



Step 2

Give the Team
AI Basic Literacy



20 AI Keywords All Agilists Must Know

Basics of AI

Here are 20 key basic terms of AI that you should know:

1 AI (Artificial Intelligence)

A technology that enables computers to simulate human intelligence and problem-solving capabilities.

2 Generative AI

Generative AI refers to artificial intelligence models that can create new content, such as text, images, audio, or video. Examples are ChatGPT or DeepSeek.

3 ANI (Artificial Narrow Intelligence)

A type of AI designed to perform a single specific task. Examples are self-driving cars or defect detection in a factory line.

4 AGI (Artificial General Intelligence)

The futuristic potential form of AI that can perform any intellectual task a human can. We are not there yet.

5 Agentic AI

AI systems designed to autonomously make decisions, plan, and execute tasks to achieve goals with minimal human intervention.

6 AI Model

An AI system that has been trained on a dataset to recognize patterns, make predictions, or generate new content. Examples are GPT-4 model by OpenAI or Gemini 1.5 Flash model by Google.

7 Algorithm

A set of rules or instructions designed to enable machines to learn from data, make decisions, or perform tasks.

8 LLM (Large Language Model)

A type of AI trained on vast amounts of text data to understand and generate human-like language.

9 Multimodal

Models that can process and generate multiple types of data, like text, images, audio, and video. An example is Gemini 1.5 by Google that can process and generate text, images, audio, video, and code.

10 ML (Machine Learning)

A subset of AI that enables systems to learn from data, identify patterns and make decisions or predictions.

Basics of AI

Here are 20 key basic terms of AI that you should know:

11 DL (Deep Learning)

A subset of machine learning that uses multi-layered neural networks to learn complex patterns and create highly accurate outputs.

12 Prompt

A text input given to an AI model to guide its response or generate desired output.
Example: Write an email to stakeholders, inform them of the scope and goal of the Sprint, and invite them to the upcoming Sprint Review.

13 Prompt Engineering

The process of designing effective prompts to generate better and desired responses.

14 Supervised Learning

A model training approach to learn from paired input-output labeled data to predict outputs for new, unseen inputs.

15 Unsupervised Learning

A model training approach where a model discovers patterns and structures within unlabeled data.

16 Reinforcement Learning

A model training approach where a model learns to optimize decisions based on receiving rewards and penalties for actions. An example is when computers learn to play chess.

17 Diffusion Model

A generative AI technique that learns by iteratively adding noise to data and then denoise it to create new outputs like images. An example is that most AI images are generated by Diffusion Models.

18 RAG (Retrieval-Augmented Generation)

An AI technique that enhances large language models by retrieving information from external sources to generate more accurate and context-aware responses.

19 Fine-tuning

The further training of a pre-trained AI model on a smaller, domain-specific (like medical or legal domains) dataset to enhance its performance for a particular domain.

20 Token

The fundamental unit of text that an AI model processes, which can be a word, part of a word, or even a single character.

To learn more about the basics of AI,
watch this short video



[Watch Video](#)

Responsible & Ethical AI

Responsible & Ethical AI Principles



Transparency & Explainability



Privacy, Security, and Data Protection



Fairness,
Non-Discrimination
(Bias mitigation)



Inclusiveness, Sustainable
Development and Well-being



Safety, Robustness, and
Reliability



Human Oversight & Control



Accountability & Responsibility



Environmental Impact

1

Transparency & Explainability

Transparency in Ethical AI refers to openly sharing how AI systems are developed, trained, and deployed, ensuring stakeholders understand the processes and data involved.

Explainability focuses on making AI decisions interpretable to users, providing clear reasons for outcomes to build trust.

AI systems should not operate like a "black box." Instead, users and stakeholders should understand how and why an AI system makes its decisions or predictions.

Example:

You are a bank using AI to decide whether a customer is eligible for a loan or not. In one case, a customer asks you for the reason why you rejected their loan request, but you cannot explain it because you don't know how your AI system evaluates loan requests.

In this case, are you upholding the Transparency & Explainability principle of Ethical AI?!



Fairness, Non-Discrimination (Bias mitigation)

Fairness ensures that AI systems treat all individuals and groups equitably, avoiding favouritism or harm based on attributes like race, gender, or age.

Non-Discrimination (Bias mitigation) involves actively identifying and reducing biases in data, algorithms, and decision-making processes to prevent unjust or prejudiced outcomes.

Example:

Imagine your company uses an AI screening tool for hiring, but you notice it consistently recommends men for engineering roles more often than equally qualified women. When a highly skilled female candidate is rejected without clear justification, you question the system's decision.

In this case, are you upholding the Fairness, Non-Discrimination principle of Ethical AI?!



3 Safety, Robustness, and Reliability

Safety in Ethical AI ensures that AI systems operate without causing unintended harm to users, society, or the environment.

Robustness means the AI performs consistently under varying conditions, including adversarial attacks, edge cases, or noisy data, without failing unpredictably.


