

Adobe Generative AI: Redefining Productivity in Creative Imaging

About this research

This report presents the findings of a research project conducted by Pfeiffer Consulting for Adobe, focusing on the impact of recently introduced generative AI functionality, such as Generative Fill in Photoshop, and the impact its use can have on productivity in creative workflows.

Benchmarks were executed using *Pfeiffer Consulting's Methodology for Productivity Benchmarking*, which has been fine-tuned over more than a decade and measures the time that experienced operators take to execute specific real-world tasks. Please refer to the Methodology section on the last page of this document for more information.

About Adobe generative AI technology

Generative AI has been introduced by Adobe in two different ways: **by feature additions to key Creative Cloud apps** such as Adobe Photoshop and Adobe Illustrator, and through **the dedicated web service Adobe Firefly**, that provides image generation from text prompts and other generative tools.

More specifically, Generative Fill and Generative Expand in Photoshop can have a significant impact on creative possibilities in retouching. This research shows that these features can increase productivity in common retouching tasks very significantly: In our benchmarks of real-world retouching tasks, **Generative Fill and Generative Expand were on average over ten times faster** than traditional tools and provide creative options impossible without these features.

Key benchmark results: Average of 10 workflow benchmarks




Chart based on the average of 10 different real-world workflow benchmarks.
Reference value: Average time when working with traditional retouching tools. **Shorter is better.**

Executive Summary

- Adobe has introduced generative AI technology, **available through the Adobe Firefly website**, in Adobe Express, as well as in several **generative AI features in Creative Cloud applications** such as Adobe Photoshop and Adobe Illustrator.
- Generative AI **can significantly extend the creative workflow** and make creative professionals more productive in every phase of a creative project.
- The newly introduced Generative Fill feature in Photoshop **significantly speeds up common retouching, composing and editing tasks**, and provides new perspectives for ideation and creative image manipulation.
- In benchmarks conducted for this project, **Generative Fill and Generative Expand were on average over ten times faster** than using traditional retouching methods. (See chart below.)
- Generative Recolor, available in Illustrator and through Adobe Firefly, **can speed up adapting colours of vector art significantly.**

Redefining Productivity in the Creative Retouching Workflow

Human creativity in the age of AI

In the creative industries, AI has become a touchy subject. Ever since generative AI systems, trained on millions of online images, have become capable of generating images that would have taken a skilled artist hours if not days to produce in a few seconds, controversies about the role of AI with regards to human creativity have soared.



In this loaded context, Adobe's singular approach to AI is standing out. Ever since the company started introducing AI-based tools in the creative workflow many years ago, **Adobe's focus has always been to help creative professionals work faster and better – not to replace them.** And the latest developments, Adobe Firefly and generative AI functionality introduced in apps like Photoshop and Illustrator, show that **AI can indeed provide impressive – and in some cases, spectacular – productivity gains** to creative professionals, help them work faster and better, without trying to replace their creativity. This research was designed to analyse and quantify these productivity gains and to look at how generative AI fits in the creative workflow.

Major Points

- Adobe Generative AI focuses not on replacing human creativity, but aims to provide **tools that make professionals more productive and extend their creativity.**
- Generative AI functionality is available through the dedicated web-service **Adobe Firefly** and also in **key new features of Creative Cloud apps** such as Photoshop and Illustrator.
- Generative AI can **increase productivity in every phase of the creative workflow.** (See table below.)

Productivity Gains of the AI-Assisted Creative Workflow

	Ideation	Asset Creation	Creative Elaboration	Finalisation
AI Feature	Use Adobe Firefly to ideate faster and explore different alternatives more quickly.	Create entire visual assets using Adobe Firefly. Enhance and adapt existing assets using generative AI in Photoshop and Illustrator.	Use generative AI in Photoshop to quickly implement requested changes. Use Generative Recolor to quickly iterate vector art.	Use generative AI in Photoshop alongside traditional retouching functionality and tools to quickly finalise assets.
Benefit of AI	Generative AI can extend and accelerate the human creative process.	Generative AI can produce images that are not easily found elsewhere. Adapting existing assets is accelerated significantly.	Stakeholder-requested changes can be implemented much faster using generative AI in Photoshop and Illustrator.	Generative AI features in Photoshop significantly increase productivity in production tasks.



