



SAP AI Ethics Handbook

**Applying SAP's Global
AI Ethics Policy across the
Business AI Lifecycle.**



External



SAP.com

Authors

Camila Lombana Diaz
Saskia Welsch

Special Thanks to:

Carolyn Christian, Claus Holzknecht,
Rainer Kerth, Bettina Laugwitz,
Elena Lucarelli, Vikram Nagendra,
Abhishek Pant, Nicole Sanchez Cordero,
Julia Schöndienst, Sebastian Wieczorek.

Updated version: September 2024

Table of contents.

SAP's Definition of AI	0.08
SAP's Ethical AI Guiding Principles	0.12
The AI Ethics Policy	0.18
Partner Ecosystem Implementation	0.20
Red Lines & High-Risk Cases	0.22
SAP's AI Ethics Assessment Process	0.24
A Human-Centered Approach	0.28
SAP's Business AI Lifecycle	0.32
01. Ideation	0.34
02. Validation	0.38
03. Realization	0.42
04. Productization	0.46
05. Operations	0.50

Disclaimer.

This document is intended to serve as a general guide for applying the SAP Global AI Ethics Policy in everyday work and is not meant to advise, implement, or replace either existing or new processes, policies, or procedures. This guide is not intended to be comprehensive and is not a substitute for legal or data protection and privacy advice.

Foreword.

Dear colleagues and stakeholders,

We are thrilled to introduce the AI Ethics Handbook, a comprehensive guide that maps the requirements of our SAP Global AI Ethics Policy to the various stages of the AI development process. This handbook is a vital tool for our commitment to ensuring that our AI solutions are relevant, reliable, and responsible.

At SAP, we recognize the immense potential of AI to drive innovation and transformation across industries. However, we also understand the importance of ensuring that the development and deployment of AI technologies are guided by ethical principles and values. For this reason, our AI Ethics Policy is grounded in the principles laid out in the UNESCO's "Recommendation on the Ethics of Artificial Intelligence" published in 2023.

The AI Ethics Handbook serves as a practical and accessible resource for SAP and partners involved in the development and deployment of AI solutions at SAP. It provides clear guidance on how to implement

Business AI features and solutions that align with our AI Ethics Policy, thereby fostering trust and transparency in our AI technologies.

By following the guidelines outlined in this handbook, we can ensure that our AI solutions uphold the highest standards of ethics, transparency, and accountability. Ultimately, this will enable us to deliver AI technologies that have a positive impact on our users, our customers, and society at large.

All our valued partners and SAP employees are encouraged to familiarize themselves with the contents of this handbook and to integrate these principles into their work. Together, we can leverage the power of AI in a way that is ethical, responsible, and beneficial for all.

Thank you for your dedication and commitment to upholding ethical standards in our AI solutions.

Sincerely,

Walter Sun

SVP, Global Head SAP Business AI

About the AI Ethics Handbook

The first two chapters present SAP's AI Ethics Guiding Principles, recently adopted from the UNESCO's recommendation on the Ethics of Artificial Intelligence; and SAP's Global AI Ethics Policy.

The Red Lines and High-Risk Cases chapter clarifies which AI use cases are prohibited within SAP. It defines high-risk use cases and points out the necessary approval processes. Every use case must complete the AI Use Case Impact Assessment at the beginning of the development process.

SAP's Business AI lifecycle chapter provides an overview of its different stages, including the people involved in the process. Before each stage is introduced in more detail, some general guidelines are described regarding how the AI Ethics Policy is applicable.

You can find a collection of links and sources at the end of the handbook which will guide you through the Business AI lifecycle. The sources are not meant to be binding, but to bring some checkpoints that we consider helpful depending on your use case.

SAP personnel shall approach ethical dilemmas and trade-offs related to their use case through reasoned, context-relevant, and evidence-based decision-making rather than intuition or random discretion.

SAP employees are encouraged to raise any issues related to possible policy breaches for evaluation by their immediate L1 unit, even if employees only have doubts or concerns.

If an end user or affected individual either at SAP or outside the company has concerns about an AI system adherence to this policy and wishes to raise these concerns anonymously, they can use the [SpeakOut](#) tool.

If questions or concerns remain, then employees should inform the AI Ethics Office describing the use case.

SAP's Definition of AI.

AI is typically defined as the ability of a machine to perform cognitive functions that we associate with the human mind, such as perceiving, reasoning, learning, and solving problems.

It requires a system to correctly interpret external data, learn from such data, and use that learning to achieve specific goals through flexible adaptation.

In general, an AI system is a machine-based system that, for explicit or implicit goals, infers from the input it receives how to generate outputs. Such outputs like predictions, content, recommendations, or decisions may affect physical or virtual environments. Different AI Systems vary in their levels of autonomy and adaptiveness after deployment.

For this policy, an AI system shall be deemed to be:

- AI systems are developed by SAP, whether commercialized or not.
- An AI system developed by an SAP supplier and embedded in a solution.
- AI systems developed by an SAP supplier and sold by SAP, such as supplier-branded or SAP-branded resellers.
- An AI system developed by an SAP supplier and used by SAP for internal purposes.

At SAP, the AI Ethics Assessment Process shall be performed on all AI systems during the development and subsequent phases to ensure no unintended biases.

SAP's Differentiation of AI Systems:

SAP differentiates three types of AI system: Symbolic AI, Traditional/narrow AI, and Generative AI.

01. Symbolic AI

Symbolic AI is based on a knowledge database with explicitly programmed if-then rules to process information and make decisions. Human experts define these rules, so symbolic AI is often called rule-based or expert AI.

Rule-Based AI

Rule-based AI systems are characterized by their behavior being fully defined by rules created by human experts. These systems are often described as symbolic or expert systems.

A traditional application domain of such AI systems is supporting diagnostic and therapeutic decisions in medical contexts. The advantage of symbolic AI is that its decisions are comprehensible due to the explicit programming of rules and effective for tasks with clear, structured rules and objectives.

However, it may struggle with more complex, ambiguous situations that do not fit neatly into predefined categories.

Another way to understand symbolic AI is to imagine playing chess with a computer. The computer knows all the rules, can predict your moves, and makes decisions based on a predefined strategy. However, it is not inventing new ways to play chess but rather selecting from the strategies it was programmed with.

02. Traditional/ Narrow AI

Traditional AI, often called Narrow or Weak AI, focuses on performing a specific task intelligently. It refers to systems designed to respond to a controlled set of inputs. It focuses on a single area and outperforms in that particular task. These systems can learn from data and make decisions or predictions based on data.

