



Adobe's Commitment to AI Ethics

At Adobe, our purpose is to serve the creator and respect the consumer, and our heritage is built on providing trustworthy and innovative solutions to our customers. As our technology becomes more sophisticated, our products and features have the potential to impact our customers in profound and exciting ways. However, we believe we have a role that goes beyond creating the world's best technology. We are committed to ensuring that our technology and the use of our technology benefits society. At Adobe, as we innovate and harness the power of AI in our tools, we are dedicated to addressing the harms posed by biased data in the training of our AI. AI Ethics is one of the core pillars of our commitment to Digital Citizenship, a pledge from Adobe to address the consequences of innovation as part of our role in society.

How AI is used in Adobe's products

We believe AI will enhance human creativity and drive value from the complex global digital ecosystem. For the Creative Cloud, we are focused on making it easier for everyone to tell their story with simpler and more intuitive tools. As part of our Digital Experience offerings, Adobe's enterprise customers use our AI to deliver relevant and meaningful insights and personalized digital experiences to their end customers. And with Document Cloud, AI-enabled features help understand the structure of PDFs to assist the user in viewing, searching, and editing documents on any platform.

However, we recognize the potential challenges inherent in this powerful technology. AI systems are based on data, and that data can be biased. AI systems trained on biased data can unintentionally discriminate or disparage, or otherwise cause our customers to feel less valued. Therefore, we are committed to maintaining a principled and ethically sound approach to ensure our work stays aligned with our intended outcomes and consistent with our values. And we are actively participating in government discussions around the world to shape AI Ethics regulation for the good of the consumer and effectiveness in the industry.

AI Ethics Principles

At Adobe, we believe responsible AI development is based on the following three principles:

- **Responsibility:** We will approach designing and maintaining our AI technology with thoughtful evaluation and careful consideration of the impact and consequences of its deployment. We will ensure that we design for inclusiveness and assess the impact of potentially unfair, discriminatory, or inaccurate results, which might perpetuate harmful biases and stereotypes. We understand that special care must be taken to address bias if a product or service will have a significant impact on an individual's life, such as with employment, housing, credit, and health.
- **Accountability:** We take ownership over the outcomes of our AI-assisted tools. We will have processes and resources dedicated to receiving and responding to concerns about our AI and taking corrective action as appropriate. Accountability also entails testing for and anticipating potential harms, taking preemptive steps to mitigate such harms, and maintaining systems to respond to unanticipated harmful outcomes.

Adobe's Commitment to AI Ethics

- **Transparency:** We will be open about, and explain, our use of AI to our customers so they have a clear understanding of our AI systems and their application. We want our customers to understand how Adobe uses AI, the value AI-assisted tools bring to them, and what controls and preferences they have available when they engage with and utilize Adobe's AI-enhanced tools and services.

AI development and AI ethical review is still in its infancy. With any such complex topic, errors may occur, but with the commitment of our engineers and with help from our employees, our products and features will be best-in-class while continuing to reflect Adobe's values.

Adobe's AI Ethics Principles

Responsibility

We at Adobe place a high value on taking responsibility for the impact of our company and the innovation we deliver to the world. It stems from taking pride in our work and our dedication to the best possible outcomes. Therefore, we've determined that responsibility is the critical foundational principle to underpin Adobe's commitment and efforts toward developing AI. We must understand and address the impact of introducing new technologies such as AI.

Responsible development of AI encompasses the following: designing an AI system thoughtfully, evaluating how it interacts with end users, exercising due diligence to avoid unwanted kinds of bias, and vetting the AI system to determine when the behavior of the system is unacceptable. This entails anticipating potential harms, taking preemptive steps to mitigate such harms, measuring and documenting system performance throughout the technology lifecycle, and establishing systems to monitor and respond to unanticipated harmful outcomes.

