large datasets, which are usually accessible via the internet, to produce accurate and relevant outputs. The requirement for such substantial computational resources and data access means that generative AI tools can be costly and may not be accessible to everyone. These tech requirements also raise concerns about energy consumption and environmental impact, making the efficient use of

resources an important consideration.

Technological	The hardware, software, and data necessary to develop, operate, and maintain	
Resources	AI models and systems.	
Accessibility	The ease with which individuals or organizations can use and benefit from technology, including AI tools.	
-------------------------	--	
Environmental Impact	The effect that technological operations, including AI, have on the environment, particularly in terms of energy consumption and resource use.	
Efficiency	The ability to achieve desired outcomes with minimal waste of time, energy, or resources.	

Maria-Elena, a graphic designer, tries to use a generative AI tool on her laptop, but she quickly realizes that her device doesn't have the necessary processing power. She decides to use a cloud-based service that offers the necessary computational resources, but at an additional cost.

```
Learn Task 👾
```

Access the Learner Workbook to complete the Learn Task for this skill.

Privacy



Privacy is a significant concern when using generative AI. Many AI tools automatically collect and store data to improve their models. This practice can lead to privacy violations if users are unaware or unable to control how their data is used. To limit these risks, it's essential for users to understand and manage the privacy settings of AI tools, ensuring that their personal information is protected.

Privacy	The right of individuals to control the collection, use, and sharing of their personal information.
Data Collection	The process by which AI tools gather information from users, often for the purpose of improving model accuracy.
Data Protection	The measures taken to ensure that user information is kept secure and is not exposed to unauthorized access.
Personal Information	The data that can be used to identify an individual, such as name, address, or contact details.
Privacy Settings	The options within a software or platform that allow users to manage how their personal information is collected, used, and shared.
Sensitive Information	The data that requires extra protection due to its potential impact on privacy, such as financial details or health records.
Privacy Violation	A breach or unauthorized use of personal information that infringes on an individual's privacy rights.

Project manager Yerma uses an AI tool to draft a sensitive company report. Later, she realizes that the tool could use the content for future training. Concerned about privacy, she adjusts the settings to ensure that her data is not stored or used by the AI provider in the future.

Learn	ask	-(40)-
Leann	Tusk	

Access the Learner Workbook to complete the Learn Task for this skill.

Lack of Standards

There are no universal standards governing the use of generative AI. This leads to inconsistencies in how these technologies are developed and deployed. This lack of standardized guidelines means

that different AI models may operate under varying ethical considerations, privacy protocols, and levels of transparency. The absence of universally accepted standards can create challenges in ensuring that AI tools are used responsibly and do not cause harm. As the use of generative AI becomes more widespread, the development of industrywide standards will be essential to establish trust and accountability in AI technologies. These standards could help in areas such as bias mitigation, data privacy, and the ethical use of AI-generated content.



Standards	A set of established guidelines or rules that dictate the acceptable practices and quality in a particular field.
Ethical Considerations	The principles that guide decision-making and behavior, especially in ensuring fairness, integrity, and respect for individuals.
Transparency	The practice of being open and clear about processes, decisions, and data, allowing for accountability and trust.
Accountability	The responsibility of individuals or organizations to answer for their actions and ensure they are following set standards and ethics.
Protocols	The formal procedures or systems of rules that are followed in specific situations to ensure consistency and reliability.
Bias Mitigation	The strategies and actions taken to reduce or eliminate bias in processes, decisions, or outcomes.

Danijela is frustrated by the lack of clear guidelines for implementing AI in her software development projects. She advocates for her company to adopt stricter ethical standards, emphasizing the need for responsible AI use.

<u>Lea</u>rn Task 🥳

Access the Learner Workbook to complete the Learn Task for this skill.

