



PUBLIC

Global Artificial Intelligence (AI) Ethics Policy

Table of contents

| | |
|---|-----------|
| 1. Introduction | 3 |
| 2. Purpose and Objectives | 3 |
| 3. Scope | 3 |
| 4. Terms and Definitions | 4 |
| 5. Roles and Responsibilities | 6 |
| 6. Ethical AI at SAP | 7 |
| Principle 1: Proportionality and Do Not Harm | 7 |
| Principle 2: Safety and Security | 8 |
| Principle 3: Fairness and Non-Discrimination | 8 |
| Principle 4: Sustainability | 9 |
| Principle 5: Right to Privacy and Data Protection | 9 |
| Principle 6: Human Oversight and Determination | 9 |
| Principle 7: Transparency and Explainability | 10 |
| Principle 8: Responsibility and Accountability | 11 |
| Principle 9: Awareness and Literacy | 11 |
| Principle 10: Multi-Stakeholder and Adaptive Governance and Collaboration | 11 |
| 7. Governance | 12 |
| Use case evaluation | 12 |
| Corrective Measures | 12 |
| 8. Supporting Documentation | 12 |

1. Introduction

SAP is committed to the ethical development, deployment, use, and sale of AI systems.

Foundational to SAP's approach to AI Ethics is the company's commitment through its [Global Human Rights Commitment Statement](#) to uphold and support the Universal Declaration of Human Rights. To meet the challenges posed by the rapidly evolving AI technologies and align to an internationally recognized set of values, SAP has chosen to ground this new version of the Global AI Ethics Policy in the "**Recommendation on the Ethics of Artificial Intelligence**" by UNESCO. Our principles are:

1. Proportionality and do not harm
2. Safety and Security
3. Fairness and non-discrimination
4. Sustainability
5. Right to privacy and data protection
6. Human oversight and determination
7. Transparency and explainability
8. Responsibility and accountability
9. Awareness and literacy
10. Multistakeholder and adaptive governance and collaboration

2. Purpose and Objectives

SAP believes that AI presents vast opportunities not only for businesses but also for governments and society. But, like all great technological advancements, it also has the potential to create economic, political, and social challenges, depending upon how it is used and implemented. SAP is committed to producing relevant, reliable, and responsible solutions. The policy defines an SAP group-wide ethical framework for the development, deployment, use, or sale of AI systems, which complements the rules and regulations issued by national and international governments and organizations.

To learn more about the operationalization of the policy, please refer to the [Responsible AI](#) website.

3. Scope

This policy applies to the SAP group and all actors (i.e., SAP employees who play an active role in any phase of the AI system life cycle). For SAP employees who are not actively participating in the AI system life cycle, it is relevant to understand:

- Which ethical principles we apply to AI systems and how we apply them.
- Which feedback mechanisms are available in case the AI system is not adhering to the ethical principles.
- What the governance process is for AI systems and solutions.

If any member of the SAP group has its own AI ethics policy, it must be adapted within a reasonable time to comply with the rules and terms of the effective version of this policy.

The policy defines SAP's intent, expectations, and commitment to the ethical development (including writing software code, generating model-based code, consuming and integrating open source/third party/freeware components or libraries, integrating and configuring web services), deployment, use and sale of AI systems. The ethical principles outlined in this policy apply to the solutions as defined in section 4 of this policy.

Although SAP makes every effort to ensure that AI systems used by customers comply with the ethical principles described in this policy, it is also essential that customers use it in compliance with all laws and regulations including any specific national or sectoral regulatory requirements applicable to them and within the purpose intended by SAP. SAP may also issue specific terms and conditions for the acceptable use of AI systems (certain types thereof) that are binding to customers. Further information on documentation is outlined in section 6, principle 7.2 “Documentation”.

