



GENERATIVE ARTIFICIAL INTELLIGENCE AS A TOOL FOR JEWELRY DESIGN

Michael D. Magee
Jewelry Manufacturing Arts,
Research & Development

Gemological Institute of America
Carlsbad, California

INTRODUCTION:

The allure of jewelry lies in its intricate design, the craftsmanship behind each piece, and the significance and lore that we attach to particular jewels. From the skilled hands of ancient goldsmiths to the precision of modern-day artisans, the creation of jewelry has always been a deeply human endeavor. Yet, as we stand on the brink of a new era, the time-honored traditions of jewelry design are converging with the cutting-edge advancements of artificial intelligence. This convergence has opened up a new frontier in design and creation, where generative AI is beginning to appear in the designer's workshop alongside traditional tools and techniques.

For decades, CAD/CAM and other digital design tools have made computers commonplace in the jewelry designer's workshop. However, generative AI brings us into a new realm where computers are no longer mere tools directed solely by the hands of the designer, but have become capable of generating unique and even surprising designs with seemingly little to no direction or human input. In the context of jewelry design, this technology offers the potential to redefine the boundaries of creativity and production.

Despite the exciting possibilities, this fusion of AI with the creative arts is not without controversy. Concerns range from the potential displacement of skilled artisans to the ethical implications of AI-generated art. "AI-generated content, while impressive in its ability to analyze trends and produce designs quickly, may lack the nuanced touch and emotional connection that is often conveyed through a human designer's work. This could lead to a homogenization of design styles, as AI algorithms tend to optimize for efficiency and patterns in data rather than the bespoke artistry that clients seek."¹

Moreover, how we talk about AI, as a tool or as a co-creator, adds a layer of complexity to the question of ownership and copyright. “Anthropomorphizing AI can pose challenges to the ethical usage of this technology. In particular, perceptions of human-like agency can undermine credit to the creators whose labor underlies the system’s outputs and deflect responsibility from developers and decision-makers when these systems cause harm. We, therefore, discuss generative AI as a tool to support human creators, rather than an agent capable of harboring its own intent or authorship.”² As we navigate this uncharted territory, it is imperative to understand both the opportunities and challenges that generative AI brings to the table.

AIdeation and Hallucination

One of the most powerful opportunities that this technology presents is the ability to brainstorm design ideas at incredible speed. Sometimes referred to as “AIdeation,” using AI to quickly create ideas as starting points is becoming the most common use-case for AI in design. With only a few words, AI can create many iterations from a single concept or idea.

Here is an example of a potential “AIdeation” workflow. Let’s say that you are designing a pair of earrings for a client who is passionate about both Ancient Egyptian and Art Deco styles. After entering a short description, or “prompt”³ into Midjourney, you quickly have four original images to choose from blending these styles. None of these need to be the final design, simply stepping stones to more variations. Pick the one option that you want to explore and create variations based on that image. In approximately five minutes, you can easily have 16 or more variations of the design to begin designing around.




Figure 1 – Midjourney 5.2, variations on the prompt “a pair of 18 karat yellow gold earrings featuring a fusion of art deco and ancient Egyptian styles”

You may have heard the term “hallucinating” in reference to ChatGPT or other text-based AI chatbot programs, where the program makes up patently false information in its answers. While this is understandably a significant problem if you are expecting an AI program to curate or summarize factual information for you, it is actually a feature where image generation is concerned. The images created by generative AI are not collages cut-and-pasted from different images in the training data. They are original, “hallucinated” compositions.



Figure 2 – Midjourney 5.2, prompt “a pair of white earrings gold displaying a sense of asymmetrical balance”



Figure 3 – Midjourney 5.2, prompt “a rose gold pendant with contrasting round gems”

The tools are far from perfect, and at this time are better suited for generating new ideas instead of trying to perfect an existing image or design. So far, the AI is no substitute for a trained designer and has no concept of what can be manufactured. Notice the miniature engagement rings on the ends of the baguettes in Figure 4, and the floating and irregular bezels in Figure 5.



