2025 AI Skills Glossary

accenture cisco cornerstone eightfold.ai Google

IBM indeed intel Microsoft SAP

Table of Contents

Introduction

Developed through cross-industry collaboration by The AI Workforce Consortium, the 2025 AI Skills Glossary establishes a common vocabulary for today's most in-demand AI skills, creating a shared language for workers, educators, and employers. This clarity helps align job requirements with training programs and empowers individuals to build the right skills for 2025.

The dictionary was created through a meticulous, two-phase process designed to ensure relevance and usability:

1. **Demand Analysis:** We began with an extensive analysis of 50 key job roles across G7 countries, encompassing eight distinct ICT job families and specialized support roles. This research identified 480+ AI skills which are detailed in the Job Canvas appendix of our report, the "**ICT in Motion: The Next Wave of AI Integration**".
2. **Skill Consolidation and Taxonomy Alignment:** Next, these 480+ skills were systematically consolidated into 120+ individual skills organized in 10 categories.

This 2025 AI Skills Glossary serves as a foundational blueprint for navigating our collective future. Its purpose is twofold:

- **To Establish a Universal Language:** By mapping the full spectrum of AI competencies—from core technical principles to strategic leadership and ethical governance—the dictionary creates a unified framework. It provides a common vocabulary for educators, industry leaders, and policymakers to align on what proficiency in AI truly means.
- **To Empower the Global Workforce:** It acts as a strategic guide for individuals charting their careers, for organizations building future-ready teams, and for educational institutions designing relevant, high-impact curricula. It directly connects the pursuit of knowledge to the real-world demands of an AI-powered economy.

A Call to Action

This dictionary doesn't pretend to be exhaustive or complete. It is a curated guide with a clear point of view: to map the competencies that deliver tangible value and drive the job market in mid-2025. Therefore, consider this less a neutral catalog and more a strategic playbook for action, defining what it means to be professionally proficient right now. In an era of rapid AI-driven transformation, it empowers individuals and organizations to adapt, upskill, and thrive. While the AI landscape will continuously evolve, this framework provides the essential insights needed to make informed decisions today.

About the AI Workforce Consortium

The AI Workforce Consortium is a group of ten global corporations—Accenture, Cisco, Cornerstone, Eightfold, Google, IBM, Indeed, Intel, Microsoft, and SAP—working alongside global advisors. Together, we have embarked on a collaborative endeavor to share insights and advance an AI-enabled workforce.

Our mission is to prepare today's and tomorrow's workforce with actionable insights and scalable frameworks to leverage the transformational opportunity of AI on ICT jobs across all Industries. In a world where AI is discussed everywhere, this mission has never been more critical

The AI Workforce Consortium is dedicated to equipping workers, policymakers, journalists, executives, researchers, and the public with accurate, rigorously validated, and globally sourced data from G7 economies. This initiative aims to help stakeholders effectively harness the transformative potential of AI in the ICT sector.

We are leveraging the collective insights of our members and advisors to recommend and amplify reskilling and upskilling training programs that are inclusive and can benefit multiple stakeholders – students, career changers, current IT workers, employers, and educators – in order to skill workers at scale. This vision is deeply aligned with the G7's commitment to fostering workforce development amidst the rapid adoption of AI across their economies.

accenture cisco cornerstone eightfold.ai Google

IBM indeed intel Microsoft SAP

How to Use This Glossary

1. Click a category to explore its skills
2. Within a category, select a skill to view detailed information
3. Use the “See Also” links to discover related skills
4. Use the navigation buttons to return to categories anytime
5. Access the 2025 Skills Glossary via our Chatbot

Go to Categories

Get insights with AI

AI Skills Categories 2025

Click a category to explore its skills

1. Fundamental AI Concepts

- AI Literacy
- Conversational AI
- Deep Learning
- Generative AI
- Large Language Models (LLMs)
- Multimodal AI
- Natural Language Processing (NLP)
- Transformer

Back to Skill Categories

AI Literacy

AI Literacy refers to the ability to understand, evaluate, and effectively interact with artificial intelligence systems, including awareness of AI capabilities, limitations, and societal implications.

See also: Generative AI, Large Language Models (LLMs)

Conversational AI

Conversational AI refers to technologies, such as chatbots and voice assistants, that users can interact with using natural language. Built on technologies like NLP and LLMs, these systems simulate human-like dialogue and perform tasks based on conversation.

See also: Large Language Models (LLMs), LLM API Integration, Prompt Engineering

Deep Learning

Deep Learning is a subset of artificial intelligence that utilizes neural networks with multiple layers to automatically learn complex patterns from large datasets. This approach enables tasks such as image recognition, natural language processing, and other sophisticated applications.

See also: Generative AI, Large Language Models (LLMs), Transformer

Generative AI

Generative AI is a type of artificial intelligence that uses generative models to produce new and original content, such as text, images, videos, audio, and code, in response to a prompt.

See also: Deep Learning, Large Language Models (LLMs), Multimodal AI

Large Language Models (LLMs)

Large Language Models (LLMs) are advanced AI systems built on deep learning techniques. Typically containing billions to trillions of parameters, they are trained on vast textual datasets and serve as the foundational technology powering modern generative AI.

See also: Tokens & Embeddings, Prompt Engineering, Chain-of-Thought (CoT) Prompting

Multimodal AI

Multimodal AI refers to systems that can process, understand, and integrate information from multiple types of data (“modalities”), such as text, images, audio, and video.

See also: Vision-Language Models (CLIP, BLIP), Multimodal Embeddings, Cross-Modal Retrieval

Natural Language Processing (NLP)

Natural Language Processing (NLP) is a branch of AI focused on enabling machines to understand, interpret, and generate human language, covering tasks such as translation, summarization, sentiment analysis, and question answering.

See also: Transformer, Tokens & Embeddings

Transformer

The Transformer is a deep learning architecture introduced in 2017 that uses self-attention mechanisms to process sequences of data in parallel, enabling significant advancements in natural language processing, machine translation, and generative AI. It serves as the foundation

See also: Deep Learning, Flash Attention & Optimization, Vision-Language Models (CLIP, BLIP)

2. Large Language Model (LLM) Technologies

Back to Skill Categories

Chain-of-Thought (CoT) Prompting

A prompting technique encouraging an LLM to explain its reasoning step-by-step before giving a final answer, improving performance on complex tasks.

See also: Tree of Thoughts (ToT), Self-Consistency, Prompt Engineering

Context Engineering

Systematically designing and managing all relevant information provided to an LLM at inference time, including instructions, inputs, retrieved documents, and tool outputs.

