



# GENERATIVE AI INTERVIEW PREPARATION GUIDE



Visit Our Website  
[www.gsdccouncil.org](http://www.gsdccouncil.org)



## 1. What is Artificial Intelligence?

Artificial Intelligence is the field of building systems that can perform tasks requiring human intelligence, such as reasoning, learning, perception, and decision-making. Traditional AI relied heavily on predefined rules and logic created by humans. These systems worked well only in controlled environments. As data and complexity increased, rule-based AI struggled to scale. This limitation led to the development of Machine Learning and later Generative AI.

## 2. What is Machine Learning and how is it different from traditional AI?

Machine Learning is a subset of AI where systems learn patterns from data instead of relying on fixed rules. Unlike traditional AI, ML models improve their performance as they are exposed to more data. Machine Learning focuses on prediction and pattern recognition. It uses algorithms such as regression, decision trees, and clustering. However, Machine Learning typically predicts outcomes rather than creating new content.

## 3. What is Deep Learning?

Deep Learning is a specialized branch of Machine Learning that uses neural networks with multiple layers. These layers allow the model to learn complex patterns from large amounts of unstructured data like text, images, and audio. Each layer extracts increasingly abstract features from the data. Deep Learning is the foundation behind modern breakthroughs in speech recognition and image processing. Generative AI would not be possible without Deep Learning.

## 4. What is the difference between Generative and Discriminative models?

Discriminative models focus on distinguishing between different classes of data. They learn the boundary between inputs and outputs, such as classifying emails as spam or not spam. Generative models, on the other hand, learn how the data itself is generated. Because of this, they can create new data that looks similar to real data. In simple terms, discriminative models predict labels, while generative models create content.

## 5. What is Generative AI?

Generative AI refers to AI systems that can generate new content such as text, images, audio, video, or code. These systems learn patterns and structures from large datasets during training. Instead of giving a single prediction, they generate outputs step by step. Generative AI is probabilistic in nature, meaning it predicts the most likely next output. This capability makes it useful for creative and cognitive tasks.

## 6. What are foundation models?

Foundation models are large, pre-trained AI models that can be adapted to many different tasks. They are trained on massive and diverse datasets. Instead of building a model from scratch, organizations customize foundation models using prompts or fine-tuning. This approach saves time and cost while improving scalability. Examples include large language models like GPT and PaLM.

## 7. What are Transformers and why are they important?

Transformers are a type of neural network architecture used in modern Generative AI systems. They process entire sequences of data simultaneously instead of one step at a time. The key innovation in transformers is the self-attention mechanism. Self-attention allows the model to understand context and relationships between words. This architecture enables better language understanding and long-context processing.

## 8. What is prompt engineering?

Prompt engineering is the practice of designing inputs that guide a Generative AI model's behavior. A prompt defines the task, context, constraints, and expected output. Since the model follows instructions literally, prompt clarity is critical. Well-crafted prompts significantly improve output quality. Prompt engineering is often faster and cheaper than fine-tuning models.

## 9. What are Large Language Models (LLMs)?

Large Language Models are deep learning models trained on vast amounts of text data. They learn language patterns, grammar, facts, and reasoning styles. LLMs generate responses by predicting the next word based on context. They do not truly understand meaning like humans but simulate understanding statistically. LLMs power chatbots, summarization tools, and coding assistants.

## 10. What is fine-tuning and why is it needed?

Fine-tuning is the process of training a pre-trained model further on domain-specific data. It helps align the model with organizational needs or specialized use cases. Fine-tuning improves accuracy, tone, and relevance. However, it requires quality data and additional resources. In many cases, prompt engineering or RAG is preferred before fine-tuning.

## 11. What is Agentic AI?

Agentic AI refers to AI systems that can act autonomously toward a goal. Unlike standard Generative AI, agentic systems can plan, make decisions, and execute multiple steps. They often use tools, APIs, and feedback loops. Agentic AI works in reasoning cycles rather than single responses. This makes it suitable for complex workflows and autonomous processes.

## 12. How is Agentic AI different from Generative AI?

Generative AI primarily responds to user prompts by generating content. Agentic AI goes a step further by taking actions to achieve a goal. It can break tasks into sub-steps and decide what to do next. Agentic systems also evaluate their own outputs and adjust behavior. In short, Generative AI assists, while Agentic AI acts.

## 13. What are hallucinations in Generative AI?

Hallucinations occur when a model generates confident but incorrect or fabricated information. They happen because the model predicts text based on probability, not factual verification. Hallucinations are more common when prompts are vague or data is missing. This is a major risk in regulated industries. Techniques like RAG and human review help reduce hallucinations.

## 14. Why is Human-in-the-Loop important in Generative AI?

Human-in-the-loop ensures that humans review and validate AI outputs. This improves accuracy, accountability, and trust. AI can make mistakes or reflect biases from training data. Human oversight helps catch errors before decisions are finalized. This approach is critical for ethical and responsible AI use.