Figure 4 – DALL.E 3, prompt “create a halo of tapered baguette diamonds, also set in platinum, which are attached to the gallery of the paraiba’s setting”



Figure 5 – Midjourney 5.2, prompt “watercolor style, pear-shaped paraiba tourmaline brooch set in platinum with a halo of diamonds surrounding it”

Understanding Generative AI

Before going into any of the technical details, I would like to share an important caveat. I do not have a technical background in artificial intelligence, machine learning, neural networks, or any other computer science. I come to this technology with the background of a computer-savvy CAD modeler, designer, bench jeweler, and educator. With this in mind, I will attempt to describe how generative AI works. Thankfully, the most important thing that I have learned about how generative AI models work, is that understanding the technical details of the AI model is not at all important to using the tool effectively. A better comprehension of the underlying programming will not give you better results.

Broadly speaking, generative AI refers to a type of artificial intelligence that can generate new content, be it text, images, music, or even videos, which did not exist before. Generative AI essentially operates using complex algorithms and large amounts of data. What makes the process unique is the machine learning used to find patterns within that data.

Machine learning is the science of developing models that computer systems use to perform complex tasks without explicit instructions. This is where the AI learns from the data, absorbing information and becoming better over time. The system processes large quantities of data in order to identify patterns. Then the AI analyzes the data with algorithms and makes predictions based on

that data analysis. The AI makes many attempts to process the data, testing itself and measuring its performance after each round of data processing. In this way, the AI “learns” from mistakes and gradually improves its ability to generate new content that’s increasingly sophisticated and realistic. Given enough text, images, or other data, generative AI can find patterns linking similar concepts together and then create a new result that follows the same patterns.

Creating Images

“Creating noise from data is easy; creating data from noise is generative modeling.”⁴

So how does generative AI actually create images? AI image creation models are a specific application of generative AI that focuses on creating visual content, such as pictures or graphics. These models have gained a lot of attention for their ability to create stunning, sometimes surreal images based on simple text descriptions. One of the most prevalent image creation models for generative AI is Stable Diffusion. To understand how Stable Diffusion works, it’s helpful to understand that the diffusion referred to in the name is similar to the concept of particle diffusion in physics. In an image, this “particle diffusion” is represented by each pixel moving or changing in a random direction, which slowly transforms the image into static or visual noise.

For example, let’s say that you have an image of a cat. “A forward diffusion process adds noise to a training image, gradually turning it into an uncharacteristic noise image. The forward process will turn any cat image into a noise image.”⁵ Step-by-step, the cat image is transformed into random noise.




Figure 6 – Steps of Forward Diffusion,
<https://stable-diffusion-art.com/how-stable-diffusion-work/>

Both the original cat image and the images for each step of added noise are given to another program which attempts to calculate exactly how much noise was added at each step. With enough data, this second program can create a working noise prediction model. Once this is accomplished, then the noise prediction model can be applied in reverse.

In reverse diffusion, “We first generate a completely random image and ask the noise predictor to tell us the noise. We then subtract this estimated noise from the original image. Repeat this process a few times. You will get an image of a cat.”⁵




Figure 7 – Steps of Reverse Diffusion, <https://stable-diffusion-art.com/how-stable-diffusion-work/>

With enough data, the generative AI can assign reverse-diffusion pathways for any type of image or prompt that it is trained on. Of course, there are other steps going on within the code to compress the amount of data being manipulated, streamlining the process.

Prompts

At the time of this writing, text prompts are still the most common way of interacting with any generative AI program. Somewhere between a coding language and regular sentence structure, each word in a text prompt is converted into its own set of weighted vectors, which are used to “steer” the noise toward a desired outcome. Each program has its own language requirements and quirks that can be used to optimize the user’s control over the image.

When you give the AI a text description, it uses what it learned during training to translate those words into a visual format. For example, if you ask for an image of a “sunset over the ocean,” the AI combines its understanding of “sunset” and “ocean” to create an entirely new image that fits the description. By assigning each part of a prompt different weights, the model will use a different

noise prediction. Combine those prompt weights with a randomly generated noise starting image and you get a near-infinite number of possible outcomes.

Instead of creating an image using only text as a prompt, most platforms will also allow you to use another image as a starting point, or a combination of image and text prompts together. This prompt can be an original image input by the user, or an image previously generated by the AI. The program will create new variations based on the original image’s identifiable content, color scheme, and style.



Figure 8 – Photograph of “River” brooch by Dolmen Metalworks, sterling silver and Mississippi River-stone cabochons (left) and eight necklace variations created in Dream Studio SDXL v1.0 along with the prompt “silver necklace, water motif, sandblast and high-polish, cabochons” (right)

Beyond creating entirely new images, generative AI’s abilities can be focused on a particular aspect of an image. Four of these commonly used tools are outpainting, inpainting, upscaling, and blending. Each generative AI program also has its own proprietary features, abilities, and specialties in addition to these common tools.