See also: Context Window Management, Retrieval-Augmented Generation (RAG), LangChain and LlamaIndex

Context Window Management

Techniques to handle the limited input capacity of LLMs, such as summarization and sliding windows, ensuring the most relevant information is available.

See also: Tokens & Embeddings, Prompt Engineering, Retrieval-Augmented Generation (RAG)

Foundation Model Adaptation

Foundation Model Adaptation refers to methods for customizing large pre-trained AI models for specific tasks or domains through fine-tuning, prompt engineering, or lightweight adaptation techniques.

See also: Parameter-Efficient Fine-Tuning (PEFT), Retrieval-Augmented Generation (RAG)

LLM API Integration

Integrating large language models into custom applications and workflows via APIs to leverage advanced reasoning and generation capabilities.

See also: Prompt Engineering, Tool Use & Function Calling, LangChain and LlamaIndex

LangChain and LlamaIndex

Software frameworks that simplify development of LLM-powered applications with modular components for workflows, data retrieval, and external data sources.

See also: Retrieval-Augmented Generation (RAG), Tool Use & Function Calling, Vector Databases (Pinecone, Weaviate)

Prompt Engineering

Prompt Engineering is the practice of crafting effective inputs for AI models, particularly large language models, to produce desired outputs, optimize performance, and reduce ambiguity or bias.

See also: Chain-of-Thought (CoT) Prompting, Context Engineering

Retrieval-Augmented Generation (RAG)

RAG enhances LLM responses by retrieving relevant, up-to-date information from external knowledge bases, increasing answer accuracy and allowing source citation.

See also: Vector Databases (Pinecone, Weaviate), LangChain and LlamaIndex, Context Engineering

Self-Consistency

Improves accuracy by sampling diverse reasoning paths and choosing the most consistent answer, assuming correct answers are reached by different logical steps.

See also: Chain-of-Thought (CoT) Prompting, Tree of Thoughts (ToT), Prompt Engineering

Test-Time Compute & Reasoning Models

Test-Time Compute & Reasoning Models refer to AI architectures that allocate additional computational resources during inference to enhance reasoning, problem-solving, and accuracy without retraining.

See also: Chain-of-Thought (CoT) Prompting, Tree of Thoughts (ToT)

Tokens & Embeddings

Tokens are fundamental units of data that an LLM processes, while embeddings are high-dimensional vectors representing the meaning and context of these tokens, enabling the model to perform mathematical operations on language.

See also: Context Window Management, Vector Databases (Pinecone, Weaviate), LangChain and LlamaIndex

Tool Use & Function Calling

Empowers LLMs to interact with external tools, APIs, and functions for real-time information or actions, extending their capabilities beyond text generation.

See also: LLM API Integration, LangChain and LlamaIndex, Autonomous AI Agents

Tree of Thoughts (ToT)

A framework enabling LLMs to explore multiple reasoning paths as branches of a tree, self-evaluate, and backtrack to improve answers.

See also: Chain-of-Thought (CoT) Prompting, Self-Consistency, Tool Use & Function Calling

Vector Databases (Pinecone, Weaviate)

Specialized systems for storing, managing, and searching high-dimensional vector embeddings, core infrastructure for semantic search and RAG systems.

See also: Tokens & Embeddings, Retrieval-Augmented Generation (RAG), LangChain and LlamaIndex

3. Generative & Multimodal Systems

- 3D Generation (Point-E, Shap-E)
- Any-to-Any Generation
- ControlNet & Image Editing
- Cross-Modal Retrieval
- Digital Human Creation
- Multimodal Embeddings
- Music Generation (MusicGen, AudioCraft)
- Neural Radiance Fields (NeRFs)
- Stable Diffusion/DALL-E 3/Midjourney
- Video Generation (Runway, Pika, Veo, Sora)
- Vision-Language Models (CLIP, BLIP)

Back to Skill Categories

3D Generation (Point-E, Shap-E)

Creating 3D objects and scenes from text prompts or 2D images, accelerating development in gaming and virtual reality.

See also: Neural Radiance Fields (NeRFs), Stable Diffusion/DALL-E 3/Midjourney, Video Generation (Runway, Pika, Veo, Sora)

Any-to-Any Generation

Generative AI models that take any data modalities as input and generate output in any other modality.

See also: Music Generation (MusicGen, AudioCraft), Video Generation (Runway, Pika, Veo, Sora), Digital Human Creation

ControlNet & Image Editing

A neural network architecture providing granular control over image generation models, enabling precise editing using poses, edges, or depth maps.

See also: Stable Diffusion/DALL-E 3/Midjourney, Video Generation (Runway, Pika, Veo, Sora), 3D Generation (Point-E, Shap-E)

Cross-Modal Retrieval

Enables searching for one type of data using a different type as the query, e.g., finding images based on text descriptions, powered by multimodal embeddings.

See also: Multimodal Embeddings, Vision-Language Models (CLIP, BLIP), Vector Databases (Pinecone, Weaviate)

Digital Human Creation

Utilizes AI to generate and animate highly realistic virtual humans for entertainment and customer service, combining advanced graphics with natural language.

See also: Any-to-Any Generation, Stable Diffusion/DALL-E 3/Midjourney, 3D Generation (Point-E, Shap-E)

Multimodal Embeddings

Representing different data types (text, images, audio) within a shared vector space, enabling cross-modality relationships and searches.

See also: Vision-Language Models (CLIP, BLIP), Cross-Modal Retrieval, Tokens & Embeddings

Music Generation (MusicGen, AudioCraft)

Models designed to compose and produce original, high-fidelity music from text descriptions or melodic inputs.

See also: Any-to-Any Generation, Multimodal Embeddings, Generative AI

Neural Radiance Fields (NeRFs)

Synthesizes novel 3D views of a scene from partial 2D images, creating photorealistic, navigable 3D scenes from photographs.

See also: 3D Generation (Point-E, Shap-E), Video Generation (Runway, Pika, Veo, Sora), Multimodal Embeddings

Stable Diffusion/DALL-E 3/ Midjourney

State-of-the-art text-to-image models that generate complex, high-quality visuals from natural language, revolutionizing art and design workflows.

See also: ControlNet & Image Editing, Video Generation (Runway, Pika, Veo, Sora), Any-to-Any Generation

Video Generation (Runway, Pika, Veo, Sora)

Platforms that use AI to create and edit video content from text prompts, images, or other videos, enabling new forms of digital storytelling.

See also: Stable Diffusion/DALL-E 3/Midjourney, Neural Radiance Fields (NeRFs), Music Generation (MusicGen, AudioCraft)

Vision-Language Models (CLIP, BLIP)

Models trained to understand relationships between images and text, supporting tasks like image classification, captioning, and visual question answering.