4. Terms and Definitions

| | |
|-------------------------------------|--|
| Actors | Actors are those SAP employees who play an active role in any phase of the AI system life cycle, including design, verification and validation, development, sale, operation, monitoring or use of the AI system. |
| Affected Individuals | Affected individuals encompass all individuals directly or indirectly affected by AI systems or decisions based on the output of AI systems. These individuals do not necessarily interact with the deployed solution. |
| AI (Artificial Intelligence) | Information technology that performs tasks that would ordinarily require biological brainpower to accomplish, such as making sense of spoken language, learning behaviors or solving problems. |
| AI Ethics | A set of values, principles and techniques that employ widely accepted standards of right and wrong to guide moral conduct in the development, deployment, use and sale of AI systems. |
| AI System Life Cycle | AI system life cycle phases involve: i) ‘design, data and models’; which is a context-dependent sequence encompassing planning and design, data collection and processing, as well as model building; ii) ‘verification and validation’(including testing); iii) ‘deployment’; and iv) ‘operation and monitoring’. These phases often take place in an iterative manner and are not necessarily sequential. The decision to retire an AI system from operation may occur at any point during the operation and monitoring phase. |
| AI System | <p>In general, an AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.</p> <p>For the purposes of this policy, an AI system shall be deemed:</p> <ul style="list-style-type: none"> • an AI system developed by SAP, whether commercialized or not; • an AI system developed by a SAP supplier and embedded in the solution; • an AI system developed by a SAP supplier and sold by SAP, like supplier-branded reseller or SAP branded reseller; and/or |

| | |
|--|--|
| | <ul style="list-style-type: none"> • an AI system developed by an SAP supplier and used by SAP for internal purposes. |
| Bias | An inclination or prejudice for or against one person or group, especially in a way considered to be unfair. |
| End-User | An end-user is the person who ultimately uses or is intended to ultimately use a solution. |
| Explainability | In the context of AI and machine learning, explainability is the ability of an AI system to provide clear and concise explanations of its predictions, decisions, behavior and processes. AI explainability aims to make the internal mechanisms of an AI system and its decision-making process transparent, understandable and coherent to humans. |
| Fairness | Impartial and just treatment or behavior without favoritism or discrimination. |
| Human-in-command | Refers to the capability to oversee the overall activity of the AI system (including its broader economic, societal, legal and ethical impact) and the ability to decide when and how to use the AI system in any particular situation. This can include the decision not to use the AI system in a particular situation, to establish levels of human discretion during the use of the AI system or to ensure the ability to override a decision made by the AI system. |
| Indirectly affected individuals | Indirectly affected individuals are people who have suffered consequences, other than or in addition to direct effects, over time due to disruption or changes in the economy, critical infrastructure, basic services, AI infrastructure, commerce or work, or social, health, and psychological consequences. |
| Human-in-the-loop | Refers to the capability for human intervention in every decision cycle of the AI system. |
| Human-on-the-loop | Refers to the capability for human intervention during the design phase of the AI system as well as the monitoring of the AI system's operation. |
| SAP Group | SAP SE and all its subsidiaries worldwide, i.e. all companies in which SAP SE directly or indirectly holds a majority ownership interest or which SAP SE directly or indirectly controls. |
| Solution | <p>A set of related software programs and/or services with an AI system embedded that are offered as a single package by SAP.</p> <p>Solution covers:</p> <ul style="list-style-type: none"> • Solutions offered on the SAP price list; • Solutions developed by SAP, not offered on the SAP price list, but still externally provided by SAP; |

| | |
|---------------------|--|
| | <ul style="list-style-type: none"> • Solutions developed by SAP and used for SAP internal purposes only. |
| Supplier | A person or legal entity that is providing services, including professional/consulting service, physical items as well as software, data, scripts, or code, and any goods conceived, produced, or developed in connection with the services under an order to the SAP entity which has ordered the supplies or to a third party on behalf of SAP. |
| Transparency | <p>A characteristic of an AI system that involves the justifiability of the processes that goes into its design, development and implementation to provide its outcome – i.e. the soundness of the justification of its use.</p> <p>Transparency is characterized by visibility or accessibility of information or the characteristic of being easy to see through. The principle of transparency entails that development and implementation processes are justifiable through and through. It demands as well that an algorithmically influenced outcome is interpretable and made understandable to affected parties.</p> |