It is effective for tasks with clear, structured rules and objectives, but it may struggle with more complex, ambiguous situations that do not fit neatly into predefined categories. Notable examples of Traditional AI include voice assistants such as Siri and Alexa, recommendation engines on Netflix or Amazon, or Google's search algorithm. These AIs have been trained to follow specific rules, do a particular job, and do it well, but they do not create anything new.

Learning-Based AI

Learning-based AI systems differentiate themselves by the fact that humans define the problem and the goal, but the behavior, rules, and relationships required for the system are learned in an automated way. With the help of data, AI systems learn how to solve a problem and continuously adapt their function in this process.

There are 3 categories of learning:

• Supervised Learning:

This is used when the task of the model to be trained is already known, e.g., predicting house prices based on location, square footage, and number of rooms. Based on historical data on house prices, the algorithm attempts to find the statistical correlation between input and output and constantly adjusts its parameters to maximize the accuracy of its predictions.

• Unsupervised Learning:

Algorithms are given an unlabeled data set and independently recognize correlations and patterns by grouping data points with similar properties. The developer does not need to know these patterns in advance.

• Reinforcement Learning:

Algorithms attempt to solve a given problem through constant trial and error. The algorithm is "rewarded" for good solution approaches, while bad approaches result in "punishment." By striving to maximize the reward, the algorithm constantly corrects itself and thus optimizes its strategy.

03. Generative AI

Generative AI refers to a category of AI algorithms that can create images, text, sounds, or videos that resemble human-made content.

Unlike Narrow AI, generative AI can produce new data or patterns based on the learned information. This is done by training on a large dataset of examples, allowing the AI to understand the underlying structure of the data. Once trained, it can produce original outputs that mimic the training data's style, tone, or format.

Generative AI uses deep learning models to understand input data and create new content.

What distinguishes traditional and generative AI from rule-based AI systems is that they are not interpretable. They are opaque systems with millions of billions of model parameters and dozens or hundreds of stages of inference, and it is still difficult to discern what relationships they have learned or how they reach a decision.

SAP's Ethical AI Guiding Principles.

The following principles are consistent with the AI Ethics Policy and are based on the UNESCO Recommendation on the Ethics of Artificial Intelligence.

Principles 01 to 07 apply to all employees involved in the creation of AI systems at SAP. Related requirements are described in the SAP Global AI Ethics Policy and need to be fulfilled.

All SAP employees should be aware of the company's commitment to ethical, responsible, and trustworthy AI.

01.

Proportionality and do not harm.

- SAP defines a clear scope for the AI system and ensure that the AI system's behavior matches this scope
- SAP remains true to its human rights commitment statement, the UN guiding principles on business and human rights, laws, and widely accepted international norms
- SAP puts measures in place to help prevent any inappropriate use

02.

Safety and Security.

- AI systems go through SAP's regular, continuously adapted quality assurance process
- AI systems undergo thorough testing under real world scenarios to firmly validate that they meet their purpose and product specifications
- SAP works closely with its customers and users to uphold and further improve its AI systems' quality, safety, reliability, and security

03.

Fairness and non-discrimination.

- *Bias in AI systems can harm individuals and SAP's customers by perpetuating societal patterns of marginalization, inequality, and discrimination*
- Throughout the AI system lifecycle, SAP assesses and monitors fairness by ensuring representation of all forms of diversity and making reasonable efforts to avoid discriminatory or biased outcomes
- Check SAP's mission on [business beyond bias](#).

04.

Sustainability.

- SAP strives to create AI systems that foster positive environmental, economic, and social impact
- At every stage of the AI system lifecycle, SAP ensures that its sustainability policies are being followed

05.

Right to privacy and data protection.

- Data protection and privacy are at the core of any of SAP's products and services. SAP communicates clearly how, why, where, and when customer and anonymized user data is used in its AI systems
- This commitment to data protection and privacy is reflected in SAP's compliance with all applicable regulatory requirements
- In addition, it is reflected through the research SAP conducts in partnership with leading academic institutions to develop the next generation of privacy-enhancing methodologies and technologies

06.

Human oversight and determination.

- SAP follows a human-centered approach to Business AI
- SAP puts in place appropriate governance mechanisms to allow for human intervention and control of its AI systems' decisions
- The extend of oversight necessary for a specific AI system is evaluated depending on its context of use and risks of interference with the rights and freedom of affected individuals

07.

Transparency and explainability.

- SAP clearly communicates input, capabilities, intended purpose, and limitations of its AI systems to its customers
- SAP ensures that end users are aware about their interaction with an AI system
- SAP strives to provide clear and understandable explanations of recommendations and decisions of its AI systems
- SAP actively supports industry collaboration and conducts research to improve transparency

The last three principles

Principles 08 through 10 refer to organizational requirements to provide a clear governance structure for creating and documenting AI systems that comply with AI ethics requirements.

In addition, the organization must ensure that all stakeholders are aware of the requirements and have the knowledge to meet them.

08.

Responsibility and accountability.

- AI systems are no moral agents and cannot be held accountable for their actions. Ethical and legal responsibility must always be attributable to human actors and existing legal entities
- SAP established governance mechanisms like impact and compliance assessments to ensure that AI systems are developed, deployed, sold, and used ethically and responsibly

09.

Awareness and literacy.

- SAP provides courses and learning material accessible for everyone on the topic of AI and its responsible use
- SAP sponsors and organizes educational events and encourages exchange on the topic
- SAP commits itself to promote global initiatives on digital literacy and access

10.

Multistakeholder and adaptive governance and collaboration.

- The development of ethical AI requires a broader discourse across industries, disciplines, borders, and cultural, philosophical, and religious traditions
- SAP strives to become an active voice in these debates by engaging with its AI Ethics Advisory Panel and a wide range of partnerships and initiatives

The AI Ethics Policy.

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

SAP's commitment to the Global Human Rights Commitment Statement to uphold and support the Universal Declaration of Human Rights is foundational to SAP's approach to AI Ethics.

SAP has based this updated policy version on [UNESCO's Recommendation on the Ethics of Artificial Intelligence](#) to address the challenges posed by rapidly evolving AI technologies and align with an internationally recognized set of values.

[The Global AI Ethics Policy](#) clarifies how principles based on UNESCO's Recommendation on the Ethics of AI relate to AI use cases. The policy applies to SAP and all its employees worldwide. It defines the intent, expectations, and obligations of employees involved in the design, development, deployment, and sale of AI systems.

Each actor involved in any phase of the AI system lifecycle must ensure that AI systems are designed, developed, sold, and used in accordance with the ethical principles outlined in this policy and that these ethical principles are maintained and followed throughout the AI system lifecycle. Adherence to these ethical principles is especially important so customers and the public can trust SAP's ability to develop, deploy, use, and sell AI systems responsibly and ethically.