See also: Multimodal Embeddings, Cross-Modal Retrieval, Stable Diffusion/DALL-E 3/Midjourney

4. Advanced AI Architectures & Techniques

Back to Skill Categories

Agentic AI, Agents

AI systems capable of autonomous decision-making and action. These systems, often called AI agents, proactively pursue goals with minimal human intervention, adapt to changing environments, and may use tools or interact with other systems to achieve objectives.

See also: Tool Use & Function Calling, Multi-Agent Systems

Constitutional AI & RLHF

Guides models with principles (“constitution”) and uses human feedback (RLHF) to align model behavior with desired outcomes.

See also: Direct Preference Optimization (DPO), AI Ethics, AI Alignment

Direct Preference Optimization (DPO)

A stable, efficient alternative to RLHF for aligning models to human preferences by directly optimizing on preferred responses.

See also: Constitutional AI & RLHF, AI Risk Management, Model Evaluations & Benchmarking (Evals)

Federated Learning

A machine learning technique training algorithms across decentralized devices without exchanging data, preserving privacy.

See also: Privacy-Preserving AI, Knowledge Distillation, Parameter-Efficient Fine-Tuning (PEFT)

Flash Attention & Optimization

An optimized algorithm for Transformer attention, speeding up training and inference while reducing memory usage.

See also: Transformer, State Space Models (Mamba), Quantization (GPTQ, AWQ, GGUF)

Knowledge Distillation

A training technique where a smaller “student” model learns to replicate a larger “teacher” model, transferring knowledge efficiently.

See also: Quantization (GPTQ, AWQ, GGUF), Model Merging & Ensemble, Parameter-Efficient Fine-Tuning (PEFT)

Knowledge Graphs

Represents data as interconnected entities and relationships, enhancing factual context and reasoning for AI systems.

See also: Retrieval-Augmented Generation (RAG), Federated Learning, Data Engineering

Mixture of Experts (MoE)

Architecture that activates only specific model subsets ("experts") per input, enabling larger, more efficient, and cost-effective models.

See also: State Space Models (Mamba), Parameter-Efficient Fine-Tuning (PEFT), Model Merging & Ensemble

Model Merging & Ensemble

Combines weights of multiple trained models to create a more capable model without additional training, improving or blending skills.

See also: Mixture of Experts (MoE), Quantization (GPTQ, AWQ, GGUF), Knowledge Distillation

Parameter-Efficient Fine-Tuning (PEFT)

Adapts pre-trained models to specific tasks by updating a small parameter subset, enabling powerful customization at low cost.

See also: Quantization (GPTQ, AWQ, GGUF), Knowledge Distillation, Mixture of Experts (MoE)

Quantization (GPTQ, AWQ, GGUF)

Reduces the precision of model weights to make models smaller and faster, enabling large models to run on consumer hardware.

See also: Model Merging & Ensemble, Knowledge Distillation, Parameter-Efficient Fine-Tuning (PEFT)

State Space Models (Mamba)

An emerging architecture processing sequences with high efficiency, offering an alternative to Transformers for long-context tasks.

See also: Mixture of Experts (MoE), Flash Attention & Optimization, Parameter-Efficient Fine-Tuning (PEFT)

5. AI Infrastructure & Operations

- AI Development Platforms
- AI Incident Management
- AI Observability
- AI for Network Management & Optimization
- AIOps Integration
- Chaos Engineering with AI
- ChatOps with AI
- Cloud Native AI
- Cost Optimization Strategies
- Data Engineering
- GPU Optimization & Management
- Inference Servers (vLLM, TGI, TensorRT-LLM)
- LLMOps & Monitoring
- Local & Edge AI Deployment
- Model Evaluations & Benchmarking (Evals)
- Serverless AI (Modal, Replicate)
- Streaming & Real-time Inference

Back to Skill Categories

AI Development Platforms

Integrated environments (e.g., Azure, Vertex AI, SageMaker) providing tools for building, deploying, and managing AI models.

See also: Cloud Native AI, Serverless AI (Modal, Replicate), Inference Servers (vLLM, TGI, TensorRT-LLM)

AI Incident Management

Responding to and resolving unexpected failures in live AI systems, detecting performance issues, and restoring reliability.

See also: Chaos Engineering with AI, LLMOps & Monitoring, AIOps Integration

AI Observability

Advanced system monitoring providing deep insights into AI system behavior, analyzing telemetry to debug and diagnose issues.

See also: LLMOps & Monitoring, Model Evaluations & Benchmarking (Evals), AI Incident Management

AI for Network Management & Optimization

Applying AI to monitor, analyze, and manage computer networks, including anomaly detection and bandwidth forecasting.

See also: ChatOps with AI, AIOps Integration, AI Observability

AIOps Integration

Using AI to automate and enhance IT operations, identifying and resolving infrastructure issues for improved reliability.

See also: LLMOps & Monitoring, AI Incident Management, ChatOps with AI

Chaos Engineering with AI

Proactively injecting failures into AI systems to test resilience and discover weaknesses for improved robustness.

See also: AI Incident Management, ChatOps with AI, AIOps Integration

ChatOps with AI

Integrating AI assistants into chat platforms to streamline DevOps workflows using natural language commands.

See also: AIOps Integration, AI Incident Management, LLMOps & Monitoring

Cloud Native AI

Architecting AI applications with cloud technologies like containers and microservices for scalability, resilience, and automation.

See also: Serverless AI (Modal, Replicate), AI Development Platforms, Local & Edge AI Deployment

Cost Optimization Strategies

Techniques to reduce AI model training and deployment costs, such as spot instances, right-sizing, quantization, and efficient servers.

See also: GPU Optimization & Management, Serverless AI (Modal, Replicate), AI Incident Management

Data Engineering

Designing, building, and maintaining systems and pipelines to collect, store, and transform data for AI training and usage.

See also: Vector Databases (Pinecone, Weaviate), AI Observability, AI Development Platforms

GPU Optimization & Management

Maximizing GPU efficiency for AI workloads via kernel fusion, precision reduction, and scheduling strategies.

See also: Cost Optimization Strategies, Inference Servers (vLLM, TGI, TensorRT-LLM), Local & Edge AI Deployment

Inference Servers (vLLM, TGI, TensorRT-LLM)

High-throughput engines designed to optimize LLM inference using techniques like PagedAttention for speed and efficiency.

See also: Streaming & Real-time Inference, GPU Optimization & Management, Cloud Native AI

LLMOps & Monitoring

Practices for managing the end-to-end lifecycle of LLMs in production, including deployment, monitoring, versioning, and CI/CD.

