

# GENERATIVE ARTIFICIAL INTELLIGENCE AS A TOOL FOR JEWELRY DESIGN

Michael D. Magee

While the jewelry industry is deeply rooted in traditional tools and techniques, generative artificial intelligence (AI) must be recognized today as a revolutionary tool to assist jewelers. The intersection between AI and the creative arts is especially controversial, and the use of generative AI in jewelry design is no different. This paper examines how generative AI creates designs based on text and image prompts, comparing five of the more common AI programs (Midjourney, DALL.E, Stable Diffusion, Leonardo, and Firefly) with a focus on generating realistic jewelry images. It addresses the ethical, legal, and regulatory questions that have arisen around AI-generated art. Lastly, it explores how this new technology can best be used by jewelry designers to enhance creative expression.

The allure of jewelry lies in its intricate design, the craftsmanship behind each piece, and the significance and lore 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. But 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, with generative AI making its way into the designer's workshop alongside traditional tools and techniques.

For decades, computer-aided design and manufacturing (CAD/CAM) and other digital design tools have been widely used by jewelry designers. However, generative AI reaches into a new realm where computers have become capable of generating unique 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. As one U.S.-based digital advertising agency noted, "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" (JEMSU, n.d.).

## "AIdeation" AND "HALLUCINATION"

One of the most powerful opportunities that AI technology presents is the ability to brainstorm design ideas at incredible speed. Sometimes referred to by the play on words "AIdeation," quickly creating

### In Brief

- Generative artificial intelligence programs can create highly detailed, photorealistic images of gems, jewelry, or any other subject in seconds.
- Designers may use these tools to create wholly new concepts for jewelry or to create variations on an existing idea.
- The legal and regulatory framework for AI-generated art is not yet settled. There are a wide range of copyright concerns and legal cases around source material still in question.
- This technology is rapidly advancing and already impacting the field of jewelry design.

See end of article for About the Author and Acknowledgments.

*This article was originally published in a different form by The Jewelry Symposium for the Advancement of Jewelry Manufacturing Technology and is published with permission of The Jewelry Symposium.*

GEMS & GEMOLOGY, Vol. 60, No. 3, pp. 330–347,

<http://dx.doi.org/10.5741/GEMS.60.3.330>

© 2024 Gemological Institute of America




Figure 1. Midjourney 5.2 image variations produced from the prompt “a pair of 18 karat yellow gold earrings featuring a fusion of Art Deco and ancient Egyptian styles.” At each stage in this figure, the user selects one of the images (marked here by the colored square boundary) and feeds it back into the AI, which generates four new variations. In this manner, the image is further refined to suit the tastes of the designer or the client.

ideas as starting points is becoming the most common use case for AI in design. With an input consisting of only a few words, AI can generate many iterations from a single concept or idea.

As an example of a potential “AIdeation” workflow, consider the design of a pair of earrings for a client who is passionate about both ancient Egyptian and Art Deco styles. After a short description, or “prompt,” is entered into Midjourney ([docs.midjourney.com/docs/prompts](https://docs.midjourney.com/docs/prompts)), the program quickly returns four original images blending these styles for the user to choose from (figure 1). These images are simply stepping stones to more variations. The user can choose any of the options to explore and create variations based on that image, making widely varied options at first and then refining the image through minor alterations. In approximately five minutes, Midjourney can generate 16 or more variations to begin designing around.

“Hallucination” is a common term used in reference to the popular ChatGPT or other text-based AI chatbot programs, where the program makes up patently false information in its answers. While this understandably poses a significant problem when at-

tempting to collect and summarize factual information, it is actually a benefit 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. A simple prompt can generate entirely new designs, such as the collection of white gold earrings in figure 2 or the rose gold pendants in figure 3.

The tools are far from perfect, and for now they are better suited for generating new ideas than perfecting an existing image or design. So far, AI is no substitute for a trained designer, and it has no concept of what can actually be manufactured. Examples of AI design irregularities can be seen in the odd miniature diamond rings on the ends of the baguettes in figure 4 and the “floating” bezels in figure 5.

## UNDERSTANDING GENERATIVE AI

The good news for most jewelry designers is that knowing the technical details of the AI model is not at all important to using the tool effectively. A complete grasp of the underlying programming will not give the casual user better results.



Figure 2. Images from Midjourney 5.2 using the prompt “a pair of white gold earrings displaying a sense of asymmetrical balance.”



Figure 3. Images generated by Midjourney 5.2 from the prompt “a rose gold pendant with contrasting round gems.”

Broadly speaking, generative AI refers to a type of artificial intelligence that can generate new content,

including text, images, music, or even videos that previously did not exist. Generative AI essentially oper-

Figure 4. DALL.E 3 generated this image from the prompt “create a halo of tapered baguette diamonds, also set in platinum, which are attached to the gallery of the Paraiba’s setting.” Circled in red is a design oddity: miniature diamond rings placed on the ends of the baguettes.



Figure 5. Another image output from Midjourney 5.2, using the prompt “watercolor style, pear-shaped Paraiba tourmaline brooch set in platinum with a halo of diamonds surrounding it.” Circled in red are irregularly shaped gems and “floating” bezels.