What generative AI contributes to the creative workflow

Adobe provides its latest installment of AI-based functionality in two distinct ways. **The first one is Adobe Firefly:** a web-based service that lets users create or edit images using a natural-language description (also called a text prompt) that describes the desired outcome. Unlike other AI image generation systems, **Firefly provides a wide range of presets** to guide the user and make it easier to control the result. Firefly also can create extremely sophisticated text stylisations and recolour vector graphics.

The **second way** of accessing generative features is **directly in key Creative Cloud applications.** Right now, Photoshop and Illustrator are supported, but there is little doubt that more are to follow soon. Particularly in Photoshop, the use of generative AI constitutes a watershed moment for creative retouching. **Generative Fill and Generative Expand are probably the most significant feature additions to the program in many years,** allowing retouch operations that could take hours to be completed in seconds or minutes, while allowing creative operations that would simply have been impossible. In the following section of this report, we will discuss these features in detail and also analyse their concrete impact on productivity.

Adobe Firefly in action:

Adobe Firefly provides four distinct generative tools: image generation from text prompts and style presets, generative retouching of uploaded images, the creation of sophisticated text effects from text prompts, as well as generative recolouring of uploaded vector art.

Other generative tools, such as the possibility to generate images from 3D objects are under development.

Key Advantages of Adobe Generative AI for Retouching

	Without Adobe Generative AI	With Adobe Generative AI
Simple retouching tasks (cleaning up images, removing unwanted elements...)	<ul style="list-style-type: none"> Use standard retouching tools, such as Clone Stamp Tool, Spot Healing Brush Tool, Content-Aware Fill and others, to achieve the desired result. Can be time consuming and tricky. 	<ul style="list-style-type: none"> Use Generative Fill to remove unwanted elements from an image and reconstruct the background based on the content of the image in seconds. Each Generative Fill operation creates several variants, and a Photoshop layer and layer mask that are stored in the Photoshop file.
Complex retouching tasks (change background, add/modify elements, add effects...)	<ul style="list-style-type: none"> Use the wide range of advanced selection, composition and retouching features in Photoshop to produce the desired result. Requires solid experience with advanced Photoshop features. Can be very time consuming. 	<ul style="list-style-type: none"> Use Generative Fill to add elements or change the background of an image in seconds using generative AI. Each Generative Fill operation creates several variants, and associated Photoshop layer and layer mask that are saved with the file.
Creation of complete compositions (combine several images, create elaborate compositions...)	<ul style="list-style-type: none"> Manually combine and arrange multiple assets to create the desired composition. Requires solid experience with advanced Photoshop features. Can be very time consuming. 	<ul style="list-style-type: none"> Use Adobe Firefly or Generative Fill in Photoshop to create entire images from text prompts and a wide selection of style presets. Use Adobe Firefly to create sophisticated text stylisations based on text prompts and style presets.

Generative Fill: Revolutionising Everyday Retouching Tasks


The magic of Generative Fill

In this particular case, the use of the word magic is not hyperbole: what this feature can achieve seems to leave even seasoned Photoshop users speechless. How does it work? As soon as one makes a selection, a 'Generative Fill' palette pops up, where one can type a prompt describing the desired result, such as removing or adding an item. Previously, removing unwanted items from an image could be extremely time-consuming, especially if there were no details in the image to help the retouching process. **Generative Fill, on the other hand reconstructs the image, not only using elements and colour schemes from the present image, but using generative AI to produce a result that blends in.** In addition, Generative Fill not only leaves the original image untouched by creating a new layer and layer mask, but also provides several variations the user can choose from.

Take the illustration below as an example: To clean up the image on the right, we needed to remove not only the figure of the tourist on the left, but also the yellow awning and the soda-cans lined up in front of the concession stand behind the procession of Chinese monks. To make matters worse, the figure on the left is standing behind a grid structure,

Major Points

- Manually cleaning up images can be very time-consuming. **Generative Fill revolutionises this task**, using generative AI to remove or add items to a selected portion of an image.
- Productivity gains linked to Generative Fill are very significant, **allowing retouching operations to be completed in a fraction of the time.**
- Generative Fill not only removes items, but **completely reconstructs the content of the selection** using generative AI in a way that is coherent with the nature and the style of the image at hand.



Simple generative retouching: While in some cases, Photoshop tools such as the Clone Brush or the Spot Healing Brush can be adequate for cleaning up an image, things get more complicated when elements behind the removed item need to be reconstructed, as in the case of the yellow awning. (Left picture.)