See also: Model Evaluations & Benchmarking (Evals), AI Observability, AIOps Integration

Local & Edge AI Deployment

Running AI models on local machines or edge devices, benefiting privacy, latency, and bandwidth usage.

See also: Streaming & Real-time Inference, GPU Optimization & Management, Data Engineering

Model Evaluations & Benchmarking (Evals)

Frameworks and processes for assessing AI model quality against benchmarks and datasets across metrics like accuracy, speed, and safety.

See also: LLMOps & Monitoring, AI Observability, Model Risk Management

Serverless AI (Modal, Replicate)

Platforms enabling deployment and scaling of AI models without server management, handling auto-scaling and resource allocation.

See also: AI Development Platforms, Cost Optimization Strategies, Cloud Native AI

Streaming & Real-time Inference

Generating AI outputs token-by-token in real time, essential for responsive applications such as chatbots.

See also: Inference Servers (vLLM, TGI, TensorRT-LLM), AIOps Integration, Local & Edge AI Deployment

6. AI Safety, Ethics & Governance

- AI Alignment
- AI Auditing
- AI Ethics
- AI Governance
- AI Impact Assessments
- AI Red Teaming
- AI Risk Management
- AI Security
- AI Security Architecture
- AI Threat Modeling
- AI for Accessibility
- AI for Cybersecurity
- Bias Auditing & Correction
- Hallucination Mitigation
- Human-in-the-Loop (HITL) Systems
- Jailbreak Prevention
- LLM Security & Jailbreak Defense
- Model Interpretability & Explainable AI (XAI)
- Model Risk Management
- Privacy-Preserving AI
- Responsible AI

Back to Skill Categories

AI Alignment

Ensuring advanced AI systems pursue goals and behaviors consistent with human values, preventing unintended harm.

See also: Constitutional AI & RLHF, AI Risk Management, AI Ethics

AI Auditing

Systematic evaluation of AI systems to ensure fairness, transparency, and compliance with legal and ethical standards.

See also: AI Risk Management, AI Security Architecture, Model Risk Management

AI Ethics

A framework of principles guiding responsible AI design, development, and deployment.

See also: AI Alignment, AI Impact Assessments, Bias Auditing & Correction

AI Governance

AI Governance encompasses the frameworks, policies, and oversight mechanisms to guide the responsible development, deployment, and monitoring of AI systems within organizations and society.

See also: AI Risk Management, AI Auditing

AI Impact Assessments

Formal evaluations to understand and document the societal, ethical, and economic consequences of deploying an AI system.

See also: AI Risk Management, AI Ethics, Bias Auditing & Correction

AI Red Teaming

Adversarially testing AI systems to proactively discover vulnerabilities and prevent harmful behavior.

See also: AI Threat Modeling, Jailbreak Prevention, AI Security Architecture

AI Risk Management

Identifying, assessing, and mitigating technical, ethical, and business risks associated with AI systems.

See also: Model Risk Management, AI Auditing, AI Security Architecture

AI Security

AI Security refers to the protection of AI systems against threats, including adversarial attacks, data poisoning, model theft, and misuse, ensuring reliability and trustworthiness.

See also: AI Threat Modeling, Privacy-Preserving AI

AI Security Architecture

Designing and building end-to-end security controls into AI systems, protecting data pipelines and models from tampering and theft.

See also: AI Threat Modeling, AI Risk Management, Model Risk Management

AI Threat Modeling

Structured security practice to identify and mitigate threats unique to AI systems during design, e.g., data poisoning or model theft.

See also: AI Red Teaming, AI Security Architecture, Jailbreak Prevention

AI for Accessibility

Designing AI systems usable by people with disabilities, ensuring AI benefits everyone.

See also: Human-in-the-Loop (HITL) Systems, AI Impact Assessments, AI Ethics

AI for Cybersecurity

Applying AI for cybersecurity defense, including threat detection, endpoint detection, and zero-trust architectures.

See also: AI Threat Modeling, AI Red Teaming, AI Security Architecture

Bias Auditing & Correction

Systematic examination and correction of AI models and data to remove unfair biases related to demographics.

See also: AI Ethics, Hallucination Mitigation, AI Impact Assessments

Hallucination Mitigation

Techniques to reduce or prevent AI models from generating incorrect or nonsensical information, often by grounding with external data.

See also: Bias Auditing & Correction, Model Interpretability & Explainable AI (XAI), Retrieval-Augmented Generation (RAG)

Human-in-the-Loop (HITL) Systems

Designing systems where human oversight validates AI outputs, handles exceptions, and provides feedback for accuracy and safety.

See also: Model Interpretability & Explainable AI (XAI), AI for Accessibility, AI Risk Management

Jailbreak Prevention

Developing safeguards to prevent users from bypassing AI model safety features with malicious prompts.

See also: AI Red Teaming, AI Threat Modeling, AI Security Architecture

LLM Security & Jailbreak Defense

LLM Security & Jailbreak Defense involves techniques and safeguards to protect large language models from malicious prompts, jailbreak attempts, and adversarial manipulation, ensuring safe and ethical responses.

See also: AI Red Teaming, Constitutional AI & RLHF

Model Interpretability & Explainable AI (XAI)

Methods to understand and explain internal decision-making in AI models, using techniques like SHAP and LIME for transparency.

See also: Hallucination Mitigation, Model Risk Management, Bias Auditing & Correction

Model Risk Management

Formal framework for identifying, measuring, and mitigating risks throughout a model's lifecycle, essential for regulated industries.

See also: AI Risk Management, AI Auditing, Hallucination Mitigation

Privacy-Preserving AI

Techniques like federated learning and differential privacy to train and use AI models without exposing sensitive data.

See also: Federated Learning, AI Risk Management, Model Interpretability & Explainable AI (XAI)

Responsible AI

Responsible AI emphasizes the design, deployment, and use of AI systems in ways that are ethical, fair, transparent, and aligned with human values, minimizing bias and unintended harm.

See also: AI Ethics, Model Interpretability & Explainable AI (XAI)

7. AI-Enhanced Professional Skills

Back to Skill Categories

AI Business Value Metrics & ROI Modeling

Defining and measuring the financial impact and outcomes of AI initiatives, forecasting and proving ROI.

See also: AI Strategy Development, AI Maturity Assessment, AI Financial Modeling

AI Financial Modeling

Applying AI to automate and enhance financial models, analyzing large datasets, identifying patterns, and improving financial predictions.

See also: AI Business Value Metrics & ROI Modeling, AI-Powered Analytics, Asset Management AI

AI Maturity Assessment

Systematic evaluation of an organization's readiness to adopt and scale AI, assessing strategy, data, talent, and governance.