Partner Ecosystem implemen- tation

At SAP, we believe that together with our partners and customers, we form a chain of responsibility, vital in ensuring the delivery of ethical and high-quality products and services. Upholding Responsible AI practices throughout our partner ecosystem is not only essential for maintaining customer confidence and our collective reputation but is also a fundamental aspect of our shared success.

Putting the policy into action from a partner perspective, SAP has introduced a certification program for partner applications developed on [SAP Business Technology Platform \(BTP\)](#) using [SAP Generative AI Hub](#) that includes checks for Responsible AI compliance. The [certification program](#) enables partners to offer trusted, compliant, and enterprise-ready applications powered by AI services, leveraging SAP's expertise in business data insights.

Furthermore, solutions created by our partners and resold by SAP undergo additional rigorous checks. This ensures alignment with the same stringent Responsible AI product standards we uphold throughout our own development lifecycle. SAP remains steadfast in its dedication to integrating Responsible AI practices into every facet of our business.

We firmly believe that this ethical foundation is paramount for fostering strong, trust-based relationships with our customers. We invite our partners to join us in upholding these principles and to continue the dialogue about Responsible AI. Together, we can cultivate an environment built on trust and integrity that our customers will benefit from.

SAP's Red Line Cases

No SAP AI use cases should fall under the so-called red lines defined in the SAP Global AI Ethics Policy. These are defined as highly unethical under the purposes for AI.

If your use case is built for these purposes, you must immediately stop developing, deploying, and selling that use case.

SAP's red lines are outlined on the right.

01. Personal Freedom

Human Surveillance

Human surveillance that is utilized for targeting individuals or groups, either by biometrics, facial recognition, or other identifiable features, with the purpose of disregarding or abusing the rights of the individuals or groups.

Discrimination

Purposes which cause individuals or groups to be discriminated against or excluded from equal access to AI's benefits and opportunities to the wider population.

De-anonymization

De-anonymization should never be the goal of any AI use case.

02. Society

Manipulation

Deception or unfair manipulation of individuals or groups via public forums, media, or moderation of other similar uses.

Undermine Debate

Systems which undermine human debate or democratic electoral systems.

Harmful Impact

Intentionally harmful impacts on users and/or those directly and indirectly affected by the system.

Social scoring

Evaluation or classification of the trustworthiness of individuals based on their social behavior or known/predicted personal or personality characteristics.

03. Environment

Environmental Harm

AI system development or deployment shall be conducted with minimum to no explicit damage to the environment.

SAP's High-Risk Cases

SAP classifies AI use cases as high-risk under certain circumstances. We base these criteria on what kinds of AI have negatively affected individuals or whole populations.

High-risk use cases are not prohibited within SAP; however, they must first undergo an assessment process by the AI Ethics Steering Committee before they can be further developed, deployed, and sold.

SAP understands that AI use cases may have high-risk implications, such as potential discrimination stemming from personal data or physical harm resulting from embodied automation. The following assessment covers both implications.

Please note that further changes may be necessary after the AI Ethics Steering Committee's assessment process.

Determining High Risk Cases

Processing personal data:

Does the use case process any information relating to an identified or identifiable natural person for training purposes or during productive usage? – *Use cases with anonymized data sets or only the process of anonymizing personal data does not qualify as high-risk case.*

Processing sensitive personal data:

Does the use case include processing sensitive personal data like information on sexual orientation, religion, and biometric data (including face imaging and/or voice recognition)?

Automated decision making:

Does the use case exhibit fully or partially automated decision-making?