Outpainting expands the image, effectively zooming out and then filling in what the AI expects to be in the scene around the original image. This can typically be done with or without an additional text prompt explaining what should fill the outer section.



Figure 9 – Midjourney 5.2, before and after outpainting

Inpainting requires the user to identify a portion of the image that should be replaced. The AI then only updates that selected subsection of the image.




Figure 10 – Midjourney 5.2, inpainting prompt
“More plants on the table”

Upscaling allows the user to increase the size of the image without reducing the resolution. When the image is upscaled, the AI divides each pixel in the image into the appropriate number of new pixels, adding new information while increasing the overall image size.




Figure 11 – Midjourney 5.2, before and after upscaling

Blending will take two or more existing images and create a new image that combines aspects of each of the originals.



Figure 12 – Two images created with Midjourney 5.2 (left and center) and their resulting blend image (right)

Prompt 1 “enamel and yellow gold pendant, Alphonse Mucha style, art nouveau, african american woman in profile, on a background of flowers, sketch on white paper, drawn by a master designer --s 220”

Prompt 2 “A jewelry designer and an AI computer drafting a new design on paper together in a jewelry designer’s workshop”

Review of Popular Generative AI Programs

In this next section, we will take a look at four of the most popular generative AI programs, Midjourney, Dream Studio, DALL.E, and Firefly. Each platform has its own pros and cons, which we will explore in detail.

Midjourney is currently the most popular AI image generator, and one of the most robust. The images can be highly detailed, creative, and realistic, and the prompts have a reasonable understanding of jewelry terminology.

Subscription plans range from \$10–\$120 per month depending on the level of membership. The image generation speed varies depending on the tier of membership plan purchased, but even at the basic plan tier, it is reasonable to get a set of four images within 30 seconds. All of the images created are public except on the two highest membership tiers.

The main disadvantage to Midjourney is that it can only be accessed through the Discord interface. Initially created for video gamers to interact with each other while playing games, the Discord community has branched out into many different areas of interest beyond gaming. However, the Discord interface can be daunting to new users, especially if you are unfamiliar with video

gaming chats and streaming platforms.

Dream Studio, by Stability AI, is an open source generative AI program. The main interface is easy to use and is accessed through their website. Users are given many free credits to use, and can purchase more as needed.

The most common complaint about Dream Studio is that the images are less creative than those of other programs, particularly Midjourney, which is thought to have more of an artistic flair. One area where Dream Studio excels is in making variations on an existing image. Once a starting image is uploaded, a slider can be set to determine how closely or loosely the variations should follow the original.



Figure 13 – Dream Studio SDXL v1.0, AI generated image (left) and variations with 25% and 50% image strength (right)

One of the other interesting features of Dream Studio is that as an open-source program, you can choose to install it locally on your own computer instead of using it through their webpage. While this may slow the generative process down significantly, depending on your own computer's resources, it does offer the advantage of keeping your created images private from other users.

DALL.E by Open AI is made by the same company that created ChatGPT and has recently become exceedingly popular. The latest version of DALL.E is fully integrated with ChatGPT, which means that users can prompt the program with a conversation instead of a code-like prompt.

While the images created can be extremely detailed, the main downside is that the image creation process is typically slower than other programs. Free accounts are available for both programs; however at the time of this writing, it is necessary to purchase a \$20/month subscription to access the versions of ChatGPT and DALL.E which interact with each other.

DALL.E can also remember previous conversations and can write prompts for you if you ask it to describe uploaded images. Since integrating with ChatGPT, DALL.E is fast becoming a popular contender to Midjourney.

Adobe Firefly is a relative latecomer to the generative AI field, but brings a unique set of assets to the table, mainly Adobe's existing Creative Cloud suite. The new Firefly tool is both a stand-alone image generator as well as built in to Photoshop, where it can be used to generatively fill areas of a composition with excellent in-painting and outpainting abilities.