See also: AI Strategy Development, AI Project Management, AI Business Value Metrics & ROI Modeling

AI Product Strategy

Integrating product management with AI understanding to guide AI-powered product development, including prioritization and monetization.

See also: AI Strategy Development, AI Project Management, AI Business Value Metrics & ROI Modeling

AI Project Management

Managing AI initiatives from conception to deployment, combining traditional project management with AI lifecycle understanding.

See also: AI Product Strategy, AI Maturity Assessment, AI Business Value Metrics & ROI Modeling

AI Recruitment

Developing and managing AI solutions to streamline recruitment and hiring, combining AI knowledge with HR workflows.

See also: AI Strategy Development, AI for Marketing & Sales, AI Project Management

AI Strategy

Developing and implementing organizational plans to leverage AI for competitive advantage, efficiency, and innovation.

See also: AI Business Value Metrics & ROI Modeling, AI Maturity Assessment

AI Strategy Development

Creating and implementing plans for AI adoption, including operating models, ROI analysis, team building, vendor management, and culture.

See also: Enterprise Data Strategy for AI, AI Product Strategy, AI Maturity Assessment

AI for Customer Experience (CX)

Using AI to understand and enhance the customer lifecycle, modeling journeys, and powering customer intelligence tools.

See also: AI for Marketing & Sales, Conversational AI, AI-Powered Analytics

AI for Marketing & Sales

Applying AI to optimize marketing and sales processes, including attribution modeling and campaign optimization.

See also: AI for Customer Experience (CX), AI-Powered Analytics, AI Business Value Metrics & ROI Modeling

AI-Powered Analytics

Applying AI techniques to automate and accelerate data analysis, deriving insights, predictions, and recommendations.

See also: Conversational & Natural Language BI, AI Business Value Metrics & ROI Modeling, AI Financial Modeling

Asset Management AI

Using machine learning to track, manage, and optimize assets, including predictive maintenance and inventory optimization.

See also: AI Financial Modeling, AI for Marketing & Sales, AI Business Value Metrics & ROI Modeling

Conversational & Natural Language BI

Using AI to enable natural language queries for business intelligence insights, allowing users to ask questions in plain language.

See also: AI-Powered Analytics, AI Financial Modeling, Large Language Models (LLMs)

Enterprise Data Strategy for AI

Preparing and leveraging organizational data for AI, involving data governance, breaking silos, ensuring quality, and managing proprietary datasets.

See also: AI Strategy Development, Data Engineering, AI Product Strategy

8. Emerging AI Applications

- AI Pair Programming
- AI-Driven Forecasting
- AI-Generated GUI (Generative UI)
- AI-Powered Document Analysis & Generation
- AI-Powered OSINT (Open-Source Intelligence)
- AI-Powered Prototyping
- AI-Powered RPA (Robotic Process Automation)
- AI-Powered Robotics & Embodied AI
- AI-Powered Software Testing
- Anomaly Detection
- Autonomous AI Agents
- Coding Agents (e.g., Devin, GitHub co-pilot workspace, CodiumAI, OpenDevin)
- Design Token AI
- Information Research Agents (Deep Research)
- Multi-Agent Systems
- Natural Language Automation (NLA)

Back to Skill Categories

AI Pair Programming

Tools like GitHub Copilot that function as AI assistants inside code editors, providing real-time code suggestions and problem-solving.

See also: Coding Agents (e.g., Devin, GitHub co-pilot workspace, CodiumAI, OpenDevin), AI-Powered Software Testing, AI-Powered Prototyping

AI-Driven Forecasting

Applies machine learning to historical and real-time data for accurate predictions in demand, finance, and resource planning.

See also: AI-Powered Analytics, Anomaly Detection, AI-Powered OSINT (Open-Source Intelligence)

AI-Generated GUI (Generative UI)

Uses generative AI to dynamically create and adapt graphical user interfaces, generating code or adaptive layouts in real time.

See also: AI-Powered Prototyping, AI-Powered RPA (Robotic Process Automation), AI-Powered Document Analysis & Generation

AI-Powered Document Analysis & Generation

Automates processing of unstructured documents, including review, extraction, contract generation, and technical documentation.

See also: Natural Language Automation (NLA), AI-Powered RPA (Robotic Process Automation), AI-Generated GUI (Generative UI)

AI-Powered OSINT (Open-Source Intelligence)

Leveraging AI to gather, process, and analyze information from public sources for trends and insights.

See also: Information Research Agents (Deep Research), AI-Powered Document Analysis & Generation, AI-Powered Robotics & Embodied AI

AI-Powered Prototyping

Uses generative AI tools to rapidly create mockups and functional prototypes of applications and websites from text or sketches.

See also: AI Pair Programming, AI-Generated GUI (Generative UI), Natural Language Automation (NLA)

AI-Powered RPA (Robotic Process Automation)

Enhances traditional RPA by integrating machine learning, automating workflows involving unstructured data and decisions.

See also: AI-Powered Document Analysis & Generation, Natural Language Automation (NLA), AI Pair Programming

AI-Powered Robotics & Embodied AI

Integrates advanced AI models into robots to perform complex tasks, with agents learning through physical interaction.

See also: Autonomous AI Agents, Anomaly Detection, Design Token AI

AI-Powered Software Testing

Automating the software testing process with AI, generating test cases, identifying bugs, and predicting at-risk application areas.

See also: AI Pair Programming, Coding Agents (e.g., Devin, GitHub co-pilot workspace, CodiumAI, OpenDevin), AI-Powered Prototyping

Anomaly Detection

Identifying rare items or events that deviate from most data, widely used for fraud detection, cybersecurity, and maintenance.

See also: AI-Driven Forecasting, AI-Powered Software Testing, AI for Cybersecurity

Autonomous AI Agents

Sophisticated systems designed to perceive environments, break down goals, use tools, and take actions with minimal human oversight.

See also: Multi-Agent Systems, Coding Agents (e.g., Devin, GitHub co-pilot workspace, CodiumAI, OpenDevin), AI-Powered Robotics & Embodied AI

Coding Agents (e.g., Devin, GitHub co-pilot workspace, CodiumAI, OpenDe

Autonomous AI systems automating software development, handling planning, coding, debugging, and testing under human supervision.

See also: AI Pair Programming, AI-Powered Software Testing, AI-Powered Prototyping

Design Token AI

Uses AI to automate creation and maintenance of design tokens for brand consistency in design systems.

See also: AI-Generated GUI (Generative UI), AI-Powered Prototyping, Anomaly Detection

Information Research Agents (Deep Research)

AI systems that autonomously research topics, synthesizing information from multiple sources and delivering structured reports with citations.