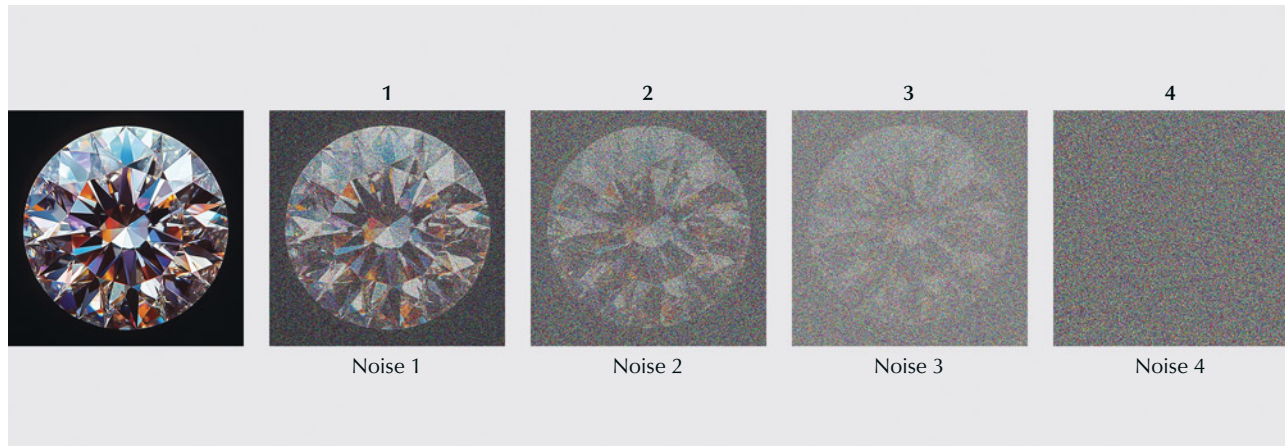


Figure 6. Steps of forward diffusion. Noise is added to the image at each step until it is completely randomized.

ates 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 performing 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 increasingly sophisticated and realistic new content. 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 (Sanderson, 2017; see also <https://www.youtube.com/@3blue1brown>).

**Creating Images.** So how does generative AI actually produce images? AI image creation models are a specific application of generative AI that focuses on visual content, such as pictures or graphics. These models have captured global media 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 the diffusion model. To understand how this model works, it is 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, slowly transforming the image into static or visual noise.

For example, consider the image of a diamond in figure 6. A user might start with this to train a model to create other diamond or gem images. Step by step, a little bit of noise is added to the training image. Eventually, the training image is completely unrecognizable as a diamond, showing only randomness. This forward diffusion process takes the original image and adds noise, gradually transforming it into an unrecognizable noise image (Stable Diffusion Art, 2024).

Both the original diamond image and the images for each step of added noise are fed into a neural network model, which attempts to calculate exactly how much and what kind of noise was added at each step. With enough data, this neural network model can create a working noise prediction model. Once this is accomplished, the noise prediction model can be applied in reverse.

In reverse diffusion, the process begins with a randomly generated image. The noise predictor then calculates the amount of noise that has been added to the image. That noise is subtracted from the original randomized image. The process is repeated until it resolves into the image of a diamond, as shown in figure 7 (Stable Diffusion Art, 2024).

These models are trained not on a single image, but on trillions of images. With enough data, the generative AI can assign reverse diffusion pathways for any type of image or prompt it is trained on. A properly trained model will produce many different images of diamonds, or any other subject, based on the learned pathways from each of the training images,




Figure 7. Steps of reverse diffusion. Predicted noise is removed step by step from a randomly generated image until it resolves into a recognizable image.

and a model trained on jewelry images will be able to generate entirely new images of gems and jewelry.

**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 idiosyncrasies that can be used to optimize the user’s control over the image.

When given a text description, the AI applies what it learned during training to translate those words into a visual format. For example, if the user prompts it 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 a weight to each part of a prompt, the model will use a different noise prediction. Combining those prompt weights with a starting image of randomly generated noise will return a near-infinite number of possible outcomes.

Figure 8. Left: A sterling silver “River” brooch by Dolmen Metalworks featuring Mississippi River stone cabochons. Photo by Michael Magee. Right: Eight necklace variations created in Dream Studio SDXL v1.0 using the prompt “silver necklace, water motif, sandblast and high-polish, cabochons.” The sixth image was generated without stones.





Figure 9. Pendant design from Midjourney 5.2 using the prompt “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.”

Rather than creating an image using only text, most platforms allow the user to upload an image as a starting point or use a combination of image and text prompts. This prompt can be an original image input by the user or an image previously generated by an AI. In figure 8, we see an example of a photograph used in conjunction with a text prompt to create variations based on the original image’s identifiable content, color scheme, and style.

Beyond creating entirely new images such as the one in figure 9, generative AI’s capabilities can be focused on a particular aspect of an existing image. Four




Figure 10. Images from Midjourney 5.2, before and after outpainting.

of the 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 found in the scene around the original image (figure 10). This can typically be done with or without an additional text prompt explaining what should fill the outer portion.

**Inpainting** requires the user to identify a portion of the image that should be replaced (figure 11). The AI then updates that selected subsection of the image. This can also be done either with or without an additional text prompt to define what should change within the image.

Figure 11. Images from Midjourney 5.2 using the inpainting prompt “more plants on the table.”