As a standalone product, Firefly currently costs \$4.99/month and is bundled into the cost of Photoshop or Creative Cloud. The Firefly website is easy to use to quickly generate and modify images. One of the main advantages to Firefly is also its main disadvantage for creating jewelry, Adobe's ethical use of training data. Most of the other generative AI programs have been trained on trillions of images pulled from the internet, which includes a vast quantity of jewelry images. In comparison, Firefly has been trained on hundreds of millions of Adobe stock images, openly licensed content, and public domain images. Therefore, the regulatory landscape is clearer for Adobe Firefly, which we will discuss in depth later in this paper. However, this is a considerably smaller dataset than the other programs, and an even smaller proportion of it is trained on jewelry. Unfortunately, Firefly has less of an understanding of jewelry forms and terminology.

To better compare these four programs, the same prompt was used to create four variations of a tanzanite ring on each program. For DALL.E, a more natural sentence structure was used for the same prompt.



Midjourney 5.2



Figure 14 – Midjourney 5.2, prompt “platinum three-stone ring, tanzanite center gem, diamond side gems, photorealistic”



Dream Studio SDXL v1.0



Figure 15 – Dream Studio SDXL v1.0, prompt “platinum three-stone ring, tanzanite center gem, diamond side gems, photorealistic”



Figure 16 – DALL.E 3, prompt “Create an image of a platinum three-stone ring, with a tanzanite center gem and diamond side gems, photorealistic.”



Adobe Firefly 2



Figure 17 – Adobe Firefly 2, prompt “platinum three-stone ring, tanzanite center gem, diamond side gems, photorealistic”

MAGEE

New Ethical, Legal, and Regulatory Challenges

As with any technological revolution, the introduction of generative AI into jewelry design brings a new set of challenges that must be navigated with care. This section will delve into the ethical considerations, legal implications, and the evolving regulatory landscape surrounding AI-generated art. By understanding these issues, jewelry designers can make informed decisions about how to integrate AI into their processes responsibly.

Ethical Considerations

The use of AI in creative industries raises important ethical questions. One primary concern is the authenticity of the designs. When a piece of jewelry is created with the heavy involvement of

AI, who becomes the true author? This is a crucial question for designers who pride themselves on the uniqueness and originality of their work. Moreover, there's the concern about the potential for AI to create and proliferate designs that closely resemble original artworks without proper attribution or compensation to the original designers. "Users of AI art generators can even directly ask the AI to generate art in the style of a human artist."⁶ As a response to this criticism, many programs are beginning to put in guardrails against directly referencing living artists in their prompts.

Another ethical issue is the impact of AI on craftsmanship. Jewelry design is an art form that traditionally requires a high degree of skill and years of practice. If anyone can use AI to generate intricate designs in a matter of minutes, then what becomes of the skilled artisans? There's a risk that the value placed on manual skills and craftsmanship could diminish, leading to a loss of heritage and tradition within the field.

Thankfully, the jewelry industry is fundamentally about tangible goods. While AI may be able to create extraordinary images, at some point the design still needs to be manufactured. It may be possible that by making the tools of design more easily accessible, greater value will be placed on the skills needed to turn those designs into reality.

Legal Landscape and Regulatory Implications

Legal issues predominantly revolve around copyright and intellectual property (IP) rights. The current legal frameworks were established in a pre-AI era and are thus not fully equipped to handle the nuances of AI-generated works. For instance, if an AI creates a jewelry design based on a prompt from a human designer, who owns the copyright? At the time of this writing, the US Copyright Office brings some clarity to this scenario:

"In the Office's view, it is well-established that copyright can protect only material that is the product of human creativity. Most fundamentally, the term "author," ... excludes non-humans. ... For example, when an AI technology receives solely a prompt from a human and produces complex written, visual, or musical works in response, the "traditional elements of authorship" are determined and executed by the technology—not the human user."⁷

In short, art solely generated by an AI cannot be copyrighted or attributed to any author. However, "an artist may modify material originally generated by AI technology to such a degree that the modifications meet the standard for copyright protection. In these

cases, copyright will only protect the human-authored aspects of the work.”⁷

Moreover, there is the complication of derivative works. If an AI is trained on a dataset that includes copyrighted jewelry designs, there’s a possibility that its outputs could infringe on the IP rights of the original creators. To address these remaining questions about copyright and ownership, the US Copyright Office is currently conducting a study regarding the copyright issues raised by generative AI. In addition, this same issue is being addressed in the courts, where there are currently several high-profile intellectual property lawsuits in progress.⁸