See also: Autonomous AI Agents, Multi-Agent Systems, AI-Powered Document Analysis & Generation

Multi-Agent Systems

Applications where multiple AI agents collaborate, negotiate, or compete to solve complex problems.

See also: Autonomous AI Agents, Coding Agents (e.g., Devin, GitHub co-pilot workspace, CodiumAI, OpenDevin), Information Research Agents (Deep Research)

Natural Language Automation (NLA)

Uses AI to understand and act on instructions in everyday language, automating complex workflows across applications.

See also: AI-Powered Document Analysis & Generation, AI-Powered RPA (Robotic Process Automation), AI-Powered Analytics

9. Key AI Platforms & Models

- Anthropic Claude 4 Series
- Cursor & Windsurf IDEs
- Google Gemini 2 Series (Flash, Pro, Ultra)
- Meta Llama 4 Series
- Mistral Large and Codestral
- OpenAI GPT-5 & “o” Family Models
- Perplexity Pro Search
- Together AI Inference
- xAI Grok 4

Back to Skill Categories

Anthropic Claude 4 Series

Released in Q2 2025, the Claude 4 series focuses on enterprise use cases, excelling in long-context tasks and sophisticated instruction following.

See also: OpenAI GPT-5 & “o” Family Models, Google Gemini 2 Series (Flash, Pro, Ultra), Meta Llama 4 Series

Cursor & Windsurf IDEs

AI-powered IDEs serving as collaborative research assistants, with contextual code generation and intelligent debugging.

See also: Coding Agents (e.g., Devin, GitHub co-pilot workspace, CodiumAI, OpenDevin), AI Pair Programming, Together AI Inference

Google Gemini 2 Series (Flash, Pro, Ultra)

Native multimodal models, with Flash for high-volume tasks, Pro for general purpose, and Ultra as the flagship.

See also: Anthropic Claude 4 Series, OpenAI GPT-5 & “o” Family Models, Meta Llama 4 Series

Meta Llama 4 Series

Released mid-2025, Llama 4 sets the standard for open-source models with a new, efficient architecture.

See also: Mistral Large and Codestral, OpenAI GPT-5 & “o” Family Models, Google Gemini 2 Series (Flash, Pro, Ultra)

Mistral Large and Codestral

Mistral Large is for complex reasoning; Codestral is a state-of-the-art open-weight model for code generation in 80+ languages.

See also: Meta Llama 4 Series, OpenAI GPT-5 & “o” Family Models, xAI Grok 4

OpenAI GPT-5 & “o” Family Models

The next generation from OpenAI, including GPT-5 and multimodal o1 (Omni), advancing reasoning, autonomy, and efficiency.

See also: Anthropic Claude 4 Series, Google Gemini 2 Series (Flash, Pro, Ultra), Meta Llama 4 Series

Perplexity Pro Search

Conversational search service leveraging powerful models to deliver in-depth, reliably cited answers.

See also: Cursor & Windsurf IDEs, LLM API Integration, OpenAI GPT-5 & “o” Family Models

Together AI Inference

Cloud platform for fast, scalable inference of open-source generative AI models.

See also: Inference Servers (vLLM, TGI, TensorRT-LLM), LLMOps & Monitoring, Cloud Native AI

xAI Grok 4

Released mid-2025, xAI’s multimodal LLM is built for superior reasoning, real-time access, and advanced STEM performance.

See also: Mistral Large and Codestral, Anthropic Claude 4 Series, Google Gemini 2 Series (Flash, Pro, Ultra)

10. Future Frontiers

- Quantum Machine Learning (QML)

Back to Skill Categories

Quantum Machine Learning (QML)

QML leverages quantum principles like superposition and entanglement to potentially accelerate complex ML tasks far beyond classical computers.

See also: Deep Learning, Federated Learning, Multimodal AI

Appendix – Summary of AI Skills 2025

1. Fundamental AI Concepts

- AI Literacy
- Conversational AI
- Deep Learning
- Generative AI
- Large Language Models (LLMs)
- Multimodal AI
- Natural Language Processing (NLP)
- Transformer

2. Large Language Model (LLM) Technologies

- Chain-of-Thought (CoT) Prompting
- Context Engineering
- Context Window Management
- Foundation Model Adaptation
- LLM API Integration
- LangChain and LlamaIndex
- Prompt Engineering
- Retrieval-Augmented Generation (RAG)
- Self-Consistency
- Test-Time Compute & Reasoning Models
- Tokens & Embeddings
- Tool Use & Function Calling
- Tree of Thoughts (ToT)
- Vector Databases (Pinecone, Weaviate)

3. Generative & Multimodal Systems

- 3D Generation (Point-E, Shap-E)
- Any-to-Any Generation
- ControlNet & Image Editing
- Cross-Modal Retrieval
- Digital Human Creation
- Multimodal Embeddings
- Music Generation (MusicGen, AudioCraft)
- Neural Radiance Fields (NeRFs)
- Stable Diffusion/DALL-E 3/Midjourney
- Video Generation (Runway, Pika, Veo, Sora)
- Vision-Language Models (CLIP, BLIP)

4. Advanced AI Architectures & Techniques

- Agentic AI, Agents
- Constitutional AI & RLHF
- Direct Preference Optimization (DPO)
- Federated Learning
- Flash Attention & Optimization
- Knowledge Distillation
- Knowledge Graphs
- Mixture of Experts (MoE)
- Model Merging & Ensemble
- Parameter-Efficient Fine-Tuning (PEFT)
- Quantization (GPTQ, AWQ, GGUF)
- State Space Models (Mamba)

5. AI Infrastructure & Operations (LLMOps)

- AI Development Platforms
- AI Incident Management
- AI Observability
- AI for Network Management & Optimization
- AIOps Integration
- Chaos Engineering with AI
- ChatOps with AI
- Cloud Native AI
- Cost Optimization Strategies
- Data Engineering
- GPU Optimization & Management
- Inference Servers (vLLM, TGI, TensorRT-LLM)
- LLMOps & Monitoring
- Local & Edge AI Deployment
- Model Evaluations & Benchmarking (Evals)
- Serverless AI (Modal, Replicate)
- Streaming & Real-time Inference

6. AI Safety, Ethics & Governance

- AI Alignment
- AI Auditing
- AI Ethics
- AI Governance
- AI Impact Assessments
- AI Red Teaming
- AI Risk Management
- AI Security
- AI Security Architecture
- AI Threat Modeling
- AI for Accessibility
- AI for Cybersecurity
- Bias Auditing & Correction
- Hallucination Mitigation
- Human-in-the-Loop (HITL) Systems
- Jailbreak Prevention
- LLM Security & Jailbreak Defense
- Model Interpretability & Explainable AI (XAI)
- Model Risk Management
- Privacy-Preserving AI
- Responsible AI