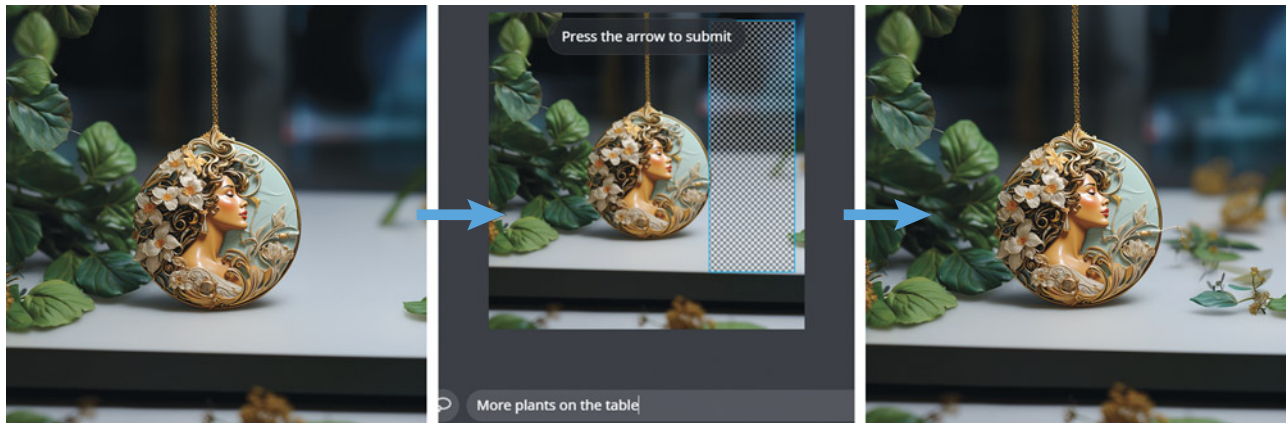




Figure 12. Images from Midjourney 5.2, before and after upscaling.

**Upscaling** allows the user to enlarge an image without reducing its resolution (figure 12). When the image is upscaled, the AI divides each pixel into the appropriate number of new pixels, adding new information while increasing overall image size. The end result is an image that is both larger in overall size and has increased definition. No additional text prompt is used, as the AI is not adding or subtracting any prompted subject matter. Instead, the AI uses the image itself as the prompt to fill in the details at a finer resolution.

**Blending** takes two or more existing images and creates a new image that combines aspects of each of the originals (figure 13), with or without an additional text prompt.

## REVIEW OF POPULAR GENERATIVE AI PROGRAMS

This section briefly looks at five of the most popular generative AI programs: Midjourney, Stable Diffusion, DALL.E, Leonardo.AI, and Firefly. Each platform has its own pros and cons, described below and compared in box A.

**Midjourney** is currently one of the most popular and robust AI image generators available. The images can be highly detailed, creative, and realistic, and the responses to prompts reflect a reasonable understanding of common jewelry terminology (see figures 1–3 and 9–13).

Subscription plan pricing varies depending on the membership tier. The speed of image generation varies

Figure 13. Two images created with Midjourney 5.2 (left and center) and their resulting blend image (right). Prompt 1 (left): “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.” Prompt 2 (center): “a jewelry designer and an AI computer drafting a new design on paper together in a jewelry designer’s workshop.”






Figure 14. Dream Studio SDXL v1.0 AI-generated image using the prompt “platinum three-stone ring, tanzanite center gem, diamond side gems, photorealistic” (left) and variations with 25% and 50% image strength (right).

accordingly, but even the basic plan can generate a set of four images within 30 seconds. Typically, all of the images created are made publicly available on Midjourney’s website and are visible to all other users. The two highest membership tiers offer a privacy mode that does not share generated images.

The main disadvantage of Midjourney is that currently it can only be accessed through the Discord interface. Initially created as a chat and messaging app for video gamers to interact with one another, the Discord community has branched out into many areas beyond gaming. However, the Discord interface can be daunting to new users, especially those unfamiliar with video gaming chats and streaming platforms. To alleviate this challenge, Midjourney is developing a web-based interface. At the time of this writing, however, the web version is in a closed development stage and only open to select users.

**Stable Diffusion**, by Stability AI, is an open-source generative AI program. The online interface is easy to use and is accessed through the website interface, called DreamStudio. Users are given free credits to start with and can purchase more as needed.

The most common criticism of DreamStudio is that the images are less creative than those generated by other programs, particularly Midjourney, which is thought to have more of an artistic flair.

One area where DreamStudio excels is in making variations on an existing image (see figure 8). Once a starting image is uploaded, a slider can be set to determine how closely or loosely the variations should follow the original (figure 14). A higher image strength

value will make any new images adhere more strictly to the original, while a lower image strength will give the AI more “creative license” to make variations.

Stable Diffusion is offered as an open-source program and can be installed locally on a computer rather than using it through the DreamStudio webpage. While running Stable Diffusion locally on a computer may slow down the generative process, depending on the computer’s specifications, it does offer several key advantages. One is that all the images created are kept private from other users and the Internet at large. Most other programs charge users a premium to keep their data private. The other significant advantage is that an advanced user can customize the diffusion process with samplers. A sampler is an add-on to the diffusion model that can nudge the image creation toward specific styles or results (Gilgamesh, 2023). This grants the user greater control over the specific qualities and details of the images generated.

**DALL.E** is made by OpenAI, the same company that created the well-known ChatGPT. The latest version of DALL.E is fully integrated with ChatGPT, which means users can prompt it with a written or spoken conversation instead of a code-like prompt (see figure 4).

While the images created can be extremely detailed, the process is typically slower and less customizable than other programs. Free accounts are available for both programs; however, at the time of this writing, a monthly subscription is necessary to access the latest versions of ChatGPT and DALL.E.



## BOX A: COMPARISON OF AI PROGRAMS

To compare the five most widely used generative AI programs, the same prompt was used to create four variations of a tanzanite three-stone ring using each program (figures A-1 through A-5). The prompt used was "platinum three-stone ring, tanzanite center gem, diamond side gems, photorealistic." For DALL.E, a complete sentence structure was used for the same prompt.