Using Generative Fill, removing the tourist at the left (1), the soda-cans behind the monks (2) as well as the yellow awning (3) took just over five minutes. Manually retouching the same image took over fifty minutes in our benchmarks.



Original

Removing elements: Removing unwanted elements from images can be time-consuming. We compared removing the cacti on the table and lamp in the corner. Generative Fill was almost three times faster than traditional retouching.

Adding to a composition: While in traditional retouching, adding an element requires manually adapting the added element to the target image, this process is completely automated using Generative Fill. Adding a grand piano to the picture of the hall took less than two minutes, compared to 22 minutes using traditional methods.




Generative Fill

2 min. 49 sec.



Trad. Retouching

10 min. 37 sec.



Generative Fill

1 min. 37 sec.



Trad. Retouching

22 min. 02 sec.

which needs to be preserved. Producing the cleaned-up image **took just over five minutes**, working with three different selections, one of the figure of the tourist, another one of the soda-cans and a last one for the yellow awning. **In our benchmarks, removing these items manually took over 50 minutes.**

How does Generative Fill compare to traditional retouching in simpler tasks? **Even manually retouching a simpler scene is three to four times slower than Generative fill.** (See illustration above.)

Then there is the case where one needs to **add an element to a composition.** This can be complex, since even if a picture of the required element is available, chances are that it does not correspond to the perspective and lighting of the target image, requiring adapting the image manually (if possible.)

Generative Fill, on the other hand, can generate an element that takes into account not only lighting and perspective but also physical properties such as shiny surfaces. For our benchmark, we decided to add a grand piano to the image of an empty hall – **which took less than two minutes using Generative Fill, also generating the piano's reflection on the ground in the process.** (See Illustration above.) Producing a

comparable image using traditional retouching took 22 minutes (including five minutes looking for the appropriate image of a grand piano on Adobe Stock.)

In other words, **the productivity gains provided by Generative Fill are very significant.** Yet speed is not the only issue at stake here. Old hands at Photoshop will appreciate how this feature is implemented in terms of the overall workflow. As mentioned before, each time a generative fill operation is executed, Photoshop produces a new layer with layer mask, as well as 3 variants of the generated content. What is important, however, is the fact that **these layers, complete with the generated variants, are saved in the Photoshop file.** This means the user has access to them even days or weeks later and can change to a different variant if necessary.

Finally, it is important to point out that Generative Fill is not restricted to Photoshop users: Adobe Firefly allows users to upload an image and edit it using generative AI, albeit without the sophistication of Photoshop's selection tools. In any case, it is fair to say that **as far as image retouching goes, Generative Fill marks the beginning of a new era.**

How Generative Fill Expands the Boundaries of Retouching

Pushing the envelope of creative retouching

So far, we have been analysing use-cases that lie within the boundaries of typical creative image retouching. Generative AI, however, can go much further to assist creative professionals in their work and does so without involving a steep learning curve.

Take a simple example: You have the perfect shot of a model, but the client doesn't like the outfit, instead of a tweedy jacket they would like her wearing denim, for instance. **In a traditional workflow, this implies organising an additional photo-shoot**, which not only adds significantly to the cost but also takes considerable time. **Generative Fill can complete this change in seconds**: all it takes is to select the jacket, type 'replace with denim' as a prompt and click the Generate button.

It is important to note that generative AI doesn't always get it right. Artifacts can appear, but that is not a problem, however. Usually at least one of the variants generated fits the bill and if it doesn't, others can be generated in a few seconds. In addition, it is always possible to refine the generated image by additional generative operations.

Major Points

- Generative Fill and Generative Expand **usher in a new era in creative imaging** and can provide **very significant productivity gains** in the creative retouching workflow.
- Generative Expand can **extend images by reconstructing visual content** that is coherent with the source images.
- For all generative operations, detailed **text prompts can be used to fine-tune** the generated content.



Beyond traditional retouching: Starting with the top right image, Generative Fill was used with different selections and prompts to generate a picture of the young woman dressed as a chef in a kitchen (left image) or as a princess in a fantasy environment (right image.) The entire creative process took just over two minutes and under nine minutes respectively.

Integration with image characteristics: Combining different generative options in Photoshop can completely change an image in a few minutes. (Picture on the left.)