7. AI-Enhanced Professional Skills

- AI Business Value Metrics & ROI Modeling
- AI Financial Modeling
- AI Maturity Assessment
- AI Product Strategy
- AI Project Management
- AI Recruitment
- AI Strategy
- AI Strategy Development
- AI for Customer Experience (CX)
- AI for Marketing & Sales
- AI-Powered Analytics
- Asset Management AI
- Conversational & Natural Language BI
- Enterprise Data Strategy for AI

8. Emerging AI Applications

- AI Pair Programming
- AI-Driven Forecasting
- AI-Generated GUI (Generative UI)
- AI-Powered Document Analysis & Generation
- AI-Powered OSINT (Open-Source Intelligence)
- AI-Powered Prototyping
- AI-Powered RPA (Robotic Process Automation)
- AI-Powered Robotics & Embodied AI
- AI-Powered Software Testing
- Anomaly Detection
- Autonomous AI Agents
- Coding Agents (e.g., Devin, GitHub co-pilot workspace, CodiumAI, OpenDevin)
- Design Token AI
- Information Research Agents (Deep Research)
- Multi-Agent Systems
- Natural Language Automation (NLA)

9. Key AI Platforms & Models (as of mid-2025)

- Anthropic Claude 4 Series
- Cursor & Windsurf IDEs
- Google Gemini 2 Series (Flash, Pro, Ultra)
- Meta Llama 4 Series
- Mistral Large and Codestral
- OpenAI GPT-5 & “o” Family Models
- Perplexity Pro Search
- Together AI Inference
- xAI Grok 4

10. Future Frontiers

- Quantum Machine Learning (QML)