Consistency



Generative AI models often face challenges with consistency, meaning that the outputs they produce can vary even when given similar inputs. This inconsistency can be problematic in situations where reliable and uniform results are required, such as in content creation or data analysis. The issue arises because AI models may interpret inputs slightly differently each time they are processed, leading to variations in the output. Understanding this limitation is crucial for users who need consistent and repeatable outputs from AI tools.

Consistency	The degree to which an AI model produces the same output when given the same input across multiple instances.
Variability	The differences in output that can occur when an AI model processes similar inputs.
Uniformity	The quality of being consistent and the same across different instances.

Dragomir, a filmmaker, notices that when he inputs the same prompt into an AI tool multiple times, it generates different versions of a video clip. He must manually adjust the outputs to ensure they align with his vision.

```
Learn Task 👹
```

Access the Learner Workbook to complete the Learn Task for this skill.

Obsolescence

Obsolescence refers to the process by which AI models and technologies become outdated or less effective over time as newer and more advanced models are developed. This is a significant concern



in the rapidly evolving field of AI, where technological advancements can render existing models obsolete quickly. As a result, users and developers must continually update and adapt to new technologies to maintain competitive and effective AI solutions. This cycle of obsolescence can lead to increased costs and resource investment, as older systems may need to be replaced or significantly upgraded. It's important for organizations to plan for obsolescence by staying informed about technological trends and ensuring that their AI strategies are flexible and scalable.

Obsolescence	The process of becoming outdated or no longer useful due to advancements in technology.	
Obsolete	No longer in use or replaced by a newer, more effective alternative.	
Evolution	The gradual development or progression of something, particularly in terms of improvement or adaptation over time.	
Technological Advancements	The progress and improvement in technology that can lead to the development of newer, more effective tools and models.	
Flexibility	The ability to adapt or change in response to new conditions or demands.	
Scalability	The capacity of a system or technology to expand and manage increased demand or workload effectively.	
Competitive Edge	The advantage that an individual or organization has over its competitors, often due to superior technology, innovation, or strategy.	

Davi, a business owner, invests in an AI tool for customer service. Within a year, a new and more advanced version is released. He faces the dilemma of whether to upgrade to the new version or continue using the older, now less competitive, tool.

Learn Task 🥳	Access the Learner Workbook to complete the Learn Task for this skill.
Assessments 🔰	Access the Learner Workbook to complete the Practice Exercise and Practice Questions.
Unit Assessment 🛐	Access the Learner Workbook to complete the Objective Assessment and Create Project.

Unit 1 Key Terms

Term	Definition
Accessibility	The ease with which individuals or organizations can use and benefit from technology, including AI tools.
Accountability	The responsibility of individuals or organizations to answer for their actions and ensure they are following set standards and ethics.
Actionable	Clear, specific, and able to be directly used to make decisions or take steps.
Adobe Express	A tool that uses AI to assist with graphic design and content creation.
Adobe Firefly	An AI-powered tool by Adobe designed to enhance creative workflows by generating images, effects, and other visual content from text-based inputs.
AGI (Artificial General Intelligence)	A type of artificial intelligence that aims to perform any intellectual task that a human can do, still largely theoretical compared to current narrow Al.
AI Hallucination	An instance where an artificial intelligence model generates content that is entirely fabricated and has no connection to reality.
Al System	A computer-based system that uses algorithms and data to perform tasks typically requiring human intelligence.
Algorithm	A set of rules or instructions given to an AI system to help it learn from data and make decisions.
Analytical AI	A type of artificial intelligence focused on processing and interpreting complex data to extract meaningful insights and support decision-making.
Anthropic	An AI research company focused on developing AI systems that are safe, ethical, and aligned with human values.
Anthropic Claude	An AI model developed by Anthropic, designed with a focus on safety, ethics, and alignment with human values.
API (Application Programming Interface)	A set of tools and protocols that allow different software applications to communicate with each other.
Artificial Intelligence	The simulation of human intelligence processes by computer systems.
Attention Layers	The components of a transformer that allow the model to focus on specific parts of the input data when making predictions.
Audio	Sound recordings, including speech or music, used as input for processing by AI models.
Automated Customer Support	Systems that handle consumer inquiries without the need for human intervention, often through chatbots or automated messaging.