Figure A-1. Images generated by Midjourney 6. These designs feature intricate details, including organic scrolls and fine milgrain. The gems are in the appropriate color range for tanzanite and diamonds. The images also feature a certain amount of depth of field, making them more photorealistic.

Figure A-3. Images generated by DALL.E 3 using the prompt "create a photorealistic image of a platinum three-stone ring, with a tanzanite center gem and diamond side gems." These images show some detail and creativity in the designs and feature a wider variety of gem shapes. However, in the bottom left image, some of the details are poorly defined, such as the milgrain along the channel setting, the engraving on the inside of the shank, and the detail on the sides of the pear setting. In all of the images, the colors are reasonably accurate for the materials. The images are less photorealistic and more closely resemble digital renderings from a CAD program. This stylistic rendering seems to be consistent across images created with DALL.E.



Figure A-2. Images generated by DreamStudio SDXL v1.0. These designs are less detailed but still mainly accurate depictions of jewelry. The metal and gemstones show the appropriate color range for their materials. The image on the bottom left shows some inaccuracies on the prong settings, blurring between prongs and diamonds on the center setting and adding extra prongs to the side setting on the left. Again, a certain amount of depth of field is apparent, making the images more photorealistic.





*Figure A-4. Images generated by Leonardo.AI Phoenix. The rings in these images are remarkably similar, with little variation in design or positioning. The designs are also very simple in comparison to the previous three examples. A greater variety of prompts would be needed to create more design variations. All the images show some degree of depth of field, and the bottom left image also displays a particular blurring of lights in the background, creating a photorealistic effect.*



*Figure A-5. Images generated by Adobe Firefly 2. These rings are less accurate depictions of the “three-stone ring” concept given in the prompt—particularly the top left image, which has no side stones at all. The desired materials, tanzanite and platinum, are also poorly depicted and could easily be confused for blue sapphire and white gold. Upon closer inspection, the center “tanzanite” gems show an oversimplified faceting pattern not common in commercial cuts. The bottom right ring also appears to be slightly out of round. These inaccuracies are likely due to the smaller amount of jewelry-specific training data available to Firefly.*

DALL.E can also remember previous conversations and write its own prompts when asked to describe uploaded images. Since integrating with ChatGPT, DALL.E has quickly become a popular contender to Midjourney.

**Leonardo.AI** is another image generator that has recently become a very popular rival to Midjourney. While Leonardo has a more complicated web interface than Dream Studio, DALL.E, and Firefly, it offers more variety and control over both the types of prompts available and the style of the images created.

Leonardo’s base memberships are free with a limited number of generations per month. Several tiers of paid memberships are available, offering greater generation speed and other options. At the lower membership tiers, all images uploaded to and created

by Leonardo become part of the program’s dataset, with private accounts available for a monthly fee.

Like Midjourney, Leonardo has a well-developed understanding of jewelry terminology and is likely to create images that technically match the prompt. Part of what makes Leonardo unique is the ability to combine text prompts with a live sketching program, called Realtime Canvas. The generated image constantly updates, within the parameters of the text prompt, as the user sketches through their web interface or through the iOS app. Leonardo also allows users to “train” their own image models by uploading up to 40 images. The AI then uses this custom dataset to fine-tune its results for the user.

**Adobe Firefly** is a relative latecomer to the generative AI field but offers a unique set of assets, namely Adobe’s existing Creative Cloud suite. The Firefly

---

tool is both a stand-alone image generator and a built-in plugin to Photoshop, where it can be used to generatively fill areas of a composition with excellent inpainting and outpainting capabilities.

Firefly can be purchased as a stand-alone product or bundled into the cost of Photoshop or Creative Cloud. The Firefly website is easy to use to quickly generate and modify images.

One of Firefly's main advantages is also its biggest disadvantage for creating jewelry: Adobe's unique sources of training data. Most other generative AI programs have been trained on trillions of images pulled from all over 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 will be discussed in depth in the next section. However, this is a considerably smaller dataset than the other programs, and an even smaller proportion of it is trained on jewelry. Consequently, Firefly has less of an understanding of jewelry forms and terminology.

## 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 evolving regulatory landscape surrounding AI-generated art.

How we think about AI—whether as a tool or as a cocreator—adds a layer of complexity to the question of ownership and copyright. Anthropomorphizing AI (attributing human qualities to it) can pose additional philosophical and moral dilemmas. It is helpful to consider generative AI as a tool used to support human creators rather than as an independent creator with its own ideas (Epstein et al., 2023). In navigating this uncharted territory, it is imperative to understand both the opportunities and challenges posed by generative AI. By understanding these issues, jewelry designers can make informed decisions about how to responsibly integrate AI into their processes.

**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 in-

volvement 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 are concerns over the potential for AI to create and proliferate designs that closely resemble original artworks without proper attribution or compensation to the original designers. Indeed, users can prompt the AI to generate images in the style of a specific artist (Yup, 2023). As a response to this criticism, many programs are beginning to place guardrails against directly referencing living artists in their prompts.

Another ethical issue involves 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 artisan? There is a significant risk that the value placed on manual skills and craftsmanship could diminish, leading to a loss of heritage and tradition within the field.

Fortunately, 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. By making the tools of design more easily accessible, greater value could potentially be placed on the skills needed to turn those designs into reality. A designer's understanding of gems and precious metals, manufacturing processes, ergonomics, and typical wear are instrumental in creating jewelry that is both beautiful and functional.