AI Workforce Consortium

Our vision is to enable the success of the ICT workforce in the AI era

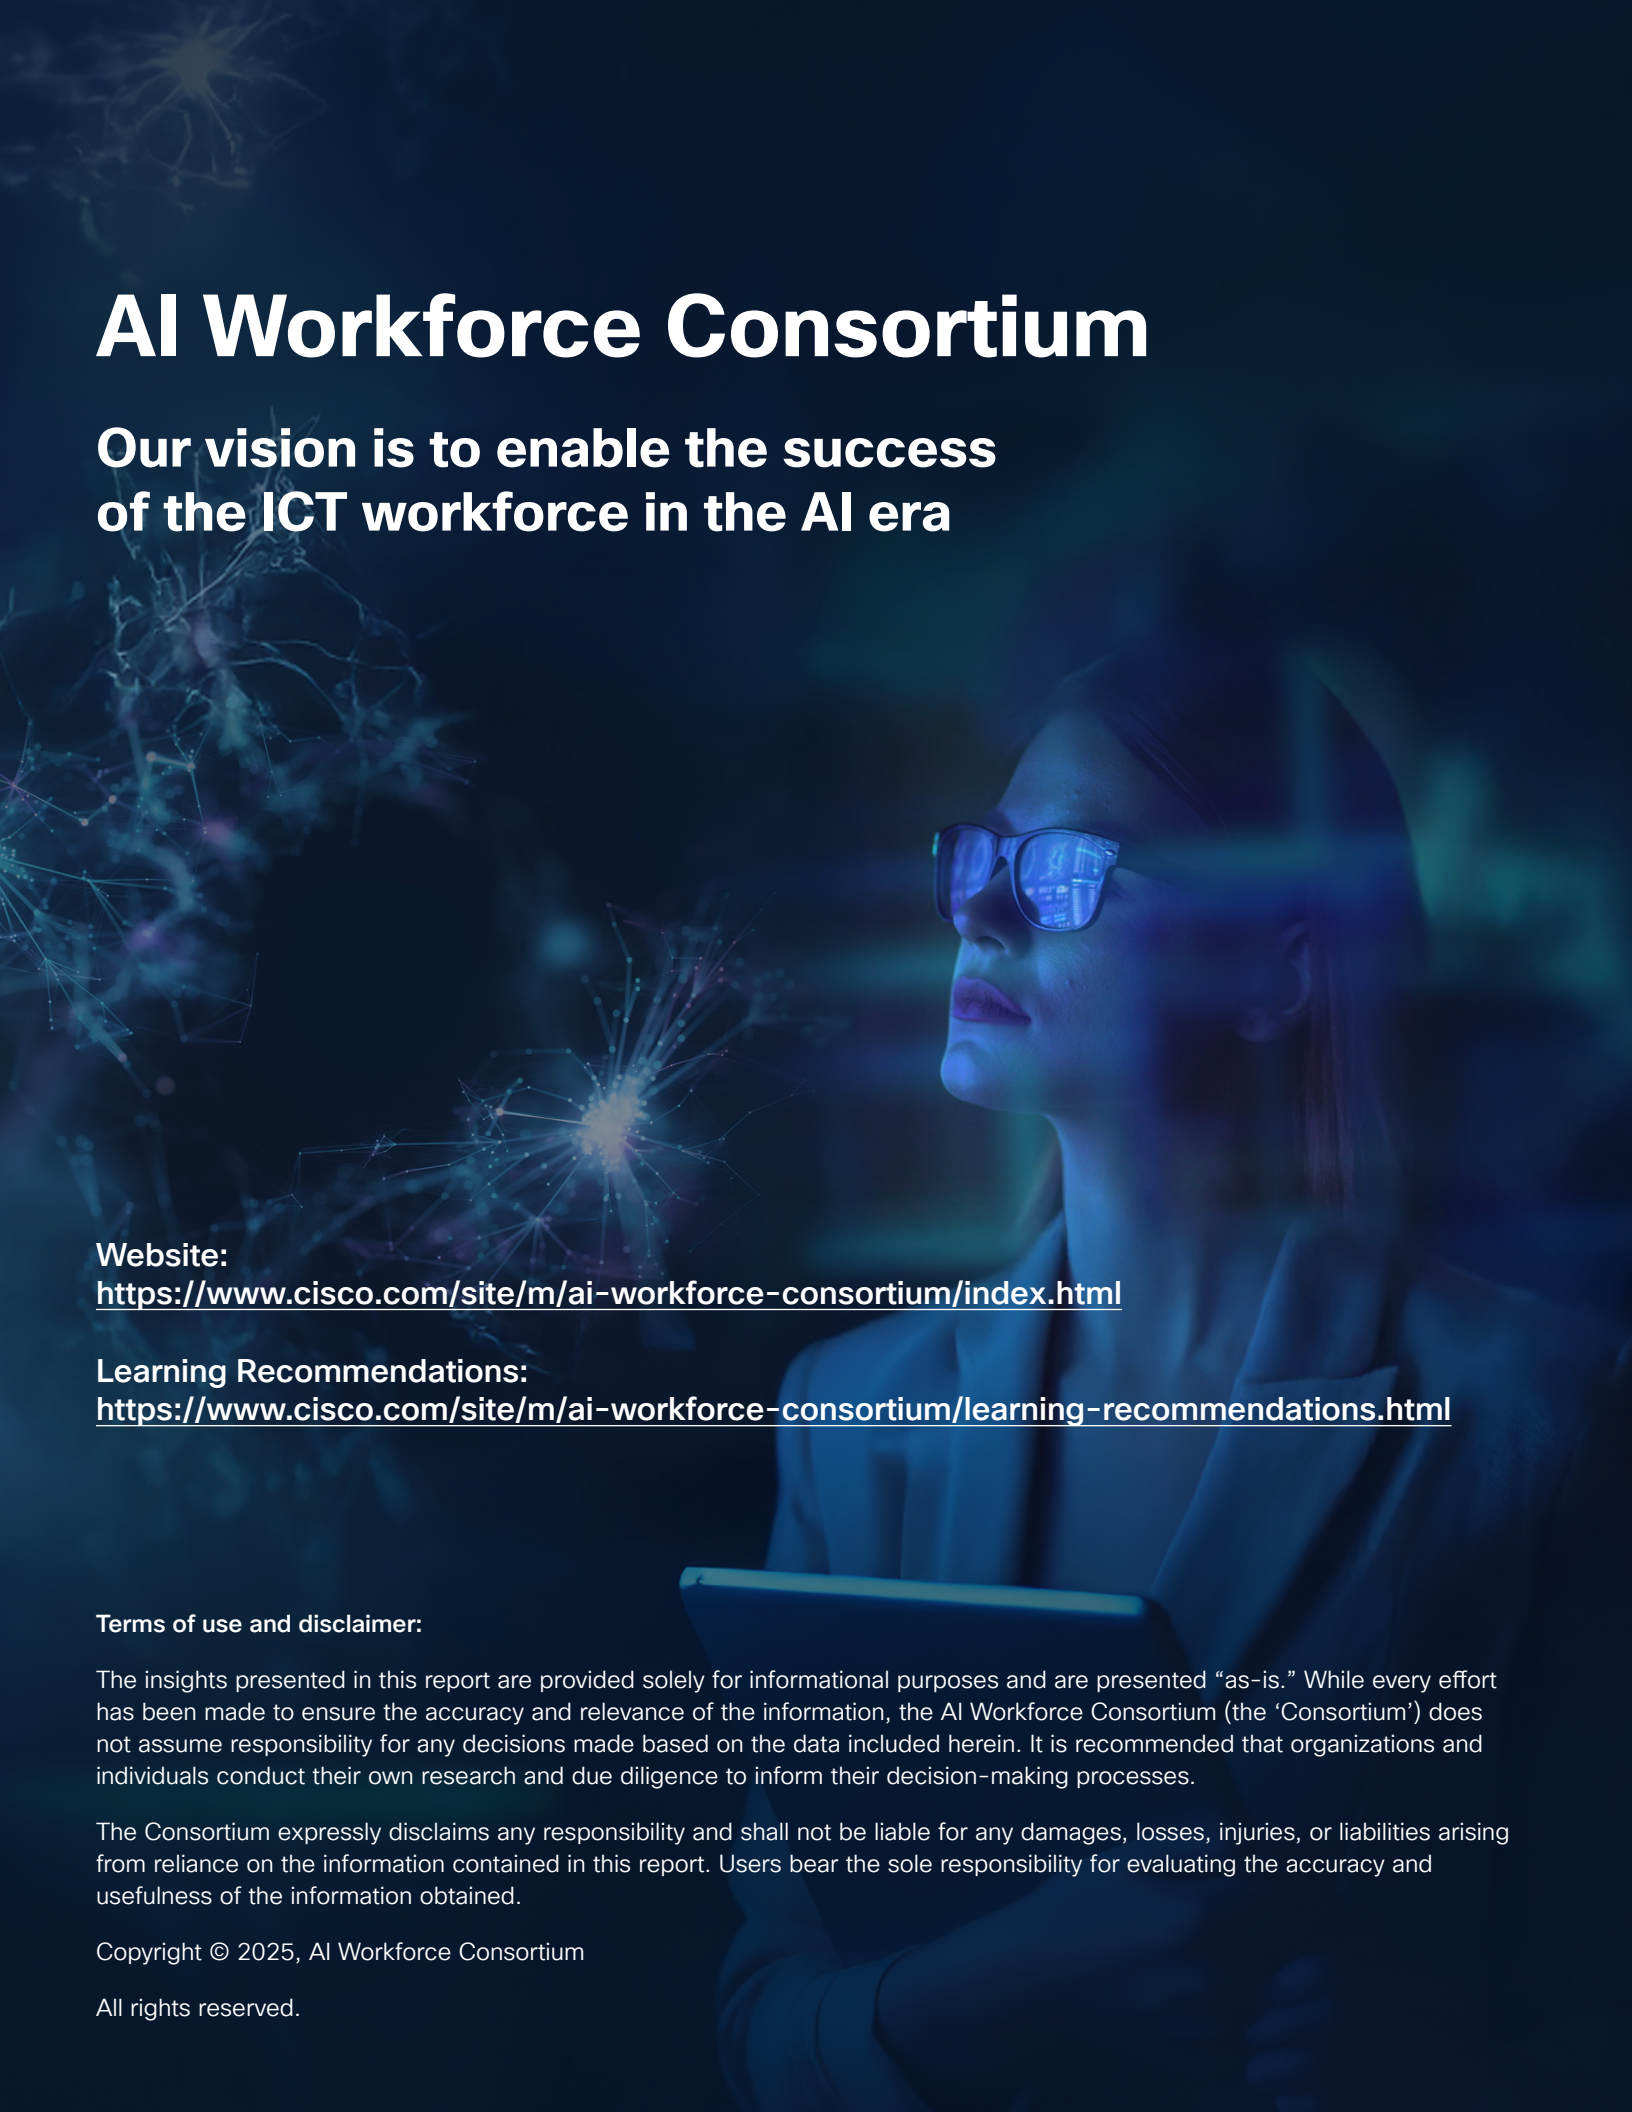

Website:
https://www.cisco.com/site/m/ai-workforce-consortium/index.html

Learning Recommendations:
https://www.cisco.com/site/m/ai-workforce-consortium/learning-recommendations.html