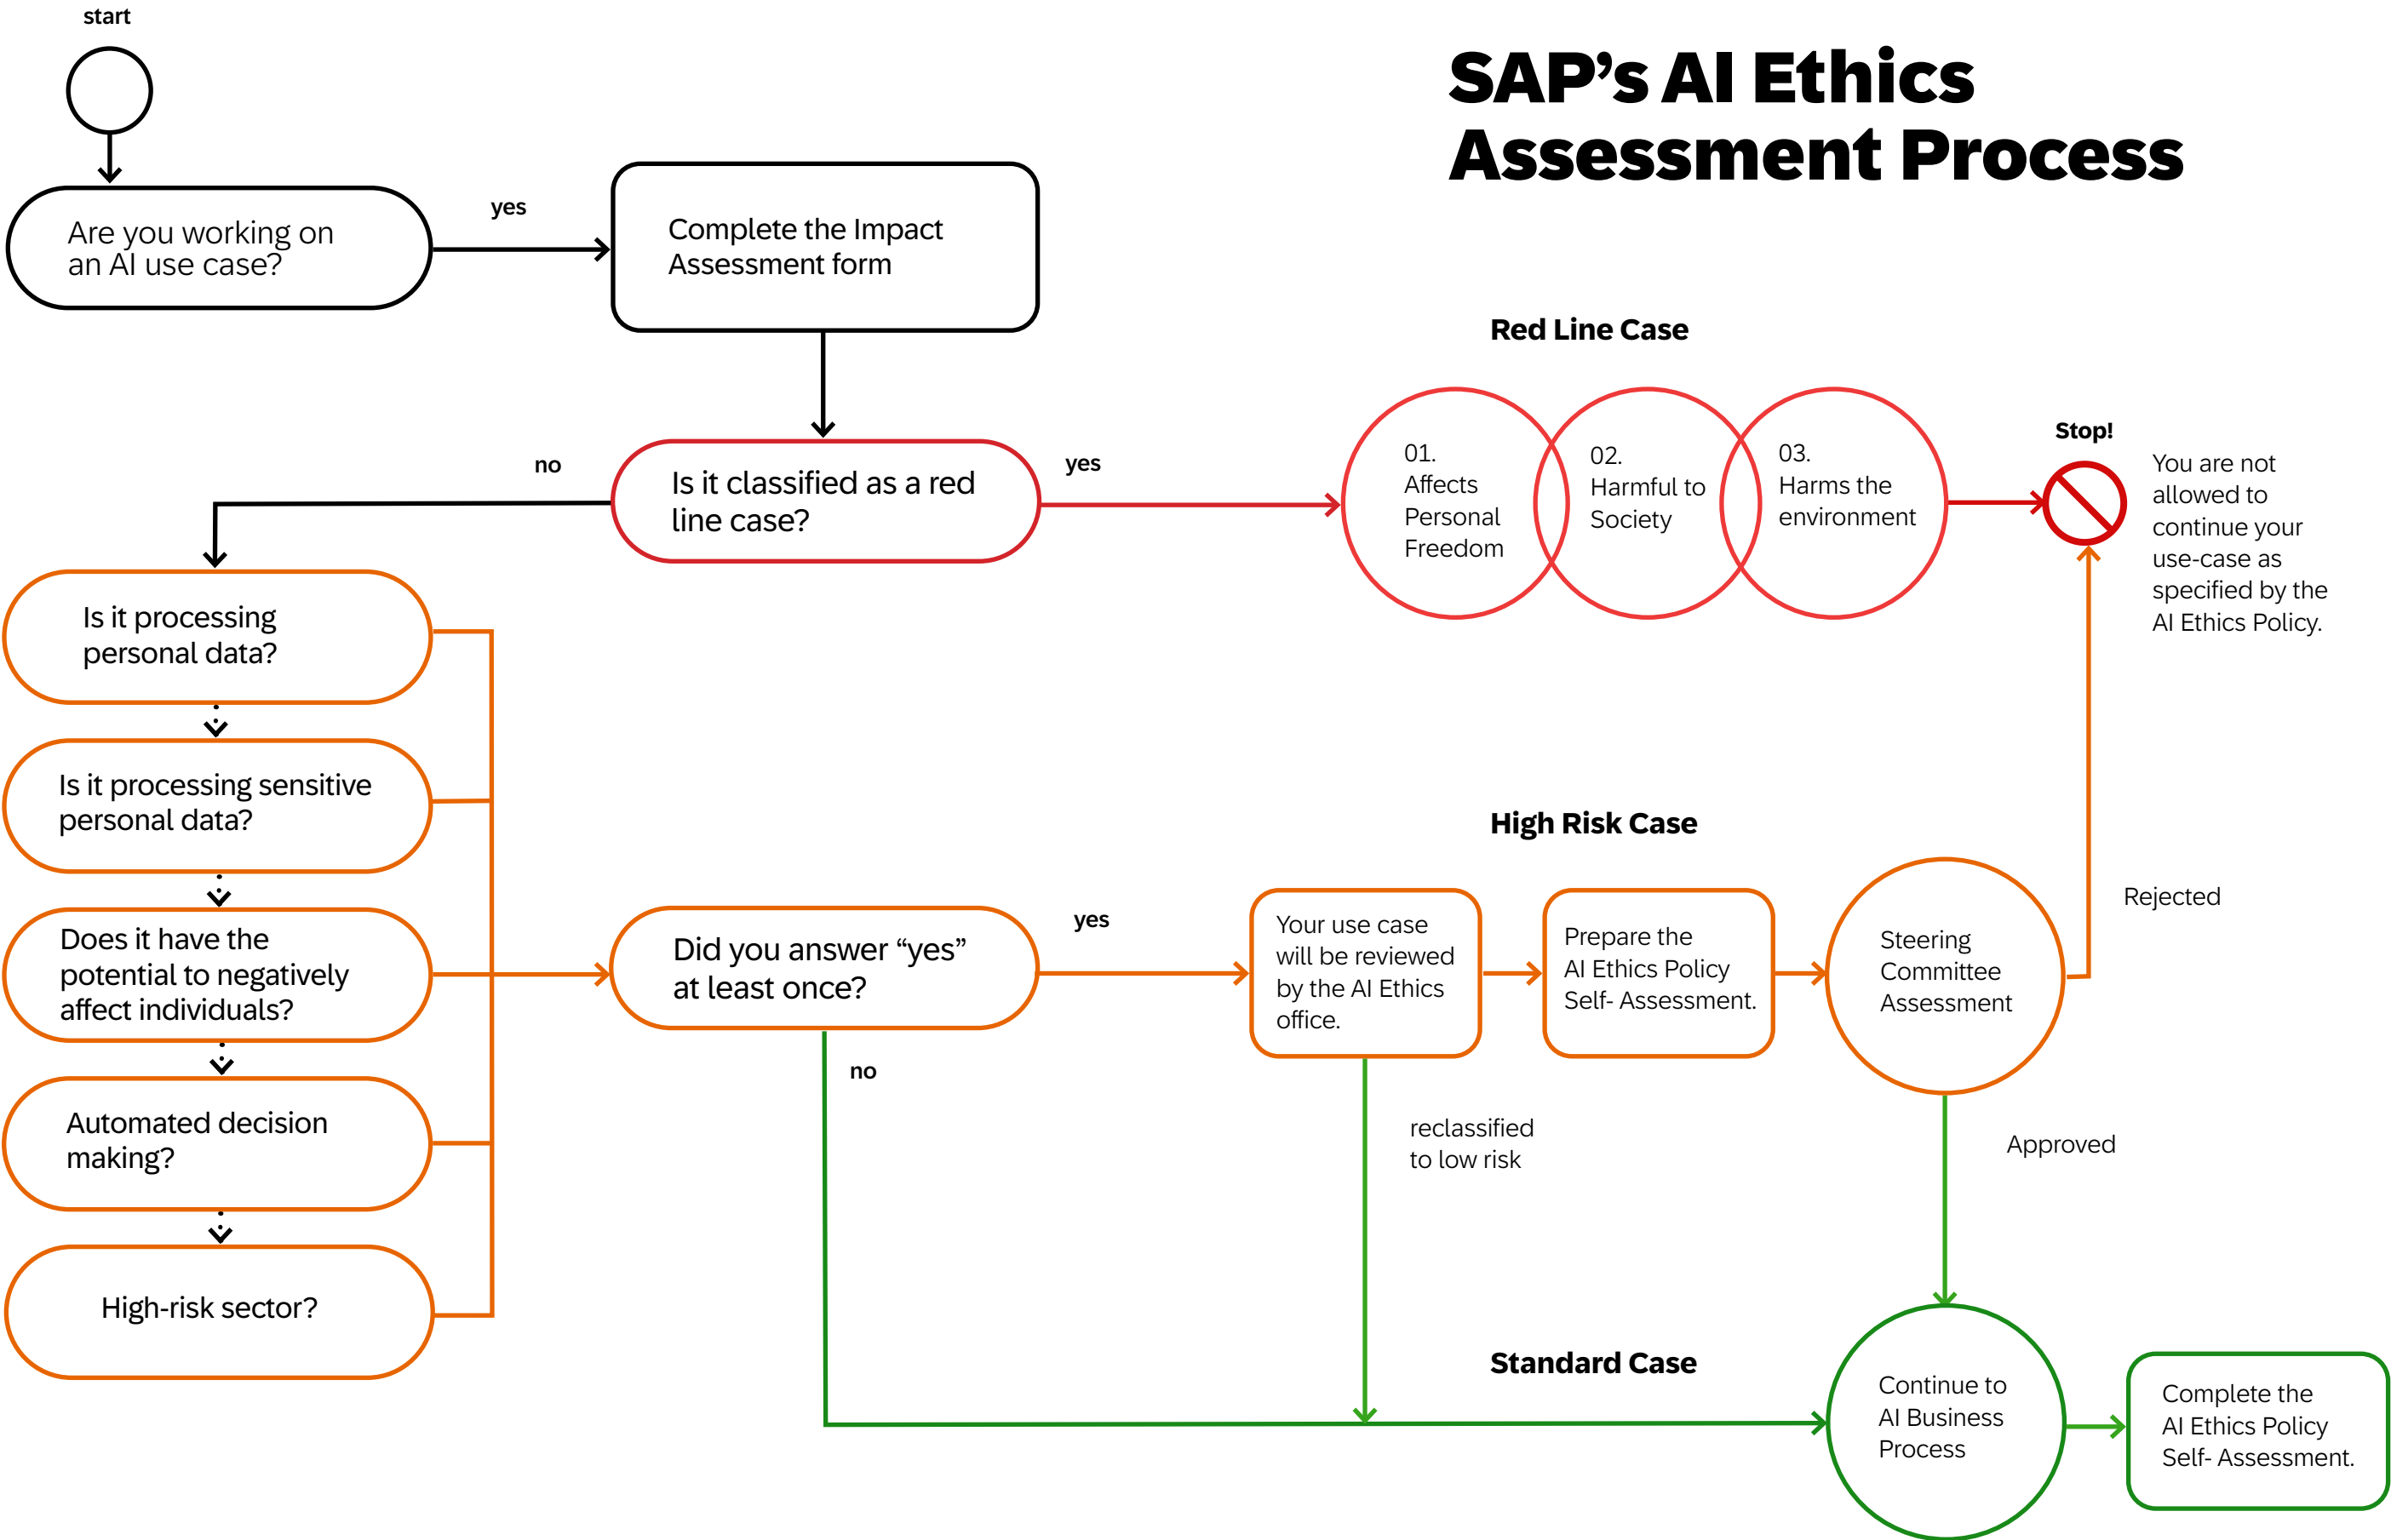
Does it include cases where no human intervention or human supervision takes place? – *This excludes any recommender system.*