**Legal Landscape and Regulatory Implications.** The current legal issues for generative AI revolve predominantly 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 U.S. Copyright Office has brought 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. (United States Copyright Office, 2023)

---

In short, art generated solely by an AI program cannot be copyrighted or attributed to any author. However, the U.S. Copyright Office went on to say that “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” (United States Copyright Office, 2023).

There is also the complication of derivative works. If an AI is trained on a dataset that includes copyrighted jewelry designs, there is the possibility that its outputs will infringe on the IP rights of the original creators. To address these remaining questions about copyright and ownership, the U.S. Copyright Office is conducting a study regarding the copyright issues raised by generative AI. This same matter is being addressed in U.S. district courts, where several high-profile intellectual property lawsuits are currently in progress (Appel et al., 2023).

Some of the larger AI companies are taking steps to help their customers address the risks of copyright infringement. In 2023, OpenAI announced that it would pay the legal costs incurred by enterprise-level customers who face claims around copyright infringement as it pertains to ChatGPT and DALL.E (Wiggers, 2023).

One notable exception to the intellectual property battles in generative AI is Adobe Firefly. As previously mentioned, Adobe has trained its generative AI model using only licensed Adobe Stock and public domain images where copyright has expired (<https://www.adobe.com/products/firefly/enterprise.html>). This means Adobe is able to guarantee that its AI images do not violate any existing copyrights, and it can offer limited IP indemnification for enterprise-level customers using Firefly (<https://helpx.adobe.com/legal/product-descriptions/adobe-firefly.html>). However, it has more recently come to light that the Firefly model was trained partially on images generated by other AI platforms—about 5% in all (Metz and Ford, 2024). While this somewhat complicates the issue of intellectual property practices, Adobe stands by its moderation process and insists that its training data does not include intellectual property.

Adobe also automatically adds specific metadata to any AI-modified or -generated image created with its software. These “content credentials” can also store the content creator’s information, establishing a more durable record of ownership. To further this goal of content transparency, Adobe founded the Content

Authenticity Initiative (<https://opensource.contentauthenticity.org>), a nonprofit group for creating tools to verify the provenance of digital images.

Adobe is also a cofounder of the Coalition for Content Provenance and Authenticity (C2PA, <https://c2pa.org>), an organization that develops open global standards for sharing information about AI-generated media across platforms and websites beyond Adobe products. Other members of C2PA include Microsoft and Google, both of which are currently the subject of lawsuits for allegedly extracting data indiscriminately from the Internet to train their own AI 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, new compliance requirements may affect how designers can use AI tools. Designers will need to stay abreast of these changes to ensure that their use of AI remains within legal boundaries. Furthermore, there is a clear 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.

## AI-ENHANCED DESIGNERS

As with any new technology, early adopters are already using generative AI to enhance their design and sales workflows. Many of the tools listed in this paper are being used to drive sales, either by creating marketing content, inspiring the final designs, or producing images to close sales. At the time of this writing, AI-generated images are still error-prone, and designers may wish to create the images in private to cull any unworkable results before showing them to a client. However, as this technology advances, it will undoubtedly open up new options for designers who enjoy working in live sessions in front of their clients.

One of the most promising AI methods being pioneered to enhance a designer’s workflow is called “sketch to render.” This approach takes a loose sketch of a design as an image prompt and then uses



Figure 15. Hand-sketched bracelet design by Devyani Shah (left) and a Dzine.ai “rendering” (right) created using the sketch along with the text prompt “white gold bracelet with pave set diamonds in leaf-shaped clusters. A large oval emerald is set in the center of the bracelet. Art Deco inspiration. Photorealistic rendering. Dazzling light reflection on the gems. Neutral background.”

a generative AI program to create a fully rendered image, filling in the color, materials, lighting, and shading of a 3D object based on the original sketch and an accompanying text prompt (figure 15). A talented designer or salesperson can use “sketch to render” to quickly turn their sketched ideas into realistic images for a client to review. Some promising examples of this technology are Dzine.ai and Vizcom.ai. In figure 15, a hand sketch is used along with a detailed text prompt to create a more photorealistic version of the original sketch in Dzine. Looking closely at the AI-generated image shows some blend-

ing between the diamonds and the metal materials in the bracelet links. Dzine recently began training their models on jewelry-specific images in order to remove these kinds of inaccuracies.

The “sketch to render” method introduces significant enhancements to the design workflow. First, it brings the focus of the design process back into the hands of the human designer or salesperson. Designs can be drawn by hand and scanned, or drawn on a digital application on a tablet, computer, or even phone. Once drawn, an AI program can quickly process the sketch into a more photorealistic



Figure 16. Digital sketch (left) and a Dzine.ai “rendering” (right) using the original image and adding the text prompt “intricate yellow gold earrings, adorned with scintillating diamonds, dangle elegantly from the ear, their wavy design evoking a sense of fluidity, as if dancing in the light; photorealistic rendering, the warm golden hue and flecks of diamond sparkle radiating a sun-kissed glow, perfect for a summer evening soiree or a vintage-inspired bridal look, with hints of Art Deco sophistication and bohemian flair, every detail meticulously captured in the rendering’s subtle nuances.”

tic image, similar to the digital rendering of a CAD model (figure 16).

With the prevalence of CAD programs in the industry, custom jewelry purchasers have come to expect photorealistic renderings of their jewelry before giving approval. This means a designer often spends hours developing a CAD model for approval without first receiving a deposit on their time and work. When a design goes through many iterations before client approval, this CAD modeling time can add up to a significant cost.