Reliability guarantees that the system functions accurately and dependably over time, maintaining correct and stable outputs to prevent errors that could lead to mistrust or dangerous outcomes.

Example:

A healthcare diagnostic AI used to detect cancer from X-rays initially performs with high accuracy. However, over time, it starts misclassifying cases because it was trained mostly on data from one hospital's imaging equipment.

When deployed in a new clinic with different machines, the AI fails to adapt, producing inconsistent and incorrect diagnoses.

In this case, are you upholding the Reliability principle of Ethical AI?!



4 Accountability & Responsibility

Accountability in Ethical AI ensures that organizations and individuals behind AI systems are answerable for their outcomes, including errors or harms, with clear mechanisms for redress.

Responsibility mandates that developers, deployers, and users uphold ethical standards throughout the AI lifecycle—from design to deployment—by proactively addressing risks and biases.

These two address the critical question of who is ultimately answerable when an AI system causes harm or makes errors. This principle ensures that humans maintain control and ownership over the outcomes generated by autonomous systems.

Example:

Imagine you work in a self-driving car product development company. One of your self-driving cars, operating under AI control, causes an accident resulting in pedestrian injuries. Who is primarily accountable for the incident?

In this case, Accountability for your AI systems lies with its creator (your company), who must ensure safety, ethical design, and compliance with regulations.



Privacy, Security, and Data Protection

Privacy in Ethical AI ensures that personal data is collected, used, and stored with respect for individual rights, minimizing unnecessary data exposure and obtaining informed consent.

Security involves safeguarding AI systems and their data from breaches, attacks, or misuse through robust encryption, access controls, and threat monitoring.

Data Protection guarantees compliance with legal standards (e.g., GDPR), ensuring data is handled transparently, retained only as needed, and anonymized where possible to prevent harm or re-identification.

Example:

With ChatGPT, you can change the settings of the tool to align it with your expectations as follows:

- Prevent your data from being used in the model training and improvement by turning off the "Improve the model for everyone" option.
- Use Temporary Chat to prevent storing your data.



Inclusiveness, Sustainable Development and Well-being

Inclusiveness in Ethical AI ensures that AI systems are designed for and accessible to diverse populations, including underrepresented groups, to prevent exclusion and promote equitable benefits.

Sustainable Development means AI solutions should support long-term environmental, social, and economic health, avoiding harmful shortcuts like excessive energy use or exploitative labor practices.

Well-being emphasizes that AI must prioritize human dignity and flourishing, safeguarding mental/physical health.

Example:

Imagine your company is developing a voice recognition AI tool that has been initially trained mostly on accents from one region, causing it to fail for users with different accents or speech patterns, like older adults, non-native speakers, or people with speech impairments.

In this case, are you upholding the Inclusiveness principle of Ethical AI?!



Human Oversight & Control

Human Oversight & Control ensures that humans retain ultimate authority over AI systems, with the ability to intervene, override, or stop decisions, especially in high-stakes scenarios like healthcare or criminal justice.

It mandates that AI operates as a tool to augment human judgment, not replace it, with clear protocols for monitoring outputs and addressing errors.

Example:

A hospital uses an AI diagnostic tool to recommend treatments for cancer patients. While the AI analyzes data quickly, doctors must review and approve every recommendation before it's applied.

If the AI suggests an aggressive treatment with high risks, the doctor can reject or modify the plan based on their expertise.



Environmental Impact

Environmental Impact addresses the ecological footprint of AI systems, from energy-intensive training processes (e.g., large language models) to hardware waste, urging sustainable practices like renewable energy use and efficient algorithms.

It emphasizes minimizing harm—such as reducing carbon emissions from data centers or avoiding e-waste from obsolete AI hardware—while ensuring AI advancements align with global climate goals.

Example:

Training a single AI model can emit as much CO₂ as 300 round-trip flights.



AI Fluency Framework

AI Fluency Framework*

An AI-aware Scrum Team member leverages AI within the boundary of the 4D AI fluency framework to ensure their interactions with AI are effective, efficient, ethical and safe.

This framework has four aspects:



* Developed by Anthropic

1 Delegation

Setting goals and deciding whether, when and how to engage with AI.

2 Description

Effectively describing goals to prompt useful AI behaviors and outputs.

3 Discernment

Accurately assessing the usefulness of AI outputs and behaviors.

4 Diligence

Taking responsibility for what we do with AI and how we do it.



Professional Scrum Master™ – AI Essentials

Now, it is decision time. If you want to go further and prepare yourself as a Scrum Master to support your team in their AI Transformation journey, join my upcoming Scrum.org official PSM-AI Essentials class.

[Enroll Now](#)



Professional Scrum Product Owner™ – AI Essentials

Now, it is decision time. If you want to go further and prepare yourself as a Product Owner to support your team in their AI Transformation journey, join my upcoming Scrum.org official PSPO-AI Essentials class.

[Enroll Now](#)

See you in Step 3 ...