Negatively affected individuals:

Could the use case negatively affect the well-being of individuals or groups (social, safety, financial, and/or physical harm) or intrude/restrict an individual's fundamental rights and freedom, if the AI system is insufficiently trained and tested, incorrectly used by customers or used for a different purpose than the original scope?

High-risk application:

Does the use case belong to one of the following high-risk sectors? Categorization of natural persons, management and operation of critical infrastructure, employment/HR, healthcare, private services and public services and benefits, law enforcement, migration, or democratic processes?



A Human Centered Approach

People Involved



At SAP, emphasizing Human-Centered AI is crucial for fostering beneficial and sustainable AI processes and services. Enhancing our understanding of this concept will magnify the benefits for AI end users and society within business contexts while mitigating fears and existential threats our users face.

Consequently, our AI system lifecycle must be guided by a profound comprehension of our users' and stakeholders' needs. Employing methodologies such as user research, design thinking, and accurately defining user stories are essential to our development process.


This approach ensures that our products closely align with the genuine needs of our target groups, thereby reducing unconscious bias and minimizing the risk of unintended harm.

People Involved

AI Use Case Owner.

- Responsible for coordinating the delivery of an AI-enabled product or feature.
- Their responsibility is found end-to-end across all phases of the Business AI lifecycle.


Process Activity:



AI Data Scientist.

- Develops, implements, or applies AI methods to derive solutions to business problems that can be translated into AI functions.

Process Activity:



AI Engineer.

- Responsible for the design, implementation, and maintenance of product AI functions.
- Works according to specifications and project plans.

Process Activity:



User Assistance Developer.

- Responsible for designing, developing, reviewing, and maintaining content for product documentation, user interface messages, and conversational applications.
- Collaborates with internal stakeholders to design customer-facing content and implements embedded help delivery mechanisms that improve the delivery of technical information to customers.

Process Activity:



Designer.

- Provides the Human-Centered understanding of the product and/or use case.
- Develops a user-centered strategy that includes all stakeholders, affected end users, and product roadmap.
- Is obliged to conduct user research for possibly affected users.
- Creates flows and conducts usability evaluation (including voice response and conversational user interaction).
- Applies Fiori-based AI patterns.


Process Activity:



Product Developer.

- Develops the business logic of a product, or at least parts of it. This task includes the integration of the provided AI functions into the product context.

Process Activity:



AI/LLM OPS Engineer.

- Operates the AI-enabled products or individual AI functions.
- Manages lifecycle aspects like version updates.

Process Activity:



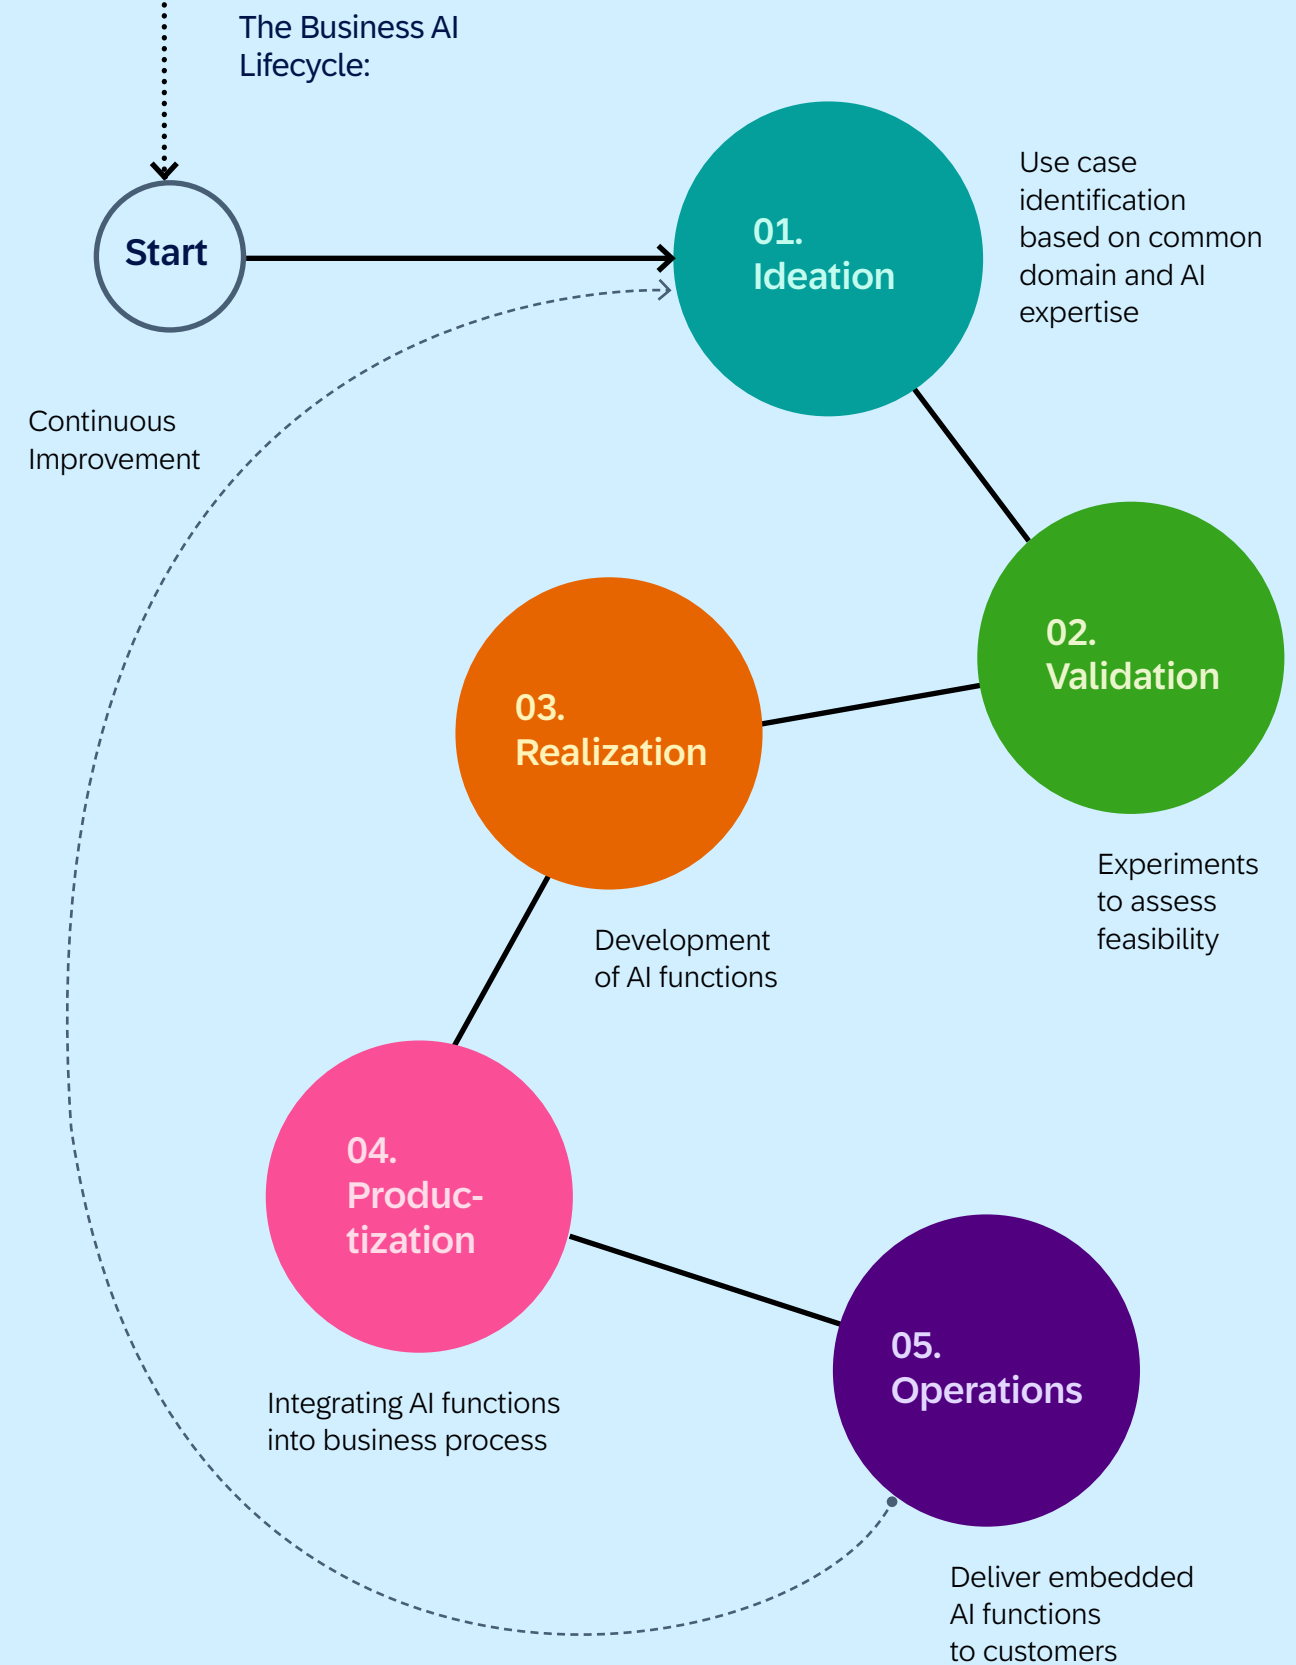
Customer.

- Connects the development team with the potential end users for user research, insights, and data.

Process Activity:




SAP's Business AI Lifecycle



01 Ideation

People Involved



AI Use Case Owner



AI Data Scientist



Designer



Customer

Example: Versatile Writing Assistant

The SAP SuccessFactors team aims to use generative AI to help users communicate more clearly, efficiently, and effectively. Their project, the Versatile Writing Assistant, is designed to refine, condense, and translate user texts and interpret and summarize system text from various sources.

The team ensures that their project aligns with ethical guidelines by referring to the Global AI Ethics Policy. They check for any violation of SAP’s Red Lines under sections **1.1 ‘Personal freedom’** and **1.2 ‘Society’** of the policy. They also discuss SAP’s Ethical AI Guiding Principles, being aware that each principle can become relevant in the development of the use case. They identify two principles that are particularly important to consider in the ideation phase, namely **Principle 1 (Proportionality and do no harm)** and **Principle 6 (Human Oversight and Determination)**.

To fulfill **Principle 1 (Proportionality and do no harm)**, they plan to predefine which instructions can

be given to the underlying AI model and implement them through the user interface, rather than permitting direct user interaction with the AI model.

Following **Principle 6 (Human Oversight and Determination)**, the team opts for a human-in-the-loop model to allow users to review and change AI-generated texts.

Before moving on by developing a minimum viable product, the team consults their data protection officer to see whether the use case is feasible from a data protection and privacy perspective.

The project owner also completes the AI use case impact assessment questionnaire to evaluate ethical risks. Given that the project is intended for HR (Human Resources) solutions, it receives a high-risk classification and must be reviewed by the AI Ethics Steering Committee. With help from the AI Ethics Ops team, the SAP SuccessFactors team is preparing for this review.

Ideation Checklist.

The ideation process in the AI lifecycle refers to the initial phase where ideas for AI solutions are generated, evaluated, and refined before moving on to development and implementation. This phase is crucial as it sets the foundation for the entire AI project.