5. Roles and Responsibilities

| | |
|-------------------------------------|--|
| AI Ethics Steering Committee | <p>Description:</p> <ul style="list-style-type: none"> • Assembly of SAP senior leaders representing all relevant company functions. <p>Responsibilities:</p> <ul style="list-style-type: none"> • Overseeing SAP's AI Ethics approach. • Providing strategic decisions and directions. • Accountable for approving high-risk use cases. |
| AI Ethics Advisory Panel | <p>Description:</p> <ul style="list-style-type: none"> • External experts from academia, policy development and industry. <p>Responsibilities:</p> <ul style="list-style-type: none"> • Providing independent and unbiased feedback and recommendations on SAP's AI Ethics processes, policies, use-cases, and decisions. |
| AI Ethics Office | <p>Description:</p> <ul style="list-style-type: none"> • Central function coordinating the AI Ethics governing bodies. <p>Responsibilities:</p> <ul style="list-style-type: none"> • Organizing and hosting the assemblies of the AI Ethics Steering Committee and the AI Ethics Advisory Panel. • Reviewing and updating the AI Ethics policy and guiding principles. • Acting as the first point of contact for any AI Ethics related queries. • Orchestration of the high-risk use case reviews. |

| | |
|--------------------------------------|---|
| Trustworthy AI Workstream | <p>Description:</p> <ul style="list-style-type: none"> • Coordinating AI Ethics related activities across the company. <p>Responsibilities:</p> <ul style="list-style-type: none"> • Defining and implementing measures to ensure adherence to the AI Ethics policy. • Providing education and training material to employees and public. • Setting up and operating a global community for responsible AI. |
|--------------------------------------|---|

6. Ethical AI at SAP

Any actor involved in any phase of the AI system life cycle must ensure that AI systems are designed, developed, sold, or used in compliance with the ethical principles specified in this policy and that these ethical principles are ensured and followed throughout the AI system life cycle. Adherence to these ethical principles is especially important so that customers and the public continue to trust in SAP's ability to develop, deploy, use, and sell AI systems in a responsible and ethical way.

Principle 1: Proportionality and Do Not Harm

[Human rights and fundamental freedoms](#) must be respected during all stages of an AI system life cycle. Impact classification mechanisms and measures to prevent harm shall be defined to ensure the safeguarding of human rights and fundamental freedoms. AI systems shall not exceed what is necessary to achieve the goal that has been set and should be used appropriately to the context.

How this principle is implemented at SAP

- Actors who develop an AI system define its purpose and tasks. The scope defined should match the behavior of the solution deployed.
- The context and environment of the AI system must be analyzed, and appropriate measures must be taken to prevent misuse or unintended harm to humans.
- In addition, SAP defines a set of “red lines” along the categories of personal freedom and society, that, if identified, prohibits continuation of the design and development of the AI system. Below, the details are outlined.

1.1 Personal Freedom

- AI systems must not be developed or deployed, used, and sold for **human surveillance** that is utilized to target individuals or groups, either by biometrics, facial recognition, or other identifiable features, to disregard or abuse the human rights of the individuals or groups.
- AI systems must not be developed, deployed, used or sold to **discriminate** against individuals or a group of people, or to exclude individuals or a group of people from equal access to AI's benefits and opportunities, in accordance with SAP's Global Export and Sanctions Compliance Policy and SAP's Code of Ethics and Business Conduct.

- AI systems must not be developed, deployed, used, or sold to **deanonymize** already anonymized data which may result in the identification of individuals or groups.

1.2 Society

- AI systems must not be developed or deployed for **deception or unfair manipulation** of individuals or groups via public forums, media, or moderation of other similar uses.
- AI systems must not be developed or deployed to undermine debates or democratic electoral systems.
- AI systems must not have any **harmful impacts** on users.
- AI system shall not be developed or deployed to be used for the evaluation or classification of the trustworthiness of affected individuals over a certain period of time based on their social behavior or known or predicted personal or personality characteristics, such as social scoring.
- The context and environment of the AI system must be analyzed, and appropriate measures must be taken to prevent misuse or unintended harm to humans.

Principle 2: Safety and Security

Safety and security measures must be guaranteed through the AI system life cycle to ensure protection from attacks and unwanted harm.

How this principle is implemented at SAP

- Compliance with SAP's Global Security Policy must be ensured.

For high-risk use cases additional measures apply:

- AI systems must be tested extensively to avoid unintended behavior. Tests cases and results should be documented as prescribed by SAP development processes (SAP Global Development Policy).
- Potential safety and security risks must be identified, a range of acceptable errors should be defined, and mitigation strategies should be in place and documented throughout the AI system life cycle.
- Fallback/default mechanisms: in case of failure of the AI system or some errors detected within the output generated by the AI system, the AI system should be reverted to predefined behavior to minimize risks or harm that could be caused by the failure or error.
- An AI system should be monitored continuously so that, in conditions of normal use, foreseeable use or misuse, or other adverse conditions, the AI system functions appropriately and does not pose unreasonable safety risk.