Some of the larger AI companies are taking steps to help their customers address the risks of copyright infringement. Recently, OpenAI announced that they will pay the legal costs incurred by their enterprise-level customers who face claims around copyright infringement as it pertains to ChatGPT and DALL.E.⁹

One notable exception to the intellectual property battles in generative AI is Adobe Firefly. As noted above, Adobe has trained their generative AI model using only licensed Adobe Stock and public domain images where copyright has expired.¹⁰ This means that Adobe is able to guarantee that their AI images do not violate any existing copyrights, and offer limited IP indemnification for their enterprise customers using Firefly.¹¹

Adobe also now automatically adds specific metadata to any AI modified or generated image created with their software. These “Content Credentials” can also store the content creator’s information, creating a more durable record of ownership. To further this goal of content transparency, Adobe founded the Content Authenticity Initiative¹², a non-profit group creating tools to verify the provenance of digital images.

Adobe is also a co-founder of the Coalition for Content Provenance and Authenticity¹³, a standards development organization developing open, global standards for sharing information about AI-generated images across platforms and websites beyond Adobe’s products.

The regulatory environment for AI in creative arts is still in its infancy. Governments and international bodies are only beginning to grapple with the implications of AI and are in the process of developing laws and guidelines to govern its use. This includes ensuring that AI-generated works respect copyright and IP laws, and setting standards for transparency and accountability in AI use.

For jewelry designers, staying informed about these regulations is critical. As laws evolve, there may be new compliance requirements that affect how designers can use AI tools. Designers will need to keep abreast of these changes to ensure that their use of AI remains within legal boundaries. Furthermore, there is a need for education within the jewelry design community about the legal and ethical implications of AI. Designers should be equipped with the knowledge to use AI responsibly, respecting both the letter and the spirit of intellectual property laws.

Conclusion

How can we, as jewelry designers, best use this new technology to enhance our creative expression? First, we must acknowledge that generative artificial intelligence is already having a profound effect on the fields of art and design. Its impact on the jewelry design industry is a matter of when, not if.

At this time, the two greatest assets of generative AI are its accessibility and speed. Anyone can now create jewelry images in a matter of seconds using only a few words and the right program. However, this isn't enough. Going forward, the two greatest assets of a designer using these tools will be the ability to translate images into actual models, acting as a "reality-check" for generated designs, and the ability to personalize or customize, pushing the designs further toward a specific direction or purpose.

The images from generative AI may be eye-catching, but they still need to be made into real objects to become jewelry. A designer is needed to translate any 2D jewelry image into a 3D model by using CAD or traditional hand-carving or fabrication methods. To be effective, designers require a deep understanding of the metal and gem materials, manufacturing processes, the dimensions and tolerances needed for durability, and the ergonomics of how jewelry interacts with the wearer's body.

Personalization is where a human touch is added. Whether brainstorming ideas for a new jewelry line using prompts crafted from marketing predictions for the upcoming year, or designing a one-of-a-kind, bespoke jewelry piece for a single client, the AI images may only be the beginning. A good designer can take these AI-generated starting points and refine the design, whether to tie it in to an existing brand, tell the design story for one individual, or simply to push the boundaries of design even further. These personal touches also ensure that the final design is unique and copyrightable. As for which AI tool to use, there is enough overlap between them to be able to pick and choose whichever one works best for you.

However, jewelry designers using AI tools must be vigilant to ensure that their use of AI does not inadvertently violate the IP rights of others. Using AI to create variations of your own pre-existing work is an excellent way to ensure that the main input to the generated image came from your own creativity.

For a larger company willing to invest greater resources, there may be another interesting pathway. A company with a very large collection of designs (10,000+) could potentially create their own private generative AI platform, trained using only on their own designs. This would allow that company to generate new images based solely off their own existing portfolio, solving any copyright or intellectual property concerns.

Lastly, there should be an industry-wide push for transparency in AI tools. Designers should have a clear understanding of how AI models are trained, the data they use, and the potential for bias or infringement within these models. This transparency will help designers make more informed choices about the tools they use. There are still many unanswered legal and regulatory questions regarding this new technology waiting for their respective days in court. If 2023 was the year of AI, then 2024 may very well be the year of AI regulation.

Ultimately, like any new technology, it is up to each of us individually to decide to be an early adopter, taking on the possibilities of both risk and rewards, or to wait, watching and learning from the accomplishments and missteps of others.