Term	Definition
Automation	The use of technology to perform tasks without human intervention, often powered by AI.
Bias	A tendency of AI models to reflect and perpetuate prejudices found in the training data, leading to unfair or unbalanced outputs.
Bias Mitigation	The strategies and actions taken to reduce or eliminate bias in processes, decisions, or outcomes.
Canva	An online design tool that incorporates AI to simplify graphic design tasks.
Class	A distinct group or category that data points are organized into based on shared characteristics or features.
Classification	The task of predicting the category of a given data point.
Competitive Edge	The advantage that an individual or organization has over its competitors, often due to superior technology, innovation, or strategy.
Computational Resources	The hardware and software tools required to train AI models, including processors, memory, and storage.
Consistency	The degree to which an AI model produces the same output when given the same input across multiple instances.
Content Moderation	The process of monitoring and managing user-generated content to ensure it complies with standards and policies.
Contextual Relevance	The ability of an AI model to generate or interpret text that is appropriate to the given context.
Conversational AI	Artificial intelligence models designed to interact with users in natural language, simulating human dialogue.
Convolutional Layer	A component in a neural network that applies filters to input data to capture important visual features like edges and textures.
CNN (Convolutional Neural Network)	A type of deep learning model designed to analyze visual data by automatically detecting and learning features through layers.
Cost	The financial expenditure required to acquire, maintain, and operate a tool or system.
Custom GPT	A tailored AI chatbot that customizes OpenAI's GPT models to perform specific functions for users.
Customizability	The ability of a tool or system to be tailored to specific user needs or preferences.
Customization	The process of tailoring AI models to perform specific, individualized tasks.

Term	Definition
DALL-E	An AI model developed by OpenAI that generates images from textual descriptions, allowing users to create detailed and imaginative visuals based on prompts.
Data Analysis	The process of examining and interpreting information to extract useful facts and insights, often assisted by AI tools.
Data Collection	The process by which AI tools gather information from users, often for the purpose of improving model accuracy.
Data Privacy	The protection of personal and sensitive information from unauthorized access or disclosure.
Data Protection	The measures taken to ensure that user information is kept secure and is not exposed to unauthorized access.
Data Variability	The extent to which points in a set differ from each other or from a central value, indicating diversity or inconsistency within the set.
Dataset	A collection of information used to train AI models, often consisting of images, text, or other types of information.
Decision-Making	The ability to make choices or recommendations based on data.
Deployment	The process of making a software application or AI model operational and accessible to users.
Diffusion	A generative process where an image is gradually refined from random noise into a coherent and detailed visual by iteratively removing the noise, guided by a trained neural network.
Discriminative AI	A type of artificial intelligence that focuses on distinguishing between different classes or categories within a dataset.
Discriminator	The component of a GAN that evaluates the generated data and helps the generator improve by distinguishing between real and fake data.
Ease of Use	The degree to which a tool or system is user-friendly and intuitive to operate.
Efficiency	The ability to achieve desired outcomes with minimal waste of time, energy, or resources.
Energy Consumption	The amount of energy required to power the computational resources used during model training.
Environmental Impact	The effect that technological operations, including AI, have on the environment, particularly in terms of energy consumption and resource use.
Ethical Considerations	The principles that guide decision-making and behavior, especially in ensuring fairness, integrity, and respect for individuals.
Evolution	The gradual development or progression of something, particularly in terms of improvement or adaptation over time.

