


Building responsible generative AI solutions for the new creative era.

Generative AI can fulfill its greatest creative potential when developers and businesses handle the associated risks responsibly. Adobe is committed to developing and deploying AI systems thoughtfully, taking seriously our responsibility to use technology in ways that minimize harmful bias.



How Adobe builds ethical AI solutions.

Guided by our AI ethics principles of **accountability**, **responsibility**, and **transparency**, Adobe's standardized process depends on **real people at every step**.



Mitigate risk to maximize quality.

All features go through Adobe's AI Ethics governance process, a system of dedicated resources that includes three bodies:

- The **AI Ethics Committee** works with our development teams to understand and incorporate our AI Ethics Principles.
- The **AI Ethics Review Board** reviews features before release.
- The **AI@Adobe Working Group** explores internal generative AI adoption.

These groups ensure the proper steps are taken to mitigate risks:

Datasets

During the AI development build and train phase, training is not a one-and-done process. The model typically must be "tuned" by training on a larger, more diverse dataset or on a completely different dataset than originally expected.

Security

Adobe security testing includes working with industry researchers and ethical hackers.


Model bias and ethics

We design for inclusiveness by testing and remediating potentially discriminatory or inaccurate results using an iterative training process.



1. Plan

The **product team** submits an ethics assessment to the **AI Ethics Team**, focusing on features that have the highest potential impact to our customers. The two teams collaborate to identify and mitigate risks. This process repeats until the features meet Adobe's AI Ethics principles.



2. Build and train

The **product team** trains the model, evaluates the output, and adds to the dataset as necessary.



3. Integrate

The **product team** integrates models into an application for a specific purpose or function.




4. Test

Adobe teams conduct testing to reduce the potential for biased and harmful outcomes in our generative AI products. **Adobe Security** tests these features and their integrated solutions at the same time.



5. Review

The **AI Ethics Review Board** may conduct a final review. The board's diversity helps them identify issues before launch and deployment.



6. Launch and use

Typically, **customers** in a private enterprise beta program use features first and provide feedback before they are released in a public beta testing scenario.

7. Feedback and refine

We welcome and encourage **your** feedback on our AI-powered features and technologies. Adobe Firefly includes a built-in feedback mechanism to report biased or inaccurate results.



Explore the power of generative AI in Adobe Creative Cloud.

Discover the tools behind a powerful, efficient approach to creative collaboration.

[Learn more](#)

[Get in touch](#) with an Adobe team member to explore what Adobe Firefly can do for you.



Adobe, the Adobe logo, and Adobe Firefly are either registered trademarks or trademarks of Adobe in the United States and/or other countries. All other trademarks are the property of their respective owners.

© 2024 Adobe. All rights reserved.