By using AI to create a photorealistic image from only a sketch, a salesperson can get approval and potentially a deposit on a design before investing any time modeling in CAD. The photorealism of the AI-generated image can help the client better visualize the final product, while the designer can save the labor-intensive CAD modeling for the final design. The salesperson can even use AI to create variations of the original sketch to quickly explore design options.

Another significant benefit of this “sketch to render” method is that it creates a copyrightable foundation for the design early in the process. Instead of creating an uncopyrighted AI design with copyrightable human-made modifications, the designer creates a copyrighted original sketch and later adds AI modifications that may or may not be copy-

rightable. While not a guarantee of copyright protection, starting with a human-created drawing can help establish the “human authorship” requirement for a copyright.

Much of the effort of custom sales is the selling of jewelry from pictures—sketched by hand, digitally sketched, or CAD-rendered—as the desired jewelry does not yet exist. Now a designer can use AI to enhance their natural creativity, reduce unnecessary labor on unwanted designs, protect their creative ideas, and help their clients realize their ideal design faster.

## FUTURE DEVELOPMENTS

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 accepted for jewelry design, they may not fundamentally alter how this technology changes our industry. However, there are several growth areas that bear watching: *text to video*, *image to depth map*, and *text to mesh* generation.

**Text to Video** takes the text-to-image format one step further, creating short videos solely from a text prompt. It often requires multiple images from dif-