When considering the ideation phase from an ethical perspective, the focus shifts to ensuring that the AI solutions being conceived and developed align with ethical principles, safeguard against potential harms, and promote fairness, accountability, and transparency.

AI-driven solutions are not just about meeting compliance standards; they should aim to contribute to a world where everyone feels empowered and included. Read more below to ensure bias is limited and that the use case follows a human-centric design:

01: Understanding of AI Ethics Policy, guidelines, requirements and processes.

- Consult SAP's Ethical AI Guiding Principles and the AI Ethics Policy to become familiar with the company's red lines.
- Align with the SAP Global Data Protection and Privacy Policy and SAP Data Protection and Privacy Guidance documents.
- Suppose humans could be directly impacted by a decision made by SAP's AI system. In that case, human oversight would safeguard the AI system from undermining human autonomy or introducing unintended consequences.
- The project owner should complete the AI Use Case Impact Assessment Form to evaluate ethical risks. If it receives a high-risk classification, wait until the AI Ethics Steering Committee reviews it. Contact the AI Ethics Ops team to prepare for this review.

02: Understanding of the people who may be affected by the use case.

- Who are the people involved in the use case, and how might they be affected if applicable?
- Have a user-centric understanding of the possible risks and represent them in a personal story.

03: Define the degrees of freedom, fairness, transparency & explainability.


- Define a degree of freedom of an AI system and choose an appropriate governance mechanism; consider consulting the resources below:
 - **SAP Product Standard** – Ensure automated decision processes can be operated compliantly.
 - **SAP DPP Guideline AI** – Automated Decision Making
 - **NIST AI RMF - MAP 3.5** – Processes for human oversight are defined, assessed, and documented in accordance with organizational policies from GOVERN function.
- Define what fairness is for your use case. Consider the resources below:
 - **Fairlearn: A toolkit for assessing and improving fairness in AI.**
 - **User Guide** — Fairlearn 0.11.0.dev0 documentation.
- Decide on the degree of transparency and explainability per **Principle 7** of the AI Ethics Policy.

04: Available tools for ideation.

- AppHaus: [Generative AI Explore Workshop](#).
- Fiori for Intelligent Systems.

02 Validation

People Involved



AI Use Case
Owner



AI Data
Scientist



Designer



Customer

Example: Versatile Writing Assistant

The SAP SuccessFactors team moves forward with the Versatile Writing Assistant use case considering the instructions of the AI Ethics Steering Committee.

Although the Steering Committee had lowered the ranking of the use case to a standard (i.e., low-risk) use case, it suggested implementing additional transparency measures to make users aware of potential shortcomings.

Despite the project’s low-risk status, the group is committed to complying with the AI Ethics Policy, including, but not limited to, **Principle 7 (Transparency & Explainability)**. The team has decided on the degree of transparency and explainability they want to provide for the writing assistant’s suggestions. They’ve labeled the AI functionali-

ties in the user interface as ‘Use AI Capabilities’ to clarify to the user that they’re interacting with an AI system.

In line with **Principle 3 (Fairness and Non-Discrimination)**, the team is exploring ways to prevent the writing assistant from making discriminatory, inappropriate, or offensive suggestions. They’re testing the effectiveness of security filters that can review suggestions made by the AI and block any inappropriate ones before they reach the user.

Validation Checklist.

Validation involves assessing the performance of your AI model to ensure it meets your project’s objectives and requirements. This phase involves assessing systems or processes to ensure they meet ethical standards, align with user needs, respect cultural diversity, and evaluate a positive long-term impact.

01: Make sure your MVP (Minimun Viable Product) is compliant:

- Align with SAP DPP.
- Get familiar with the SAP Development schedule (if applicable).

02: Ensure there is a user validation.

- Always validate with users who may be affected by considering fairness, inclusion, and cultural sensitivity. Learn more on including accessibility and inclusivity measures and how AI should be used to expand, and not restrict.
Consider the following tools:
 - [User Guide — Fairlearn 0.11.0.dev0 documentation.](#)
 - [SAP's Accessibility standards for product design.](#)
 - [Fairlearn: A toolkit for assessing and improving fairness in AI.](#)

03: Validate your use case’s data & models.

- If your use case uses GenAI: LLM model selection, integrate content filtration.
- If your use case is based on Symbolic and/or Narrow AI: Find and validate the data for the use case.
- Understand security and privacy implications.
 - Is your data being stored securely?
 - Is your data being shared securely?
- Understand role-based application integration.
 - Have you identified the user roles in the system?

04: Fill out AI Ethics Policy Self-Assessment questionnaire.

05: Validate section 6.2 of the [UNESCO Ethical impact assessment tool kit](#).


- Has the algorithm been tested with different groups?
- Does the design allow all people, especially marginalized groups, to access and interact with the AI system?
- To which segment of the population will the AI system be applied?
Is the population affected particularly marginalized?

06: Other considerations and available tools for Validation.

- High-risk cases: must be evaluated by the Steering Committee.
- Partner applications: partners must adhere to our guidelines.
- [7010-2020 IEEE Recommended practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being | IEEE Standard | IEEE Explore.](#)
- [Deon: An Ethics Checklist for Data Scientists - DrivenData Labs.](#)

03 Realization

People Involved



AI Use Case
Owner



AI Data
Scientist



Designer



AI
Engineer



User Assistance
Developer

Example: Versatile Writing Assistant

The SAP SuccessFactors team has successfully built a prototype of their Versatile Writing Assistant and is now entering the realization stage.

While developing the central AI functionalities, the team also implements the ethical requirements defined in the ideation and validation phases.

Additionally, the AI engineer starts documenting all information required by **Principle 7 (Transparency & Explainability)**, as detailed in 7.2: ‘Documentation.’ This includes, but is not limited to, the capabilities, intended use, and limitations of the AI functionality, as well as the name, type, and workings of the foundation model being used.

The use case owner starts to fill out the AI Policy self-assessment questionnaire to ensure that the writing assistant adheres to the AI Ethics Policy when it’s delivered with SAP SuccessFactors.

Even if not all questions can be answered at this stage of development, it encourages the team to think again about SAP’s Ethical AI Guiding Principles. They reassess whether the use case still complies with all Guiding Principles.

Realization Checklist.

The realization step typically refers to the implementation phase, where the designed system or algorithm is built or expressed into a tangible form, such as software. It's the stage where the theoretical concepts, algorithms, and models are translated into executable code, physical systems, or applications that can perform the intended tasks.

01: Integrate fairness and explainability.

- [Explainable AI | SAP Fiori for Web Design Guidelines.](#)
- Design interfaces with Fiori AI Controls for Transparency.