Responsibility and Bias

The behavior of AI features is strongly dependent on the data used to train them. We understand that AI models can result in unintended biases in training data, whether produced by the selection of data to include or by the customer actions that produce the data, can result in correspondingly biased behavior in the system. Therefore, we are committed to building and curating AI training sets in order to avoid harmful bias for the instances where bias perpetuates societal stereotypes which in turn negatively impact people's lives. However, we understand that all data has bias; therefore, we are committing to ensuring that the output of our AI systems are remediated for bias, regardless of the input.

As part of developing and deploying its AI systems, Adobe will seek to mitigate unintended bias related to human attributes (e.g. race, gender, color, ethnic or social origin, genetic or identity preservation features, religion or belief, political belief, geography, income, disability, age, sexual orientation or vocation), and will apply a special focus, and strict standards of fairness and inclusiveness, to situations where the outcome would have an outsized impact on an individual's life, such as access to information about employment, housing, credit, and health.

Our ultimate goal is to design for inclusiveness rather than exclusion or discrimination. We will also determine whether the advantages of using AI outweigh the risk of harm of using AI at all.

This notion of fairness, however, does not imply a rigid uniformity of experience across customers, as some of the most typical AI use cases distinguish between individuals in ordinary and acceptable ways, as in demographic marketing or personalized product recommendations. Responsible development of AI means using AI in reasonable ways that accommodate the norms and values of our society.

Responsibility and Adobe's Digital Media Tools

It is possible, despite reasonable prevention efforts, that an outside party (customer or otherwise) might make use of Adobe's AI technology, such as with our video and photo tools, in a way that results in questioning the authenticity of content. Adobe feels a responsibility to support the creative community, and society at large, and is committing to contributing to solutions that address the issues of manipulated media.

However, we believe it is important to be clear about actions to which Adobe is not committing. We do not believe that we will develop AI that:

- Ensures egalitarian fairness, in the sense of providing identical experiences to all users. The value of AI is in its ability to differentiate, and as long as the AI is free of unfair bias, enabling customization and personalization of tools, recommendations, and technologies is valuable, important, and welcomed by our users.
- Certification that inbound training data is "unbiased". All training data encodes bias in one direction or another, as humans create the data point, and humans have inherent bias. We do not believe it is reasonable or useful to commit to a zero-bias starting point; rather, it is more important and effective to commit to having an outcome that mitigates against bias.

Accountability

Accountability means the commitment to take ownership for the outcomes of our actions. At Adobe, while anyone involved with AI has an obligation to help ensure it's being managed responsibly, business leaders are held accountable for the ethical operation of Adobe's AI technologies. We are ensuring processes are in place and resources are dedicated to meet Adobe's AI Ethics commitments, including to develop and implement the necessary engineering practices to achieve our responsibility goals, receive and respond to internal and external concerns, and to take corrective action as required.

How We Ensure Accountability:

- Establishing governance processes to evaluate and track the performance of AI algorithms, data and designs, including labeling datasets and models for any identified bias to ensure remediation can occur at product design stage;
- Requiring an AI Impact Assessment (as part of our services development process) to ensure an AI ethics review happens before deployment of new AI technologies;
- Creating an Ethics Advisory Board to oversee the promulgation of AI development requirements and be a place where any AI ethics concerns can be heard, while safeguarding ethical whistleblowers;
- Processes to ensure remediation of any negative AI impacts that are discovered after deployment;
- Education of engineers and product managers via mandatory training courses on AI ethics issues.

Transparency

Transparency is the reasonable public disclosure, in clear and simple language, of how we responsibly develop and deploy AI within our tools. Adobe values our trusted relationship with our customers and feels that transparency is integral to that relationship. This includes sharing information on how or whether Adobe collects and uses customer assets and usage data to improve our products and services.

Transparency includes the disclosure of the following:

- Our data collection practices:
 - When and if individual's data will be collected for AI training, and what controls a user will have over the collection
 - Providing notice prior to and if human-review of customer data for AI training will be implemented
- Model development:
 - How datasets are used in building AI models
- Accountability processes:
 - How Adobe is testing for and resolving issues related to unfair bias.
- General disclosure of how data and AI are used in Adobe's tools and services;
- Provide external and internal feedback mechanisms to report concerns on our AI practices