Questions for Future Study:

The technology involved in generative AI is moving at a tremendous pace. The speed of image generation, as well as the complexity and accuracy of the prompts, improve with every update released. While these changes may make AI more accessible and acceptable for jewelry design, they may not fundamentally alter how this technology changes our industry. However, there are several areas of growth that bear watching: *sketch to render*, *text to video*, and *text to mesh* generation.

Sketch to render is a generative AI program which uses a specific type of image prompt along with text. The program asks for a sketched image, even a simple line drawing will do, and a text prompt. The AI combines the two to create a fully rendered image, filling in the color, materials, lighting, and shading of a 3D object based on your original sketch and text prompt. This type of prompt could work hand-in-hand with a talented designer or sales-

person by quickly turning their sketched ideas into realistic images for a client to review. One promising example of this technology can be seen at promeai.com.

Text to video takes the text-to-image format one step further, creating short videos solely from a text prompt. Several of the major existing AI companies, such as Midjourney and Stability AI, are promising this capability in their next set of upgrades. It often takes multiple images from different angles to properly describe a jewelry object. Many of us will be familiar with the top, front, and side views found in many CAD programs as well as hand-drafting techniques. A simple video turning 360° can describe a ring more accurately than any single image.

Text to mesh may take longer to develop properly as a technology, but it has the potential to dramatically impact the jewelry industry. Instead of training on the trillions of images found on the internet, these programs are trained on multitudes of 3D models. While these programs currently exist and can be found at sites like meshy.ai, they are predominantly geared toward building props and avatars for the video game industry. The low-resolution mesh quality of these models are not yet suitable for jewelry. One of the biggest hurdles for this technology is the relative paucity of 3D models available to train the AI, compared to the number of images used to train the 2D models.

One interesting development in this field is recent news that Turbosquid is currently developing a 3D-modeling generative AI tool. Turbosquid is an online 3D-model marketplace owned by Shutterstock, with over a million 3D models available for purchase, although most of the models are not jewelry specific. Their collection of high-resolution models may be one of the few large enough to properly train a text-to-mesh program which could generate high-quality 3D meshes that could be imported into any CAD program.

References

1. “What Potential Drawbacks Could Interior Designers Face with AI-generated Content in 2024?” *JEMSU*, jemsu.com/what-potential-drawbacks-could-interior-designers-face-with-ai-generated-content-in-2024. Accessed 2 Jan. 2024.
2. Epstein, Ziv, et al. “Art and the Science of Generative AI: A Deeper Dive.” *arXiv*, 2306.04141, arxiv.org/abs/2306.04141.
3. Midjourney. “Prompts.” *Midjourney*, docs.midjourney.com/

docs/prompts. Accessed 15 Sep. 2023.

4. Song, Yang, et al. "Score-Based Generative Modeling Through Stochastic Differential Equations." *arXiv*, 2011.13456, arxiv.org/abs/2011.13456.

5. "How Does Stable Diffusion Work." *Stable Diffusion Art*, stable-diffusion-art.com/how-stable-diffusion-work. Accessed 11 Nov. 2023.

6. "What AI Art Means for Society, According to Yale Experts." *Yale Daily News*, 23 Jan. 2023, <https://yaledailynews.com/blog/2023/01/23/what-ai-art-means-for-society-according-to-yale-experts/>

7. United States Copyright Office. "Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence." 16 Mar. 2023.

8. "Generative AI Has an Intellectual Property Problem." *Harvard Business Review*, 7 Apr. 2023, hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem.

9. "OpenAI Promises to Defend Business Customers Against Copyright Claims." *Tech Crunch*, 6 Nov. 2023, techcrunch.com/2023/11/06/openai-promises-to-defend-business-customers-against-copyright-claims.

10. "Adobe Firefly for Enterprise." *Adobe*, www.adobe.com/products/firefly/enterprise.html. Accessed 19 Oct. 2023.

11. "Adobe Firefly IP Indemnification." *Adobe Help Center*, helpx.adobe.com/legal/product-descriptions/adobe-firefly.html. Accessed 21 Dec. 2023.

12. "Content Authenticity Initiative." *Content Authenticity Initiative*, opensource.contentauthenticity.org. Accessed 21 Dec. 2023.

13. "Coalition for Content Provenance and Authenticity." C2PA, c2pa.org. Accessed 21 Dec. 2023.