SAP

Global Artificial Intelligence (AI) Ethics Policy

Version Number: 1.0
Effective Date: January 1, 2022

CONTENTS

1 INTRODUCTION

SAP is committed to the ethical development, deployment, use, and sale of SAP developed Artificial Intelligence (AI) systems. The policy defines a group-wide minimum standard for the development, deployment, use, or sale of SAP's Artificial Intelligence systems. It defines requirements for SAP's business processes that involve AI and assigns clear responsibilities.

The policy is based on, and has been built upon the foundation of the SAP Guiding Principles for Artificial Intelligence, established in 2018, as laid out below:

- We are driven by our values.
- We design for people.
- We enable business beyond bias.
- We strive for transparency and integrity in all that we do.
- We uphold quality and safety standards.
- We place data protection and privacy at our core.
- We engage with the wider societal challenges of Artificial Intelligence.

2 PURPOSE AND OBJECTIVES

SAP believes that Artificial Intelligence has the potential to unlock abundant potential for businesses, governments, and society. But, like all great technological advancements, AI also has the potential to create economic, political, and social challenges, depending upon how it is used and implemented. In addition, the speed at which the technology has moved into common usage has outpaced guidance from governmental policymakers on acceptable use. For these reasons, the development, deployment, use, and sale of AI systems at SAP needs to be governed by clear rules of ethics that are aligned with SAP's established guiding principles for AI and its core organizational values, supplemental to any existing or pending legislation.

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, and to respect, promote, and support internationally recognized human rights and widely accepted international norms.

An essential part of this commitment is the prohibition of discrimination and harassment of humans based on personal factors (including for example culture, race, ethnicity, religion, age, gender, sexual orientation, gender identity, physical or mental disability etc.), in order to promote respect for human autonomy and to ensure that the moral worth and dignity of all human beings is respected. In addition, SAP's goals for AI systems include protecting people from harm, looking after the well-being of others, treating all individuals equitably and justly, ensuring the entitlement to equal freedom and dignity under the law, the protection of civil, political and social rights, the universal recognition of personhood, and the right to free and unencumbered participation in the life of the community.

3 SCOPE

This Policy applies to SAP and its employees worldwide involved in the development, deployment, and sale of SAP developed AI systems. If any SAP subsidiary has its own AI Ethics Policy, upon enactment of the Global AI Ethics Policy, it immediately must align with this Policy.

The Policy defines SAP management's intent and expectations for the ethical development, deployment, and sale by SAP employees of SAP-developed AI systems and the principles may be used to provide guidance and counsel to customers and partners.

However, in ensuring the rights of customers to own and manage their own data SAP may have limited knowledge of how they ultimately use SAP AI products and services. SAP shall endeavor to educate and advise customers and partners based on the principles in this policy, recognizing that specific use case responsibility lies with the customer or partner, acting in good faith in accordance with local laws and regulations. SAP reserves the right to take appropriate action in the event of a customer violation of the principles of this policy.

Other policies that were in place before this Policy was enacted may remain in force but must be adapted to comply with the rules of this Policy upon the next review or modification cycle.

4 TERMS AND DEFINITIONS

SAP Group	SAP SE and its subsidiaries.
SAP Subsidiary	Any legal entity where SAP SE holds more than fifty percent (50%) of the shares or voting rights or such entity is controlled by the SAP SE.

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 technologies.
Artificial Intelligence (AI)	Typically defined as the ability of a machine to perform cognitive functions we associate with human minds, such as perceiving, reasoning, learning and problem solving. A system's ability to correctly interpret external data, to learn from such data and to use those learnings to achieve specific goals and tasks through flexible adaptation. SAP differentiates between two types of AI systems: **Rule-based AI systems** are characterized by the fact that the behavior of their components is fully defined by rules created by human experts. These systems are often described as symbolic or expert systems. **Learning-based AI systems** are differentiating themselves by the fact that their initial configuration made by humans is only the basis for the final form of their functions. With the help of data, they train how to solve a problem and continuously adapt their function in this process. For learning-based AI systems, humans define the problem and the goal, but the behavior rules and relationships required for the system are learnt in an automized way. In addition, hybrid systems containing both learning-based and rule-based methods are available. Impact Assessment shall be performed on such Learning-based AI systems during its development phase and subsequent phases to ensure that there are no unintended consequences.
AI Systems	Any AI based algorithm, component and/or software embedded in SAP products and services (e.g. AI platform, databases, embedded or stand-alone applications, features and solutions, Conversational AI ('Chatbots')).
Agency	The capacity of individuals to act independently and to make their own free choices.
Application	Any program, or group of programs, that is designed for the end user.
Autonomy	The capacity to make an informed, uncoerced decision.
Bias	An inclination or prejudice for or against one person or group, especially in a way considered to be unfair.
Black Box Algorithms	Any artificial intelligence system whose inputs and operations are not visible to the user or another interested party. A black box, in a general sense, is an impenetrable system. Deep learning modeling is typically conducted through black box development: The algorithm takes millions of data points as inputs and correlates specific data features to produce an output. That process is largely self-directed and is generally difficult for data scientists, programmers and users to interpret.
Data Subject	Any individual person who can be identified, directly or indirectly, via an identifier such as a name, an ID number, location data, or via factors specific to the person's physical, physiological, genetic, mental, economic, cultural or social identity.
Deep Neural Network	A technology developed to simulate the activity of the human brain – specifically, pattern recognition and the passage of input through various layers of simulated neural connections.
Developed (as relates to AI systems)	Includes Design, SAP Standard Development, Prototyping, Co-innovation, and Innovation and Cloud Service team's development activities.