Term	Definition
Filters	Tools or features that modify or enhance the appearance of images or videos, often by applying visual effects, color changes, or overlays.
Fine-Tuning	The process of making specific adjustments to a pre-trained AI model to improve its performance on a particular task.
Flexibility	The ability to adapt or change in response to new conditions or demands.
Functionality	The range of tasks and operations that a tool or system can perform.
GAN (Generative Adversarial Network)	A type of AI model that consists of two neural networks, a generator and a discriminator, which compete against each other to create highly realistic data, such as images or videos.
Generative AI	A subset of artificial intelligence focused on creating content, such as images, text, music, or videos.
Generative Audio	Sound or music produced by an AI model based on input data.
Generative Image	Visual content generated by an AI model from textual or visual input.
Generative Text	Written content produced by an AI model based on input data.
Generative Video	Video content created by an AI model from scripts, storyboards, or visual descriptions.
Generator	The part of a GAN that creates new data samples, attempting to mimic real data from the dataset.
Google	A widely used search engine that indexes and retrieves vast amounts of information from the web in response to user queries.
Google Gemini	A multilingual AI model developed by Google, known for its capability to handle diverse language tasks across different languages.
Google Gems	A tailored AI chatbot that customizes Google's Gemini AI for specific language-related tasks.
GPT (Generative Pre-trained Transformer)	A type of AI model developed by OpenAI, designed for generating human-like text based on large-scale pre-training.
GPTx	The evolving versions of OpenAI's GPT models, where "x" serves as a placeholder for the specific version number (e.g., GPT-3, GPT-3.5, GPT-4).
GPUs (Graphics Processing Units)	Specialized hardware used to accelerate the training of AI models, especially in tasks involving large datasets and complex computations.
Human Values	The moral principles and beliefs that guide the behavior and decisions of individuals and organizations.
Hyperparameters	The settings that define the architecture and behavior of a neural network, such as learning rate and number of layers.

Term	Definition
Image	Static visual representations used as input for AI models focused on tasks like recognition or generation.
Image Classification	The task of categorizing images into predefined classes based on their visual content using models like CNNs.
Image Model	An AI representation designed to generate or interpret visual content, such as images and graphics.
Image Recognition	The process of identifying and classifying objects, people, or other elements within an image using Al.
In Parallel	Refers to performing multiple tasks or processes simultaneously rather than one after the other.
Information Retrieval	The process of obtaining relevant data or documents from a database in response to a user query.
Input	The data or information provided to an AI model for processing, such as text, images, or audio.
Integration	The process of combining different systems, tools, or data sources to work together.
Iteratively	Repeatedly applying a process or set of steps multiple times, often with the goal of refining or improving the outcome with each repetition.
Labeled Data	Images that have been tagged with descriptions or classifications, used to train AI models.
Latent Space	A compressed representation of data learned by a model, from which new variations can be generated.
Learning Rate	A hyperparameter that controls how much the model's guesses are adjusted with each step during training.
LLaMA (Large Language Model Meta Al)	A research-focused AI model developed by Meta, optimized for studying language processing and AI model behavior.
LLM (Large Language Model)	An Al system trained on extensive text data to perform advanced language tasks.
Machine Learning	A subset of AI that enables systems to learn from data and improve their performance over time without being explicitly programmed.
Meta	A technology company that develops social media platforms and Al research projects, including advanced Al models like LLaMA.
Microsoft Azure Al Studio	A platform for building, deploying, and managing AI models.
Misinformation	The false or inaccurate information that may be generated by AI, often due to flawed or incomplete input data.