Elements added using Generative Fill take into account the content of the image and create appropriate reflections on shiny surfaces, including the correct reflection of the generated chair. (Picture on the right.)



Extending Images:

Using the Crop Tool in Photoshop, Generative Expand can reconstruct entire scenes with content based on the original image. Starting with a square view of the kitchen (small image), we created the expanded images using the text prompt 'Extend kitchen counter and add plants'. The entire process took just over a minute. (Picture on the right.)



In more general terms, Generative Fill goes beyond the limitations of traditional retouching in three ways: First of all by allowing the user **to add new elements that take into account the characteristics of the target image.**

One example: In the image above, the Teddy bear in a lounge chair was added with Generative Fill. The generative process included something almost impossible to create in traditional retouching: the correct shadow or reflection of an added element. (Top right illustration.)

The second essential application of Generative Fill is **the possibility to expand existing images.** In the bottom image above, the square image was expanded with the prompt 'Extend kitchen counter and add plants'. **Creating the image took just over 1 minute.** It goes without saying that this kind of possibility was simply unimaginable before generative tools and will change profoundly the way in which creative professionals will work with source images in the future.

Finally, it is **the way in which these features can be combined that completely redefines the boundaries of creative retouching.** A good example are the images on the previous page. Starting with a simple photo of

the young woman, it was possible to present her in a different profession (in our case as a chef standing in a kitchen) or in a fantasy environment, dressed in a robe and wearing a pearl necklace. This was achieved by respectively selecting different parts of the image and applying Generative Fill with the appropriate prompts. **Creating these images took two and just over eight minutes respectively. No manual retouching was required.**

But even with less extreme examples, the combination of generative features can provide significant productivity gains. It took **just over five minutes** to subtly but completely change the photo of the woman in the yellow coat above (Top left image.) **Attempting to achieve a comparable result using traditional retouching took an hour and a half.**

As noted before, what makes these features so powerful is their tight integration with the wide array of retouching and composition features available in Photoshop. It is not an exaggeration to say **that when Generative Fill is used to its full potential, it is not a simply a new tool, but a new dimension in creative imaging.**

Pushing the Limits of Image Creation with Adobe Firefly

What Adobe Firefly brings to the table

So far we have been concentrating on the impact of generative AI in the retouching workflow and have tried to quantify the productivity gains one can expect when using these technologies in typical image composition and retouching applications. It is time now to give Adobe Firefly, the online image generation system, a closer look.

While it is possible to create complete image compositions using the appropriate text prompts in Photoshop, this is much more easily achieved using Firefly. There are several reasons for this. Like other text-to-image generators, Firefly relies on text prompts, natural language descriptions of the desired output, to generate images. But while this sounds straightforward, prompting is quickly becoming a new field of technical expertise that can be daunting for newcomers. To lower the barrier of entry, the Firefly website provides hundreds of sample images, alongside the prompts that were used to generate

Major Points

- Adobe Firefly allows users **to create complex images from detailed text prompts** and can also be used for generative retouching.
- A wide array of **style presets and text prompt suggestions simplifies learning and exploring** the potential of Adobe Firefly.
- **Generative Recolor for vector art can significantly increase productivity** by selectively applying palettes based on a text prompt and style presets.

The Adobe Generative AI Toolbox

	Adobe Firefly (website)	Creative Cloud Apps
Image generation	Firefly can create complex images from a text prompt, offering a wide variety of presets for the type of the desired image, as well as for styles and lighting, among many others.	Photoshop can create complete images from a text prompt, but currently lacks the style presets Firefly offers.
Image editing (Generative Fill)	Firefly can modify uploaded images in a variety of ways, such as changing the background, adding or removing elements and more.	The Generative Fill option in Photoshop allows very sophisticated image editing and retouching that can be used on any selection.
Generative text stylisation	Firefly can generate very sophisticated text stylisations based on a text prompt. Stylised texts can be exported (including transparency information) for use in image composition or in video.	Creative Cloud apps currently do not offer integrated generative text stylisations, but can import text effects created by Firefly.
Generative Recolor for vector art	Firefly can use generative AI to recolor vector art based on a text prompt and offers a variety of presets and styles.	Generative Recolor for vector art is available in Illustrator, including presets and styles.
Extend image Support for 3D objects	Future developments of Firefly include the possibility to extend images, as well as generating images based on 3D objects, allowing greater creative guidance of the generative process.	Photoshop already supports Generative Expand using the Crop Tool. Images can be either extended in any direction and new elements can be added using a text prompt.

them, as a way to get started. Another important aspect that makes Firefly easy to use is that **it does not solely rely on text prompts, but provides a wide array of presets**, covering the nature of the desired image, such as content type (art, photo or graphic) or the desired style (for instance 'digital art', 'line drawing' or 'layered paper', among many others), as well as presets for lighting, colour and tone, and composition. Using these presets with a text prompts allows for very rapid exploration of available possibilities.