Explainability	The ability to explain both the technical processes of an AI system and the related human decisions (e.g. application areas of a system). Technical explainability requires that the decisions made by an AI system can be understood and traced by human beings.
Fairness	Impartial and just treatment or behavior without unjust favoritism or discrimination.
Feature	A distinguishing characteristic of a software item (e.g. performance or functionality).
Human-in-the-loop	AI System functionality for human beings to intervene in every decision cycle of the system
Human-on-the-loop	Refers to the capability for human intervention during the design cycle of the system and monitoring the system's operation.
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 system in any particular situation. This can include the decision not to use an AI system in a particular situation, to establish levels of human discretion during the use of the system or to ensure the ability to override a decision made by a system.
Normative	In the context of practical ethics, 'normative' refers to a given concept, value, or belief that puts a moral demand on one's practices, i.e. that such a concept, value or belief indicates what one 'should' or 'ought to' do in circumstances where that concept, value or belief applies.
Platform	Any hardware or software used to host an application or service.
Software	A set of instructions, data or programs used to operate computers and execute specific tasks.
Solution	A set of related software programs and/or services that are sold as a package.
Subordinate	Treat or regard as of lesser importance than something else
Transparency	Characteristic of AI that involves the justifiability of the processes that go into its development and the implementation and of its outcome – i.e. the soundness of the justification of its use. Transparency is characterized by visibility or accessibility of information or the characteristic of being easily seen through and explained. 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.

5 ROLES AND RESPONSIBILITIES

AI Ethics Office	**Context:** • Different AI Ethics stakeholders need to be connected and aligned. **Structure:** • Orchestration of the AI Ethics Office is the responsibility of SAP's Chief Sustainability Officer team. **Responsibilities:** • Convenes the AI Ethics Steering Committee. • Is the contact for a L1 unit if a L1 unit is unable to make a use case decision regarding AI Ethics or questions or concerns remain.
AI Ethics Steering Committee	**Context**: • As an emerging technology with significant innovation potential without standardization so far, execution of ethical AI standards requires a cross-company approach anchored in strong LOB expertise and perspectives. **Structure**: • Orchestration of the AI Ethics Steering Committee is the responsibility of SAP's AI Ethics Office. • AI Ethics Steering Committee is a 'sub-committee' of the Sustainability Council. • In addition to representatives with AI product, legal, governmental and ethical expertise, members include representatives responsible for non-standard AI development and implementation. **Responsibilities:** • Serve to advise SAP personnel on how specific use cases are affected by this policy and the related guiding principles • Monitor, evolve and update SAP's Guiding Principles for AI. • Oversee ongoing updating of SAP's AI Ethics policy to guide implementation of ethical AI in all LOBs developing, deploying or selling AI systems. • Provide guidance/direction on AI use cases to LOBs . • Provide recommendations to SAP Executive Board on critical product/customer use cases with wider implications for SAP as a company. • Provide Sustainability Council with bi-annual update on status/progress (written). • Orchestrate AI Ethics use case review process. • In instances of unresolvable AI System use case conflict between AI Steering Committee guidance/direction and relevant Business Unit, prepare the case and recommendation for Sustainability Council review. • Align with and obtain insights from External Advisory panel on AI. • Steering committee members to offer employees, through some of the existing channels, opportunity to provide feedback and exchange dialog on AI Ethics on regular basis.
Sustainability Council	**Context**: • Sustainability Council received mandate (2018) to act as internal 'Ethics Advisory Board' for SAP, reviewing critical behavior and making recommendations to the SAP Executive Board on appropriate action. **Structure:** • Orchestration of the Sustainability Council is the responsibility of the Chief Sustainability Officer's office • The council is the highest-level governance structure for sustainability at SAP