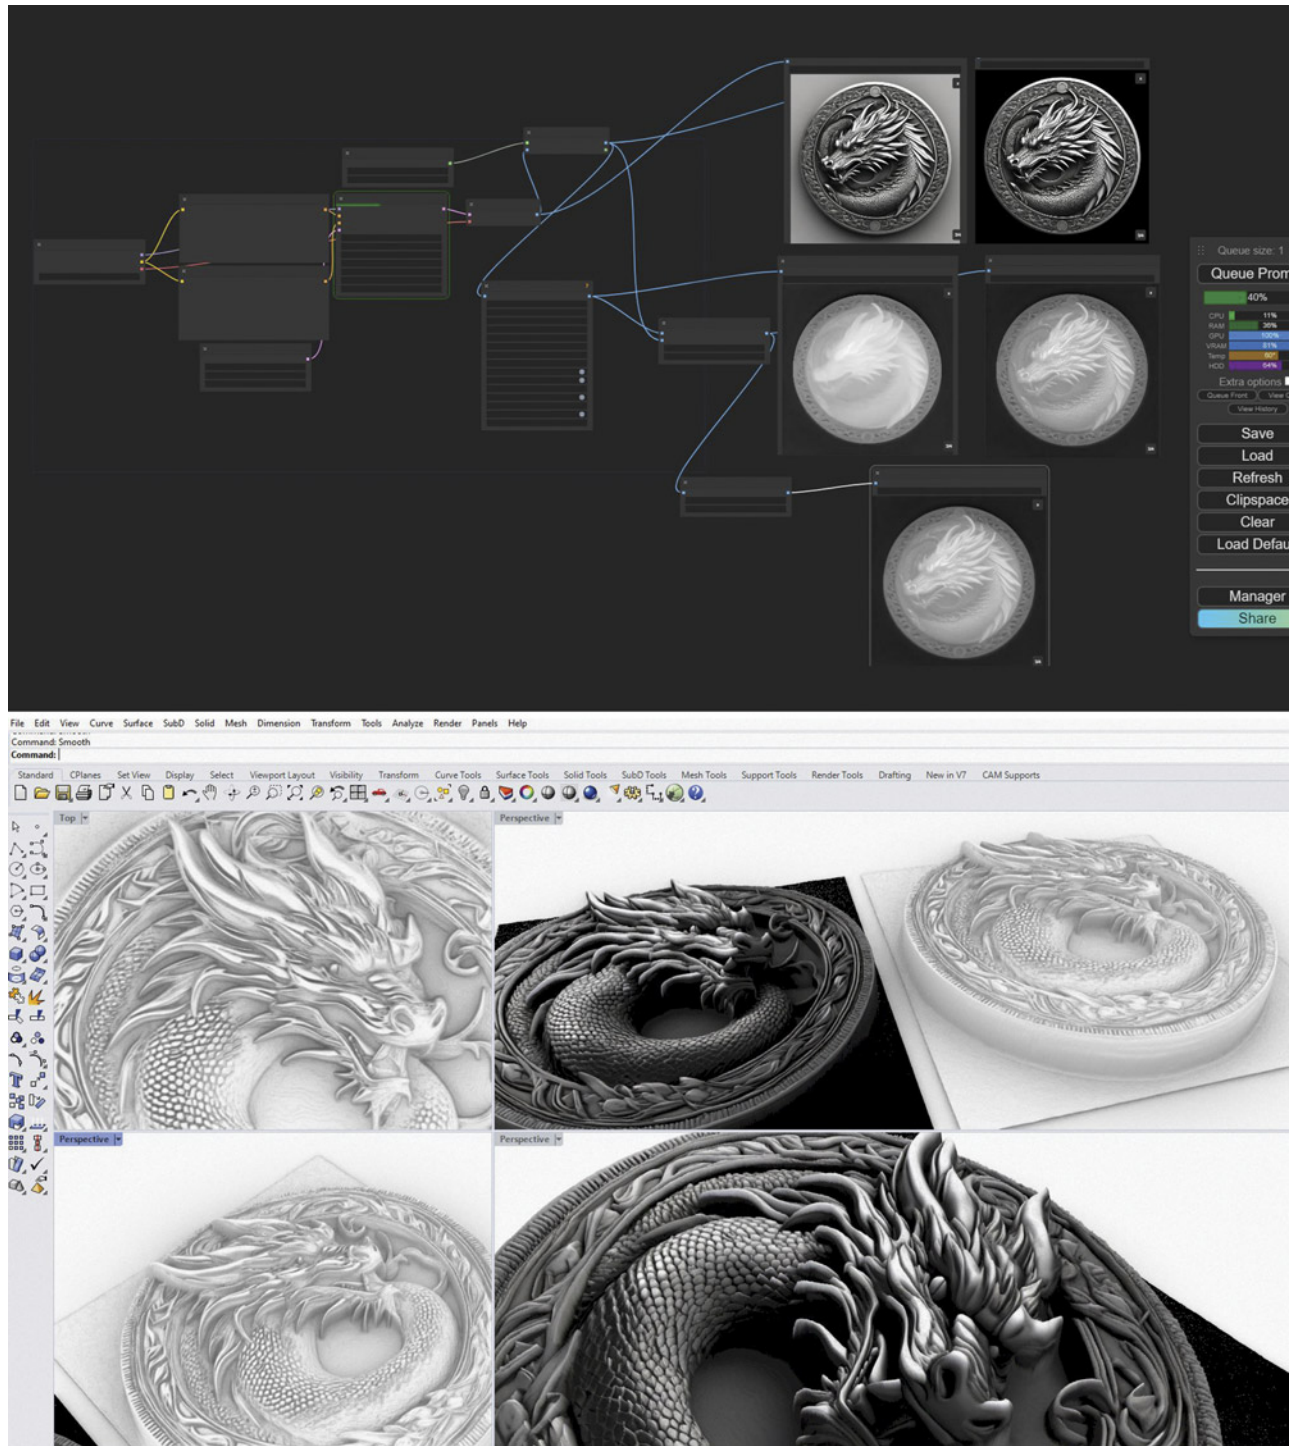


Figure 17. “Dragon Coin” by Chris Botha. Stable Diffusion image and depth map; 3D design in ZoeDepth and Rhinoceros 3D. The top image shows a customized program for Stable Diffusion using the ComfyUI user interface. ComfyUI allows the user to control Stable Diffusion with a visual programming language. Each piece of code is represented as a box with specific inputs and outputs. The boxes are connected together via these input and output connectors to modify the code, creating detailed images from the text prompt and then transforming the detailed images into heightfield images suitable for ZoeDepth. The bottom image uses Rhinoceros to show the 2.5D object created from this image by ZoeDepth.


ferent angles to properly describe a jewelry object. Many designers are familiar with the top, front, and side views found in many CAD programs, as well as hand-drafting techniques. A simple video rotating 360° can describe a ring more accurately than any single image. Several of the major AI platforms, such as Leonardo.AI and Stability AI, have already offered this capability in more recent upgrades, although at the time of this writing, these models still have a limited ability to rotate a complex jewelry design accurately (see video at [www.gia.edu/gems-gemology/fall-2024-artificial-intelligence-in-jewelry-design](http://www.gia.edu/gems-gemology/fall-2024-artificial-intelligence-in-jewelry-design)).

**Image to 2.5D Depth Map** is an interesting step toward creating producible models directly from AI. With some additional training, programs such as Stable Diffusion can be modified to create not only black-and-white images but also depth maps, which are grayscale images where the shades of gray denote the relative depth of a three-dimensional object.

Other programs, such as ZoeDepth (Bhat et al., 2023), can translate those images into actual 2.5D geometry, which can be manipulated directly in a CAD program such as Rhinoceros or ZBrush (figure 17).

**Text to 3D** may take longer to properly develop 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 such as Meshy.ai (figure 18), they are geared toward building props and avatars for the video game industry. The low-resolution quality of these models makes them not yet suitable for producing 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.

Figure 18. Images from Meshy.AI using the prompt “green emerald amulet, yellow gold, high detail, organic leaf and vine design.” This program creates fully three-dimensional models from a text prompt, as opposed to the flat 2D images of other programs. Some of the finer details are “painted” onto the model with a colored image that is applied to a simpler 3D object. Finer details suggest that the gem and pendant are separate objects (left). The same object with the colored image removed (right) reveals the actual 3D geometry created. Note that the gem and pendant are actually one object painted as different materials.





---

Several major 3D players, including TurboSquid, Autodesk, and Meta, are currently developing their own 3D modeling generative AI tools. TurboSquid, owned by Shutterstock, is the largest online 3D model marketplace, with over a million 3D models available for purchase, although most of the models are not jewelry specific. Autodesk, the maker of many of the most popular CAD programs, announced its Project Bernini for generative AI 3D shape creation in May 2024 (Autodesk, 2024). Meta (the parent company of Facebook, Instagram, and WhatsApp) recently published a paper about a new Gen3D program designed to create 3D objects from text for its Metaverse applications (Siddiqui et al., 2024). Each of these companies would have a collection of high-resolution models large enough to properly train a text-to-mesh program to generate high-quality 3D meshes that could be imported into any CAD program.

## SUMMARY

How can jewelry designers best use this new technology to enhance creative expression? Generative artificial intelligence is already having a profound effect on the fields of art and design. Its impact on the jewelry design industry has begun as well.