## 15. What are the key ethical concerns in Generative AI?

Key ethical concerns include bias, fairness, transparency, and data privacy. Models can unintentionally reproduce harmful biases from training data. There are also risks related to intellectual property and misinformation. Without governance, AI outputs may cause legal or reputational harm. Responsible AI practices help mitigate these risks.

## 16. What is self-attention in simple terms?

Self-attention is a mechanism that allows a model to focus on the most relevant parts of input data. It helps the model understand how words in a sentence relate to each other. Unlike older models, self-attention looks at the entire sentence at once. This improves context understanding, especially in long texts. It is a core reason why transformers perform so well.

## 17. What is tokenization in Generative AI?

Tokenization is the process of breaking text into smaller units called tokens. These tokens can be words, subwords, or characters. Models do not understand raw text directly and rely on tokens for processing. Tokenization affects how much context a model can handle. Efficient tokenization improves both performance and cost.

## 18. What is a context window?

A context window is the maximum amount of text a model can process at one time. It includes both the input prompt and the generated output. If the context window is exceeded, older information is dropped. Larger context windows allow better long-form understanding. However, they also require more computational resources.

## 19. What is temperature in Generative AI models?

Temperature controls how creative or deterministic a model's output is. A low temperature makes responses more predictable and factual. A high temperature increases randomness and creativity. Choosing the right temperature depends on the use case. Accuracy-focused tasks usually require a lower temperature.

## 20. What is Retrieval-Augmented Generation (RAG)?

Retrieval-Augmented Generation combines a language model with external data sources. Instead of relying only on training data, the model retrieves relevant documents at runtime. This improves accuracy and reduces hallucinations. RAG is especially useful for enterprise and regulated environments. It allows models to generate grounded and up-to-date responses.

## 21. What is few-shot prompting?

Few-shot prompting involves giving the model a few examples before asking it to perform a task. These examples guide the model's behavior and output style. It is useful when zero-shot prompts produce inconsistent results. Few-shot prompting improves accuracy without retraining the model. It is widely used in real-world applications.

## 22. What is zero-shot prompting?

Zero-shot prompting asks the model to perform a task without any examples. The model relies on its general training to understand the request. This approach is fast and flexible. However, results may vary depending on prompt clarity. It works best for simple and common tasks.

## 23. What is chain-of-thought prompting?

Chain-of-thought prompting encourages the model to reason step by step. It helps the model explain how it arrives at an answer. This improves accuracy for complex reasoning tasks. It also makes the output more transparent. However, it should be used carefully in sensitive applications.

## 24. Model fine-tuning compared to prompting

Fine-tuning changes the internal weights of a pre-trained model using new data. Prompting guides behavior without modifying the model. Fine-tuning offers more consistency but requires time and resources. Prompting is faster and more flexible. Most organizations start with prompting before fine-tuning.

## 25. Multimodal Generative AI

Multimodal Generative AI can process and generate multiple types of data. This includes text, images, audio, and video. These models understand relationships across different formats. Multimodality enables richer user experiences. Examples include text-to-image and voice-enabled assistants.

## 26. What is a hallucination risk in business use cases?

Hallucinations can lead to incorrect decisions and compliance issues. In business, inaccurate outputs may cause financial or legal damage. This risk increases when models lack access to verified data. Clear prompts and grounding techniques reduce hallucinations. Human review remains essential.

## 27. Why is data quality important for Generative AI?

Generative AI models learn patterns directly from data. Poor-quality data leads to unreliable outputs. Biased or incomplete data can amplify ethical risks. High-quality data improves relevance and accuracy. Data governance is critical for long-term success.

## 28. What is AI governance?

AI governance refers to policies and controls that guide AI usage. It ensures systems are ethical, secure, and compliant. Governance defines roles, responsibilities, and accountability. It also includes monitoring and audits. Strong governance builds trust in AI systems.

## 29. What is explainability in AI?

Explainability is the ability to understand how an AI system produces outputs. It is especially important in regulated industries. Black-box models can create trust issues. Explainability improves accountability and transparency. It helps stakeholders make informed decisions.

## 30. What is bias in Generative AI?

Bias occurs when AI outputs unfairly favor certain groups. It usually originates from training data. Bias can affect decisions and user trust. Detecting bias requires continuous monitoring. Responsible AI practices aim to reduce bias.

## 31. What is the role of cloud platforms in Generative AI?

Cloud platforms provide scalable infrastructure for AI models. They enable model deployment, monitoring, and security. Cloud services also support APIs and integrations. This makes GenAI accessible without heavy hardware investment. Cloud adoption accelerates innovation.

### 32. What is prompt iteration?

Prompt iteration is the process of refining prompts based on output quality. It involves testing, analyzing, and improving prompts. Small changes can significantly affect results. Iteration helps reduce ambiguity. It is a core skill in prompt engineering.