	• Senior leaders from all board areas and specific business units are members of the council **Responsibilities**: • Oversight of AI Ethics implementation as reported by AI Ethics Steering Committee on bi-annual basis. • Provide longer term strategic direction for AI Ethics related topics at SAP. • Review critical use cases, as elevated from AI Ethics Steering Committee. • Provide recommendations to SAP Executive Board on critical use cases with wider implications for SAP as a company.
Important Stakeholders impacted by Policy	• **Developers and Data Scientists** should implement and apply the requirements to the development processes. • **Sales or Consulting/Deployment Personnel** (e.g. Admin or Data Processors) should ensure that the systems they sell and the products and services they offer meet the requirements. • **Executive Board** should support appropriate level of governance structure to ensure the requirements are met. In addition, where necessary, the board must take a decision on whether to allow or prohibit a usecase as per the defined escalation process. • **End-users, customers, partners and the broader society** should be informed about these requirements and be able to request that they are upheld.

6 ETHICAL AI AT SAP

Whilst different groups of stakeholders have different roles to play in ensuring that the requirements are met, end-users, customers, partners, and the broader society should be informed about these requirements and be able to request that they are upheld.
AI systems shall only be developed, deployed, used or sold by SAP in accordance with the principles laid out below.

Human Agency and Oversight:
Employees need to consider the following when developing AI systems as they relate to Human Agency and Oversight:

- Before implementation, the decision-making degrees of freedom of the AI system must be defined.
- When two or more AI systems are connected to each other or embedded within each other, additional testing and control measures should be performed for the individual AI systems as well as for the overall system
- The target definition of the AI system must be given by a human.
- AI systems shall be subject to appropriate human oversight, and the rights and freedoms of a human shall exceed that of AI systems.
- 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, and shall be decided on a case-by-case basis. The use cases will be first reviewed based on the impact assessment list. Selected use cases shall be brought to the attention of the steering committee for further deliberation and decision.
- In alignment with the SAP Global DPP Policy and SAP DPP Guidance documents, in situations where humans may be directly impacted by a decision made by SAP's AI system, human oversight shall be introduced to safeguard that AI system does not undermine human autonomy or introduce unintended consequences.
- As far as is practical, a clear and simple explanation shall be provided as to how decisions were made by an AI system used in automated decision processes.
- When human-on-the-loop models are used, appropriate extensive testing and governance shall be conducted during development and deployment to ensure the system behaves as intended by the developers and does not have any unintended behavior, outputs, or usage.

Addressing Bias and Discrimination:
AI systems often gain insights from the existing structures and behavior of the societies they analyze. As a result, 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 creation of AI. Equally, because many of the features, metrics, and analytic structures of the models that enable data mining are chosen by their developers, AI systems can potentially replicate their developers' preconceptions and biases.
Finally, data samples used to train and test algorithmic systems may be insufficiently representative of the populations or the past situations from which they are drawing inferences. This may include cases where the type of business, industry, or enterprise from which the original datasets were obtained are inappropriate for the AI systems being developed and deployed.
These biases can negatively impact both the development and outputs of AI systems and, in turn, customers or end users. Particular care shall be taken when there is a risk of causing discrimination or of unjustly impacting underrepresented groups.
Employees need to consider the following as they relate to addressing bias and discrimination in SAP developed AI systems:

- In addition to the conditions laid down in SAP's Global DPP Policy, AI systems shall not be developed or deployed to de-anonymize already anonymized data which may result in the identification of individuals or groups.
- SAP shall endeavor to achieve fairness; AI systems shall not intentionally generate unfairly biased outputs.
- Where relevant, the data used to train AI systems shall be as inclusive as possible, representing as diverse a cross-section of the population or past situations as possible, and as free as possible from (or accounted and mitigated for) any historic or socially constructed biases, inaccuracies, errors, and mistakes.
- SAP shall endeavor to detect unfairly biased outputs and shall implement technical and/or organizational measures to prevent direct or indirect prejudice, discrimination, or marginalization of

groups or individuals, e.g. by reducing bias in training data.