Term	Definition
Model Training	The process of teaching an AI system to perform specific tasks by learning from data.
Multilingual	The ability of an AI model to understand and generate text in multiple languages.
Narrow Al	A type of artificial intelligence designed to perform a specific task or a narrow range of tasks.
Neural Network	Computational models inspired by the human brain, used in Al to recognize patterns and make decisions.
NLP (Natural Language Processing)	A field of artificial intelligence focused on enabling computers to understand, interpret, and generate human language.
Noise Generation	The process of adding random data (noise) to an image, which is later refined by the diffusion model to create a coherent image.
Obsolescence	The process of becoming outdated or no longer useful due to advancements in technology.
Obsolete	No longer in use or replaced by a newer, more effective alternative.
OpenAl	An artificial Intelligence research organization known for developing advanced AI models, including the GPT series.
Output	The result generated by an AI model after processing the input, such as generated text, video, images, or audio.
Output Control	The settings and parameters that allow users to influence the nature and quality of the outputs generated by a tool or system.
Parameters	The adjustable elements within an AI model that influence how it processes and interprets data.
Personal Information	The data that can be used to identify an individual, such as name, address, or contact details.
Personalization	The process of tailoring services or content to individual preferences, sometimes using AI.
Predictive AI	A type of artificial intelligence that analyzes existing data to forecast future events or trends.
Predictive Modeling	The process of creating, testing, and validating a model to best predict the probability of an outcome.
Privacy	The right of individuals to control the collection, use, and sharing of their personal information.
Privacy Settings	The options within a software or platform that allow users to manage how their personal information is collected, used, and shared.

Term	Definition
Privacy Violation	A breach or unauthorized use of personal information that infringes on an individual's privacy rights.
Probability	A measure of the likelihood that a specific event will occur.
Probability Theory	A branch of mathematics that deals with the analysis and interpretation of random events and the likelihood of different outcomes.
Productivity Tool	Software designed to help users complete tasks more efficiently.
Protocols	The formal procedures or systems of rules that are followed in specific situations to ensure consistency and reliability.
Prototyping	The process of creating a preliminary model or concept to test ideas before full-scale development.
Purpose	The specific objective or intended use of a tool or system.
Quality	The degree to which a tool or system produces reliable and accurate results.
Query	A request for information or data submitted to a search engine or database.
Query Understanding	The ability of a search engine to interpret the intent and context of a user's question to deliver the most relevant results.
Refinement Algorithm	Processes used in diffusion models to iteratively reduce noise and produce a clear, high-quality image.
Regression Analysis	A statistical process for estimating the relationships among variables, often used in predictive modeling.
Relevance Ranking	The process by which a search engine orders search results based on their relevance to the user's query.
Reliability	The degree to which an AI model consistently produces accurate and trustworthy outputs.
Scalability	The capacity of a system or technology to expand and manage increased demand or workload effectively.
Search Algorithm	A set of rules and processes used by a search engine to determine and rank the most relevant results for a given query.
Search Engine	A software system that searches a database or the internet for information based on user queries and returns relevant results.
Security	The measures taken to safeguard a tool or system from threats such as cyberattacks or data breaches.
Sensitive Information	The data that requires extra protection due to its potential impact on privacy, such as financial details or health records.

Term	Definition
Sentiment Analysis	The process of using AI to determine the emotional tone or opinion expressed in a piece of text.
Sequential Data	A type of information that is ordered or arranged in a specific sequence, such as time-series information or text, where the order of elements is important for understanding the context.
Smart Replies	Response suggestions that help users quickly reply to messages or emails.
Spam Detection	The use of AI to identify and filter out unwanted or malicious emails from legitimate messages.
Speech Recognition	The process of converting spoken language into text by recognizing and interpreting human language.
Stable Diffusion	An AI tool used for generating high-quality images from textual descriptions.
Standards	A set of established guidelines or rules that dictate the acceptable practices and quality in a particular field.
Statistical AI	A type of artificial intelligence that uses statistical models and methods to make predictions, decisions, or classifications based on data.
Style Transfer	A technique that applies the visual style of one image to another, often used in creative applications.
Summarization	The process of condensing a large amount of information into a shorter, more digestible form.
Supervised Learning	A training method where the AI model learns from labeled data with known outcomes.
Support	The assistance provided by the tool or system's provider to help users resolve issues or optimize usage.
Technological Advancements	The progress and improvement in technology that can lead to the development of newer, more effective tools and models.
Technological Resources	The hardware, software, and data necessary to develop, operate, and maintain AI models and systems.
Template	A pre-designed layout that can be customized to create various types of documents or visuals.
Text	Written words or commands provided as input to an AI model.
Text Model	An AI model designed to process, understand, and generate human language.
Text-Image Pairs	Data sets where images are matched with corresponding textual descriptions, used in training image models.