The question of copyright and intellectual property

As soon as the first text-to-image systems appeared on the market in 2022, strong controversies emerged concerning the status of the generated images and **if generated images could be considered a copyright infringement** because of the data it was trained on – a considerable concern for enterprises wanting to use generated images.

While in the case of the majority of generative image systems available today these concerns are far from being settled, Adobe has decided to address these issues right from the start. Indeed, **Adobe Firefly was designed to be commercially safe** and trained only on licensed Adobe Stock and public domain images where copyright has expired. In addition, Adobe has pledged to find **ways to compensate Adobe Stock contributors** for revenue generated by Firefly. **Adobe will also be offering enterprise customers IP indemnity**, which means that Adobe would protect customers from third party IP claims about Firefly-generated outputs*. Finally, Adobe has also committed to **clearly labelling Firefly images as AI generated**, which is in itself an important step for integrating generated images into a creative workflow that respects and supports human creativity.

* <https://techcrunch.com/2023/06/26/adobe-indemnity-clause-designed-to-ease-enterprise-fears-about-ai-generated-art/>

Recolouring vector art with AI

While for the average person, generative recolouring of vector art may not look as spectacular as the exploits of Generative Fill, for experienced Illustrator users, this new feature can significantly boost productivity.

Creating colour variations is an essential part of developing a creative concept and adapting it to different contexts. However, adapting colours to different desired 'moods' can be a complex task, since it requires changing a range of hues in a coherent fashion across a complex set of vectors.

Generative Recolor (available both in Illustrator and Adobe Firefly) can speed up this intricate process significantly. Users can use a text prompt to experiment different colour schemes, and use presets and style options to guide the recolouring process.

Generating Complex Images with Adobe Firefly



The images above were created with Adobe Firefly using the text prompt 'organic house embedded into the hilly terrain, architectural photography, style of architecture, futurism, modernist architecture' (picture on the left), and 'An ultra realistic image of teddy bear running at high speed with city

buildings in the background and sunny sky, create shallow depth of field and motion blur with long exposure' (Picture on the right). With each generative operation, four variants are created. It is important to note that even using the same prompt, the images created are never identical.

Methodology

This benchmark project was commissioned by Adobe and independently executed by Pfeiffer Consulting.

All the productivity measures presented in this document are based on real-world workflow examples, designed and executed by professionals with many years of experience with the programs and workflows involved.

How we measure productivity

The basic approach is simple: in order to assess productivity gains that a program or solution may (or may not) bring, we start by analysing the minimum number of steps necessary to achieve a given result in each of the applications or workflows that have to be compared.

Once this list of actions has been clearly established, we start to execute the operation or workflow in each solution, with the help of seasoned professionals who have long-standing experience in the field and with the solutions that are tested.

Every set of steps is executed three times, the average of the three measures is used.

About Pfeiffer Consulting

Pfeiffer Consulting is an independent technology research and benchmarking operation focused on the needs of publishing, digital content production and new media professionals.

For more information, please contact research@pfeifferreport.com

Pfeiffer Report

All texts and illustrations © Pfeiffer Consulting 2023. Images Clément Fuzier.
Reproduction prohibited without previous written approval.
For further information, please contact research@pfeifferreport.com.

The data presented in this report are evaluations and generic simulations and are communicated for informational purposes only. The information is not intended to provide, nor can it replace specific productivity research and calculations of existing companies or workflow situations. Pfeiffer Consulting declines any responsibility for the use or course of action undertaken on the basis of any information, advice or recommendation contained in this report and can not be held responsible for purchase, equipment and investment or any other decisions and undertakings based on the data provided in this report or any associated document.

Adobe, the Adobe logo, Creative Cloud, Firefly, Illustrator, Photoshop and Premiere 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.