- Wherever possible, developers shall seek to involve impacted/affected users to evaluate and check that outputs are diverse and discrimination free.
- In addition to the conditions laid down in SAP's Global Development Policy and Product Standards, processes shall be put in place to test and monitor for potential biases during the development, deployment, and use phase of AI systems.
    - It shall be trained and tested on as expansive as is feasible, representative, relevant, accurate, and generalizable datasets.
    - The model architectures shall not include target variables, features, processes, or analytical structures which are unreasonable, ethically objectionable, or unable to be validated according to the principles laid out in this document.
    - It shall be developed and deployed so that it has no intentionally harmful impacts on users and/or direct and indirect affected of the system.
    - Where feasible, a fairness function shall be applied to test AI systems for unbiased output.
- AI software shall be user-centric, addressing the widest possible range of applicable end-users, and following relevant accessibility standards, regardless of users' age, gender, abilities, or characteristics.
- As it pertains to the addressing of bias and discrimination, AI systems shall comply with or be in alignment with SAP's Global Development Policy, Product Standards Governance, Product Standards, CUX Design Principles, and Global DPP Policy.
- The use of data for the testing of AI systems shall comply with SAP's Global DPP Policy.

Transparency and Explainability:

SAP's AI systems are held to specific standards in accordance with their level of technical ability and intended usage. Their input, capabilities, intended purpose, and limitations shall be communicated clearly to our customers along with the necessary technical tools for training and prediction.

The problem is AI systems are not morally accountable agents and cannot be held accountable for their actions. It is therefore essential that mechanisms be put in place so that SAP developed AI systems are objective and viable as intended by prioritizing both the transparency of the process by which the AI system is developed, as well as the transparency and interpretability of its decisions and behaviors.

Employees need to consider the following as they relate to the Transparency and Explainability of SAP developed AI systems:

- The data sets and the processes that produce an AI system's decisions, including those of data gathering and data labelling as well as the algorithms used by the developed AI system, shall be documented to allow for traceability and transparency.
- The capabilities and limitations shall be documented as part of the development process in a manner appropriate to the use case at hand. This shall include information regarding the AI system's level of accuracy (performance metric), as well as its limitations and capabilities.
- In alignment and compliance with SAP's Global DPP Policy, products that use AI systems in the processing of personal data must provide transparency to the extent possible as to how the AI system was used in clear and simple language if requested by the data subject.
- In alignment and compliance with SAP's Global DPP Policy, AI systems that engage in profiling or automated decision-making must be able to provide explanations to the extent possible to data subjects upon request, describing the data segment the subject was placed into and the reasons they were placed there. In addition, the reasons as to why the decision was made shall be provided if requested by the data subject. The explanation must be such as to provide the data subject grounds to challenge the decision.
- The methods used for developing, testing and validating, and the outcomes of or decisions made by the AI system shall be fully documented as part of the development process according to SAP's Global Development Policy and Product Development Standards.
- Where applicable, when interacting directly with humans (including via Conversational AI or 'Chatbots'):
    - AI systems shall be made identifiable as such to appropriate end users.
    - AI systems shall be developed such that it does not encourage humans to develop attachment and/or empathy of users towards the AI system.
    - AI systems shall clearly signal to end users that its social interaction is simulated.
- SAP AI system developers shall endeavor to make the decisions, proposals, and outputs of the AI system as transparent as possible, based on the use case. This can be achieved using application logs or the user interface (UI) to allow for the best understanding and traceability of these.

- The user shall be made aware that AI systems typically work on confidence levels; the actual confidence level of a particular output shall be made available if required.
- The AI system's purpose, constraints, requirements, and decisions shall be defined and documented in a clear and transparent manner to the non-technical general reader or user.
- 'Black Box' and/or Deep Neural Network software/ conditions:
    - Where so-called 'black box' algorithms have been developed by SAP, other explicability measures shall be provided. These should include traceability, auditability, and transparent documentation and communication of the software's capabilities.
    - Wherever possible an explanation of the output shall be made available; where not possible, users shall be made aware that the output may not be fully explainable.
    - The requirement for this information will be dependent on the context and the severity of the consequences. A risk matrix approach shall be followed here.
- AI system development shall take into account the context and environment in which the system will operate such that, even with good intentions, no harm or misuse is likely to occur to humans when AI systems are deployed.
- To the extent that a 3rd Party AI system (e.g. 'TensorFlow') is embedded in SAP solutions, this policy shall apply to the overall SAP software solution.