Term	Definition
Textual Description	A written input provided to an AI model to generate corresponding images or other outputs.
Tokenization	The process of breaking down text into smaller units, like words or subwords, that can be processed by a transformer model.
Training Data	Data used to teach AI models, enabling them to learn patterns, structures, and relationships to perform tasks effectively.
Transformer	A type of neural network architecture that processes sequential data, like text, by considering entire sequences in parallel and using attention mechanisms to understand context and relationships within the data.
Transparency	The practice of being open and clear about processes, decisions, and data, allowing for accountability and trust.
Uncertainty	The degree of unpredictability in outcomes or events, often modeled and quantified in AI using probability.
Uniformity	The quality of being consistent and the same across different instances.
Unstructured Data	Information that is not organized in a pre-defined manner, making it difficult to analyze using traditional data processing methods.
Unsupervised Learning	A training method where the AI model identifies patterns in data without explicit labels.
Updates	The regular improvements or patches provided to a tool or system to enhance performance or security.
Variability	The differences in output that can occur when an AI model processes similar inputs.
Variational Autoencoder	A type of generative model that learns to encode input data into a compressed representation and then decodes it to generate new, similar data by introducing variability in the encoding process.
Video	Moving visual content, often accompanied by sound, used as input for Al models designed to handle dynamic media.
Virtual Assistant	A software agent that can perform tasks or services for an individual based on user input.
Visual Appeal	The attractiveness of a design, which can be enhanced through the use of colors, layouts, and imagery.
Visual Content	Any matter that is viewed, such as images, graphics, or videos.
Visual Perception	The ability to interpret and understand graphical information from the surrounding environment, such as images or video.
Workflow	The sequence of processes through which a piece of work passes from initiation to completion, especially in creative industries.

Thrive in an Al-driven Future

Expand Your Learning Experience with XperienceED, CCI Learning's Kinesthetic Learning Platform:



LEARN MORE: https://products.ccilearning.com

Prepare for the official Generative AI Foundations certification with this essential course.

Unlock the potential of Generative AI with a comprehensive introduction to cutting-edge AI technologies. From understanding the differences between Generative AI and other types of AI, to exploring the models that drive text, image, and video generation, this course equips you with the skills and knowledge needed to thrive in the rapidly evolving field of artificial intelligence.

You'll be guided through the core methodologies behind Generative AI models, from transformers to GANs, gaining a deeper understanding of how they generate text, images, audio, and more.

In addition to foundational knowledge, you'll dive into the art of prompt engineering—learning to craft, refine, and optimize prompts to elicit specific outputs for various tasks. The course also covers the practical applications of these tools in real-world scenarios, from creating content to transforming media formats.

Beyond technical proficiency, this course emphasizes the ethical and societal implications of Generative AI. You'll explore key concerns such as data privacy, intellectual property, and the potential biases that AI systems can introduce.

Whether you're just starting out, looking to expand your expertise to stay ahead in an Al-driven future, or aiming to understand the broader impact of these technologies, this course provides a solid foundation to help you succeed in both professional and creative contexts as Al continues to transform industries.

Key Learning Outcomes:

- Understand Core Generative Al Concepts: Explore essential Al methods, including transformers, GANs, and diffusion models, and how they produce outputs like text, images, video, and audio.
- Master Basic Prompt Engineering: Learn how to create and modify effective prompts to generate and transform content, tailoring outputs for specific tasks and media formats.
- **Refine Prompt Techniques:** Improve prompt clarity and precision using common strategies, such as zero-shot and few-shot prompting, for more accurate and targeted AI responses.
- **Recognize Ethical and Societal Impacts:** Understand key ethical considerations, such as bias, data privacy, and legal implications, and explore the broader societal impacts of AI use in various fields.





© CCI Learning Solutions Inc.

www.ccilearning.com