02: Consider metrics that are relevant to your use case.

- [OECD Catalogues of tools and metrics for Trustworthy AI](#)

03: Create use case documentation & finalize the AI Ethics Policy Self-Assessment questionnaire.

- Create documentation on capabilities and limitations
- Please fill out the AI Ethics Policy Self-Assessment questionnaire to consider requirements they may not have considered. Finish writing the AI Ethics Policy Self-Assessment questionnaire and upload it to SIRIUS, the internal quality management tool.

04: Apply anonymization.


- [SAP HANA Data Anonymization Service](#)

05: Other considerations and available tools for Validation.

- Consider having data cards/model cards (if applicable)
- [GenAI use cases: utilize the SAP Generative AI Hub](#)
- [Fairness Indicators](#) | [Responsible AI Toolkit](#) | [TensorFlow](#)
- [Fairlearn: A toolkit for assessing and improving fairness in AI: User Guide — Fairlearn 0.11.0.dev0 documentation](#)

04 Producti- zation

People Involved



AI Use Case
Owner



AI Engineer



Developer



AI/LLM OPS
Engineer



User Assistance
Developer



Designer

Example: Versatile Writing Assistant

The developed AI capabilities are now ready to be embedded into SAP SuccessFactors. At this step, the Steering Committee’s recommendations can be implemented. While making sure that all of SAP’s Ethical AI Guiding Principles are followed, they agree that they must pay particular attention to the requirements laid down in **Principle 7 (Transparency and explainability)**. Developers, designers, and user assistance developers collectively brainstorm measures to enhance system transparency.

Firstly, they decided to insert a label into the user interface (UI). This label reminds users to always critically evaluate the suggestions made by the writing assistant, as there may be incorrect or inaccurate phrasing.

Secondly, they incorporate a feedback control mechanism into the product. This feature enables users to give feedback regarding the quality of the AI recommendations. This feedback can address

any problematic suggestions the writing assistant produces during operational use. Such a continuous feedback loop also helps in improving the system over time.

Then, the team refines and supplements the existing product documentation. The user assistance developer ensures that the descriptions of the AI capabilities, data usage, and the underpinning algorithm are easy to understand for non-technical readers.

Before the Versatile Writing Assistant is shipped, the use case owner completes and uploads the AI Policy Self-Assessment questionnaire to SIRIUS, the internal quality management tool.

This refers to **Principle 8 (Responsibility and Accountability)**, which clarifies the role of individual and SAP accountability before production. The functionality cannot be shipped without fulfilling this prerequisite.

Productization Checklist.

Productization involves turning a successful prototype or model into a deployable, scalable, and market-ready product or service. It includes refining the solution, optimizing its performance, integrating it with existing systems, ensuring scalability, reliability, and usability, and often involves considerations like pricing and support.

01: Ensure transparency requirements are met.

- Include documentation requirements for customers. Refer to Section 7.2 of the SAP Global AI Ethics Policy on 'Documentation' to comply with SAP's transparency principles.
- Document all information required. This includes, but is not limited to, the capabilities, intended use, and limitations of the AI functionality as well as the name, type, and workings of the foundation model that is being leveraged.
- The user assistance developer must enhance transparency documentation and create guidance on technology risks.

02: Validate the use case through Prompt Grounding.


- Consider using SAP Generative AI Hub (Only applicable for GenAI)

03: Security checks to have into consideration.

- **System Security:** API Security, monitor anomalies, patch management, and Isolation.
- **Deployment Security:** Secure pipelines, environment segmentation, dependency Management.
- **Data Security:** Limit data access to authorized personnel, encrypt data in transit and at rest.
- **Model Security:** Protect models from tampering, implement defenses against adversarial attacks.
- **Incident Response:** Set up incident detection and a disaster recovery plan.

05 Operations

People Involved



**AI Use Case
Owner**



**AI/LLM OPS
Engineer**



**User Assistance
Developer**



Customer

Example: **Versatile Writing Assistant**

When the embedded AI function is deployed, the AI/LLM Ops engineer monitors the AI capabilities and ensures they work as intended. Furthermore, the AI/LLM Ops engineer handles the writing assistant's potential inaccurate, offensive, or objectionable suggestions and passes them on to the use case owner, if necessary.

One year after the successful product launch, the SAP SuccessFactors team has decided to enrich the writing assistant with new functionalities. Because the scope has changed significantly, they trigger a new impact assessment and self-assessment process in line with **Principle 8 (Responsibility and Accountability)**, which should be a cyclical revaluation by release. In the course of this, they also review compliance with the other guiding principles, e.g., **Principle 6 (Human**

Oversight and Determination), which ensures that humans are always in the loop.

Looking at **Principle 7 (Transparency and explainability)**, the team noticed that new foundation models have been developed to provide reasonable explanations for their output. The data scientists, AI engineers, and designers work on a proof of concept to investigate whether such foundation models can be leveraged for their use case and whether the increased explainability adds value to the use case.

Operations Checklist.

Operations is about managing and optimizing deployed AI systems. It involves monitoring performance, updating models, addressing issues, and ensuring continued effectiveness and compliance. Additionally, it may involve user support, resource allocation, and compliance with regulations or ethical guidelines.

Essentially, it's about ensuring the AI system remains operational and beneficial over time.

01: Insure AI Ethics governance checks are done.

- AI Ethics Policy Self-Assessment updates in case of feature changes.
- SIRIUS checks must be done by release. Iterative approach:
 - Changing the scope or launching new versions of a use case triggers a new impact assessment and self-assessment process.
 - Note: this may change the risk level of the use case depending on ideation.

02: Leverage post-deployment monitoring tools.

- Collect user feedback.
- Observe any changes in the code.
- Check for fairness measures.
- Define tokenization aspects (only related to GenAI)

03: Guarantee ongoing compliance.

- Potential user reports of inaccurate, offensive, or otherwise objectionable suggestions.
- Pass on to the use-case owner.

04: Ensure AI/LLM Ops compliance.

- Setting up an operational process to apply software patches for correction of algorithms on priority.
- Provide documented tests to demonstrate that the algorithms do not include discriminatory effects.
- Logging and monitoring is continuous throughout Secure SDOL (security and performance logs), applicable not only to production systems but also on dev and test systems.
- As part of governance monitoring (compliance), the compliance must be maintained throughout to ensure smooth release cycle and audits (internal or external).
- Analyze data distribution at a regular cadence to detect data poisoning, in case of ML/LLM models.

© 2024 SAP SE or an SAP affiliate company.

All rights reserved.

See Legal Notice on www.sap.com/legal-notice
for use terms, disclaimers, disclosures, or
restrictions related to this material.

SAP
AI ETHICS