Society:
SAP developed AI systems shall be developed to augment, complement, and empower human cognitive, social, and cultural skills, and not act to prevent or restrict activities commensurate with a free society. Employees need to consider the following when developing or deploying AI systems as it relates to civic society:

- In addition to the existing provisions set out in SAP's Global DPP Policy, AI systems shall not be developed or deployed for human surveillance that is utilized for the targeting of individuals or groups, either by biometrics, facial recognition, or other identifiable features, with the purpose of disregarding or abusing the human rights of the individuals or groups.
- AI systems shall not be developed or deployed for purposes which cause individuals or groups to be discriminated against or excluded from equal access to AI's benefits and opportunities that may be available to the wider population.
- AI systems shall 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 shall not be developed or deployed to undermine human debate or democratic electoral systems.
- AI system development or deployment shall be conducted with minimum to no explicit damage to the environment

7 GOVERNANCE

Use Case Review
As far as SAP developed AI systems are concerned, SAP personnel shall approach ethical dilemmas and trade-offs related to their use via reasoned, context-relevant, and evidence-based decision making rather than intuition or random discretion.

Where a use case proposed for an SAP developed AI system may breach this policy at any stage of the lifecycle of the AI system, or in order to determine whether the application of a specific use case should or should not be pursued, employees should first raise the issue for evaluation by their immediate L1 unit. This even applies if employees only have doubts or concerns.

If questions or concerns remain, or a decision is unable to be made by the appropriate L1 unit, then employees should inform the AI Ethics Office via an e-mail to ai.ethics@sap.com describing the use case.

On receipt of the e-mail the AI Ethics Office will convene the AI Ethics Steering Committee, who will review the use case and advise the employee on how a specific use case may be affected by this policy by assessing each use case submitted to it in the context in which the AI system is to be applied.

Where an employee has a concern that this policy is not being followed for a particular use case, and they are not able to obtain an appropriate response that addresses their concern(s), and wish to do so anonymously, the employee shall use the established whistleblower process to raise this issue to the attention of the AI Ethics Steering Committee.

In the exceptional situation that the AI Ethics Steering Committee cannot come to a decision on a use case query referred to it, the use case may be referred to the Sustainability Council (the 'Council') for final arbitration.

Where a use case decision may have wider implications to SAP as a company, then the recommendation made by either the AI Ethics Steering Committee, or the Council may be referred to the SAP Executive Board for review and ratification.

Where no ethically acceptable trade-offs can be identified, then the development, deployment, use or sale of the AI System shall not proceed in that form.

The Steering Committee will continually review the appropriateness of the ethical decision made to ensure that where necessary suitable changes can be made based on evolving circumstances.

Use case Review Process:

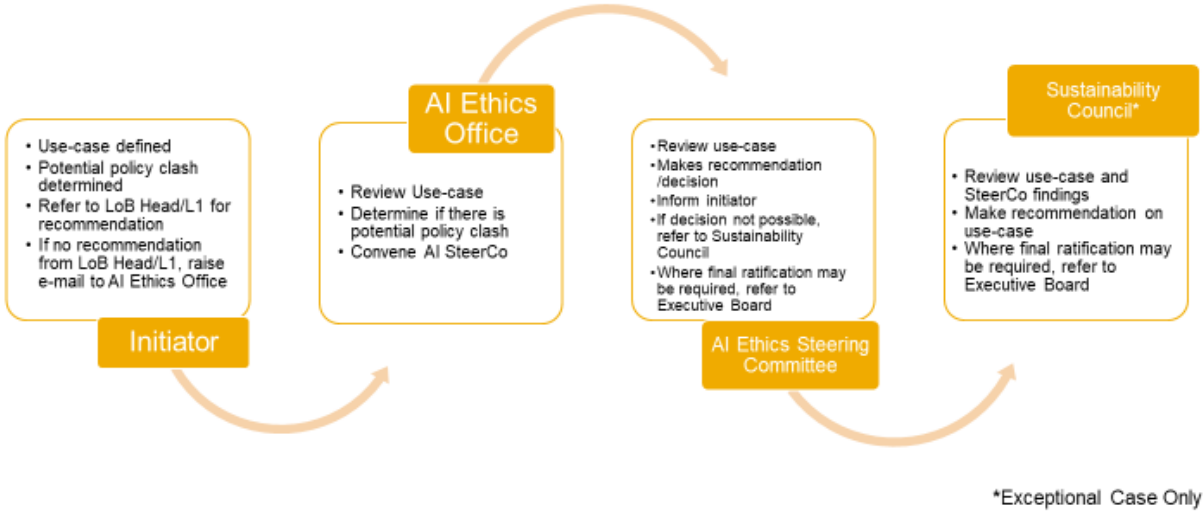

www.sap.com/contactsap

THE BEST RUN