Principle 3: Fairness and Non-Discrimination

Data-driven technologies can reproduce, reinforce, and amplify patterns of marginalization, inequality, and discrimination that exist in society and may be encoded into data sources used for the development of AI systems. These biases can impact negatively both the development and outputs of AI systems and, in turn, affected individuals. AI systems must be designed, developed, deployed, and used promoting and protecting fairness by respecting all forms of diversity and making all reasonable efforts to avoid discriminatory or biased outcomes. Care shall be taken especially when there is a risk of unjustly impacting underrepresented groups.

How this principle is implemented at SAP

3.1 Data

- The AI system shall be tested on a dataset that is as expansive, feasible, representative, relevant, accurate, and generalizable.

3.2 Fairness

- The output of the AI system is fair and free of discrimination.
AI systems shall not intentionally generate unfairly biased outputs. When appropriate with regard to the AI use case, SAP will conduct reviews of AI system outputs for fairness and equity.
- In addition to the conditions laid down in SAP's Global Development Policy and SAP Product Standards, processes shall be put in place to test and monitor for **potential biases**.
- The AI system must comply with the SAP Accessibility Standards.
- Fairness guarantees should be assessed and monitored continuously, and adapted when necessary, so that the AI systems remain aligned with changing societal values.
- Where feasible given the context and environment where the AI system shall operate, the AI system shall enable the affected individuals to evaluate the outputs and object if they are not fair or discrimination free. SAP will implement corrective measures in the AI system if the objection is deemed justified. Mechanisms to provide feedback are described in Section 7 "Governance".

Principle 4: Sustainability

AI systems should foster environmental, economic, and social sustainability. Negative social, economic and environmental impacts must be avoided.

How this principle is implemented at SAP

- The development, deployment, use, and sale of AI systems shall be conducted in accordance with [SAP's sustainability policies](#).

Principle 5: Right to Privacy and Data Protection

AI systems should be developed, deployed and used respecting privacy and data protection laws, recognizing their vital role in protecting citizens' fundamental rights.

How this principle is implemented at SAP

- The design, development, testing, and SAP's usage of AI systems must comply with the SAP's Global Data Protection and Privacy Policy, as well as SAP's Global Development Policy and SAP Product Standards.

Principle 6: Human Oversight and Determination

AI systems shall be subject to appropriate human oversight. The rights and freedoms of affected individuals shall exceed those of AI systems.

How this principle is implemented at SAP

- Human oversight shall be achieved through an appropriate **governance mechanism**. This could include but not be exclusive to Human-in-the-loop, Human-on-the-loop, or Human-in-command. An appropriate form of human oversight must be chosen for each AI system, considering the context of the use case, state of art, and risks of interference with the rights and freedoms of the affected individuals.
- To learn more about feedback mechanisms, have a look at the “Governance” section, “Feedback Mechanisms”.

Principle 7: Transparency and Explainability

AI systems capabilities, intended use, and limitations must be communicated clearly to customers and internal end-users along with the necessary technical tools for training and prediction. Transparency of the AI systems development process as well as the transparency and explainability of its decisions and behaviors should be ensured.

How this principle is implemented at SAP

7.1 Transparency

- End-users must always be aware of their interaction with an AI system.
 - Where applicable, when interacting directly with humans (including via conversational AI or ‘chatbots’), AI systems must be made identifiable as such.
 - The AI system shall be developed in such a way that it does not encourage end-users’ affection and/or empathy.
 - AI systems shall clearly signal to end-users that the social interaction is simulated.
- Developers of an AI system shall endeavor to make the decisions, proposals, and outputs of the AI system as transparent as possible. The level of transparency should be appropriate to the context of the use case, and consistent with the state of art.

7.2 Documentation

- To ensure transparency, each AI system must be accompanied by a comprehensive set of documentation. This includes:
 - Scope: capabilities, intended use and limitations.
 - Development: datasets (including data gathering, labeling), models, algorithm, training and validation procedures used.
- In addition, methods used to develop, test, and validate the outcomes or decisions made by the AI system are to be documented according to SAP’s Global Development Policy and SAP Product Standards.
- Documentation shall be adjusted, if needed, when the scope of the AI system changes (e.g. functionalities are added), or if a new version of the AI system is released.