The two greatest assets of generative AI today are accessibility and speed. Anyone can now create jewelry images in a matter of seconds using only a few words and the right program. However, there is more to real jewelry product design than just a picture. Going forward, the two greatest assets for a designer using these tools will be the ability to translate images into actual models, acting as a “reality check” on the designs, and the ability to personalize or customize, pushing designs further toward a specific direction or purpose.

Images from generative AI may be eye-catching, but they still must be made into physical objects to become jewelry. A designer’s expertise 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 must still possess 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 will continue to add a human touch. Whether a designer is 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 into an existing brand, tell the design story for one individual, or simply push the boundaries of design even further. Likewise, a creative designer can start with their own designs and use AI to explore variations and transform their quick sketches into photorealistic images. These personal touches also help ensure that the final design is unique and copyrightable.

Given the overlap between AI tools, they can be chosen based on what works best for an individual business. However, a jewelry designer 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 one’s own preexisting work ensures that the main input for the generated image comes from personal creativity.

For a company willing to invest greater resources, there may be another interesting pathway. A company with a large collection of designs could potentially create its own private generative AI platform, trained on its own designs. This would allow that company to generate new images based solely on its own existing portfolio, solving any copyright or intellectual property concerns. Even a designer with a limited number of designs could use them to train an existing generative AI model to create images that are heavily influenced by their modest portfolio.

Lastly, an industrywide push is needed 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. Many unanswered legal and regulatory questions regarding this new technology remain, waiting for their respective day in court.

As with any new technology, each individual ultimately decides whether to become an early adopter, taking on the possibilities of both risk and reward, or to wait, watching and learning from the accomplishments and missteps of others.

## ABOUT THE AUTHOR

Michael Magee is a senior subject specialist in jewelry manufacturing arts at GIA in Carlsbad, California.

## ACKNOWLEDGMENTS

The author would like to thank the jeweler community at the online forum CAD/CAMTrav and Ian McAskill in particular for being

an early adopter of AI technology and helping to pave the path. Thanks to Greg Stopka for his excellent peer review and insights. Chris Botha shared his work in figure 17, and the author's colleague Devyani Shah shared her manual design skills in figure 15. Thanks also to The Jewelry Symposium for the inspiration to write and present this paper.

## REFERENCES

- Appel G., Neelbauer J., Schweidel D.A. (2023) Generative AI has an intellectual property problem. *Harvard Business Review*, April 7, <https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem>
- Autodesk (2024) Autodesk unveils Research Project Bernini for generative AI 3D shape creation. *ADSK News*, May 8, <https://adsknews.autodesk.com/en/news/research-project-bernini/>
- Bhat S.F., Birkl R., Wofk D., Wonka P., Müller M. (2023) ZoeDepth: Zero-shot transfer by combining relative and metric depth. *arXiv*, article no. 2302.12288, <http://dx.doi.org/10.48550/arXiv.2302.12288>
- Epstein Z., Hertzmann A., Herman L., Mahari R., Frank M.R., Groh M., Schroeder H., Smith A., Akten M., Fjeld J., Farid H., Leach N., Pentland A., Russakovsky O. (2023) Art and the science of generative AI: A deeper dive. *arXiv*, article no. 2306.04141, <http://dx.doi.org/10.48550/arXiv.2306.04141>
- Gilgamesh (2023) Tutorial: What is a sampler in Stable Diffusion? *Medium*, <https://gilgam3sh.medium.com/tutorial-what-is-a-sampler-in-stable-diffusion-d5c16875b898>
- JEMSU (n.d.) What potential drawbacks could interior designers face with AI-generated content in 2024? <https://www.jemsu.com/what-potential-drawbacks-could-interior-designers-face-with-ai-generated-content-in-2024>
- Metz R., Ford B. (2024) Adobe's "ethical" Firefly AI was trained on Midjourney images. *Yahoo Finance*, April 12, <https://finance.yahoo.com/news/adobe-ethical-firefly-ai-trained-123004288.html>
- Sanderson G. (2017) Neural networks. *3blue1brown*, <https://www.3blue1brown.com/topics/neural-networks>
- Siddiqui Y., Monnier T., Kokkinos F., Kariya M., Kleiman Y., Garreau E., Gafni O., Neverova N., Vedaldi A., Shapovalov R., Novotny D. (2024) Meta 3D AssetGen: Text-to-mesh generation with high-quality geometry, texture, and PBR materials. *arXiv*, article no. 2407.02445, <https://doi.org/10.48550/arXiv.2407.02445>
- Stable Diffusion Art (2024) How does Stable Diffusion work? <https://stable-diffusion-art.com/how-stable-diffusion-work>
- United States Copyright Office (2023) Artificial Intelligence and Copyright. *59942 Federal Register*, Vol. 88, No. 167, <https://www.copyright.gov/ai/docs/Federal-Register-Documents-Artificial-Intelligence-and-Copyright-NOI.pdf>
- Wiggers K. (2023) OpenAI promises to defend business customers against copyright claims. *TechCrunch*, November 6, <https://techcrunch.com/2023/11/06/openai-promises-to-defend-business-customers-against-copyright-claims>
- Yup K. (2023) What AI art means for society, according to Yale experts. *Yale Daily News*, January 23, <https://yaledailynews.com/blog/2023/01/23/what-ai-art-means-for-society-according-to-yale-experts/>

For online access to all issues of GEMS & GEMOLOGY from 1934 to the present, visit:

[gia.edu/gems-gemology](https://gia.edu/gems-gemology)