### 33. What is human-centered AI?

Human-centered AI focuses on collaboration between humans and machines. AI supports decision-making rather than replacing humans. This approach prioritizes usability and trust. Human oversight ensures ethical use. It leads to more reliable AI systems.

### 34. What is autonomous workflow in Agentic AI?

An autonomous workflow allows AI to complete tasks with minimal human input. The system plans steps and executes actions independently. It uses feedback to adjust behavior. Autonomous workflows improve efficiency. They must be governed carefully to avoid risk.

### 35. Why is Generative AI considered transformative?

Generative AI transforms how humans interact with technology. It enables natural language interfaces and creative automation. Unlike traditional automation, it handles complex cognitive tasks. This shifts productivity across industries. Its impact is both technological and organizational.

# Scenario-Based Questions:

## Scenario 1: Hallucination in a Critical Report

**Question:** You deploy a Generative AI model to summarize regulatory documents, but it occasionally produces incorrect legal references. How would you handle this?

**Answer:** I would first identify where hallucinations are occurring and analyze the prompts used. Next, I would ground the model using verified regulatory documents through retrieval-based approaches. I would restrict creativity by lowering temperature and adding strict constraints. Human review would be added for high-risk outputs. Finally, I would monitor outputs continuously and refine prompts based on errors.

## Scenario 2: Choosing Prompting vs Fine-Tuning

**Question:** Your organization wants to use Generative AI for internal knowledge queries. How do you decide between prompt engineering and fine-tuning?

**Answer:** I would start with prompt engineering because it is faster and cost-effective. If outputs are inconsistent, I would introduce structured prompts and examples. Fine-tuning would be considered only if domain language or tone is critical. I would also evaluate retrieval-augmented generation before fine-tuning. The decision would balance accuracy, cost, and governance requirements.

## Scenario 3: AI Adoption Resistance

**Question:** Employees are hesitant to trust AI-generated outputs. How would you encourage adoption?

**Answer:** I would position AI as an assistant rather than a replacement. Training sessions would show practical benefits using real examples. Human-in-the-loop workflows would be introduced to maintain control. Transparency around limitations would build trust. Gradual adoption would help employees gain confidence.

# Scenario-Based Questions:

## Scenario 4: Biased Outputs Detected

**Question:** You notice biased responses from a Generative AI system. What steps would you take?

**Answer:** I would first document and analyze examples of bias. Then, I would review training data sources and prompts. Bias mitigation techniques such as prompt constraints and filtering would be applied. Outputs would be tested across diverse scenarios. Governance checks would be strengthened to prevent recurrence.

## Scenario 5: Deploying GenAI in a Regulated Industry

**Question:** How would you deploy Generative AI in a highly regulated environment like banking?

**Answer:** I would begin with low-risk use cases such as internal assistance. Strong access controls and data privacy measures would be implemented. Outputs would be logged for audit purposes. Human review would be mandatory for external-facing use. Compliance teams would be involved throughout deployment.

## Scenario 6: Long Document Summarization

**Question:** A model struggles to summarize very long documents accurately. How would you solve this?

**Answer:** I would break the document into smaller logical sections. Each section would be summarized independently. A final summary would be created from the section summaries. This approach reduces context limitations. It also improves accuracy and coherence.

## Scenario 7: Evaluating Model Performance

**Question:** How would you evaluate whether a Generative AI solution is performing well?

**Answer:** I would define success metrics based on business goals. Sample outputs would be reviewed for accuracy and relevance. User feedback would be collected regularly. Error trends would be analyzed over time. Continuous improvement would be part of the evaluation process.

# Scenario-Based Questions:

## Scenario 8: Data Privacy Concern

**Question:** A business unit wants to input sensitive data into a public GenAI tool. What would you do?

**Answer:** I would explain the data privacy risks clearly. Policies would restrict sharing confidential data externally. Secure enterprise solutions would be recommended instead. Data anonymization techniques could be applied. Compliance approval would be required before use.

## Scenario 9: Agentic AI Misbehaving

**Question:** An agentic AI system takes unexpected actions. How do you respond?

**Answer:** I would immediately pause the system and analyze logs. The agent's goals and constraints would be reviewed. Guardrails would be strengthened to limit actions. Human approval steps would be added for critical tasks. Testing would be expanded before reactivation.

## Scenario 10: Scaling GenAI Across the Organization

**Question:** How would you scale Generative AI adoption across teams?

**Answer:** I would standardize use cases and best practices. Shared prompt libraries would be created. Training programs would be rolled out organization-wide. Governance frameworks would ensure safe usage. Continuous feedback would guide improvements.

# GENERATIVE AI FOUNDATION CERTIFICATION



## ABOUT GSDC CERTIFICATION



### EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.



### LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.



### CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

## LEARNING OBJECTIVE

- Gain insights into autonomous decision-making processes
- Apply knowledge using ready-to-implement templates
- Demonstrate ability to work with Agentic AI models
- Validate your skills wit

Enroll now with the code **LEARN20** To avail **20%** discount

**Enroll Now**