For high-risk use cases additional measures apply

- For high-risk AI systems, the documentation must pay increased attention to the deployment and operation phase, especially monitoring, supervision, control and troubleshooting. Detected deviations of

the AI system from the intended behavior shall be promptly documented by the actor responsible for monitoring the AI system and shared with other actors.

7.3 Explainability

- To the extent appropriate to the context of the use case, a clear and simple explanation of how a decision was made by an AI system shall be provided to the end-user or affected individual upon request.
- Where applicable, information about the confidence level of the output of the AI system should be provided to the end-user.

Principle 8: Responsibility and Accountability

AI systems are not morally accountable agents and cannot be held accountable for their actions. Therefore, ethical and legal responsibility must always be attributable to actors (humans) and existing legal entities.

How this principle is implemented at SAP

- Individual actors are accountable for the responsible development, deployment, use and sale of the AI system, depending on their involvement in a particular phase of the AI system life cycle. Governance mechanisms, impact and policy compliance assessments are in place to ensure that AI systems are developed, deployed, sold, and used ethically and responsibly.

Principle 9: Awareness and Literacy

To raise awareness on the topic of AI, and help build new skills, SAP provides learning opportunities regarding AI and its responsible use, both for its workforce as well as for external stakeholders. Further, SAP commits itself to promoting global initiatives on digital literacy and access.

How this principle is implemented at SAP

- Free courses and learning materials accessible for everyone can be found in [SAP's learning site](#) or on the [SAP Community](#) page.
- In the [SAP Community groups](#), there are dedicated areas to foster discussion and encourage exchange around the topic of AI.
- SAP sponsors and organizes educational events both on site and virtually such as [SAP TechEd](#).

Principle 10: Multi-Stakeholder and Adaptive Governance and Collaboration

SAP is committed to respecting international law and national sovereignty in the use of data. Further, it is important to ensure a human-centered approach promoting interdisciplinarity and intersectionality.

How this principle is implemented at SAP

- Whenever feasible given the context and specifics of the AI system and individual use case, affected individuals should be involved in the process of validating the output. Affected individuals have a right to

challenge the output's non-compliance with the ethical principles set in the policy (e.g. discriminatory or biased nature of the output).

- The AI Advisory Panel is part of SAP's governance bodies. It supports the company in collecting feedback and getting fresh input from academia and other industries. For more information, please refer to Section 5 "Roles and Responsibilities".
- For additional information on governance, please refer to section 7.

7. Governance

Use case evaluation

- Governance processes are in place to ensure that the AI systems adhere to the ethical standards set in this policy.
- Each use case is carefully analyzed using different assessment methods. Upon the categorization of the use case, several checkpoints and further investigations are required before releasing the AI system.
- Each AI system developed, deployed, used or sold by SAP undergoes an ethical assessment. Depending upon the results, different courses of action will be taken.
 - Minimal or no risk use cases: the development, deployment, sale and use of the AI system is allowed.
 - High risk use cases: If an AI system is categorized as high risk, panels of experts will discuss the implications and eventually decide whether the use case should move forward.
 - Forbidden cases: Should an AI system be classified as forbidden, SAP will not allow further development, deployment, sale, and use.
- A periodic evaluation of the AI system's safety, stability, and adherence to the ethical principles must be performed and documented. The frequency of the review should be established in accordance with the risk categorization.
- If an end-user or affected individual both at SAP or outside the company has concerns about the adherence of an AI system with this policy and wishes to raise these concerns anonymously, they can use the [Speak Out](#) tool.

Corrective Measures

- If any of the fundamental rights of the affected individuals are hurt, corrective measures shall take place as per the pre-defined human rights due diligence process.

8. Supporting Documentation

This Global AI Ethics Policy is supported by various other policies, standards, procedures, and guidelines. The following are relevant to this policy and complement rules and governance for ethical development, deployment, use and sale of AI systems:

- [Global Code of Ethics and Business Conduct for Employees](#)
- [SAP Global Human Rights Commitment Statement](#)
- [SAP Global Antidiscrimination Statement](#)

- [SAP Global Environmental Policy](#)
- [SAP Partner Code of Conduct](#)
- [SAP Supplier Code of Conduct](#)
- [SAP Trust Center](#)