

How to use Vision FX 2.0

This program enables you to transform photos and images into anything you can imagine, from cartoon characters to futuristic fantasy, classic or modern art, science fiction and anime, and much more, backed by the power of generative AI.

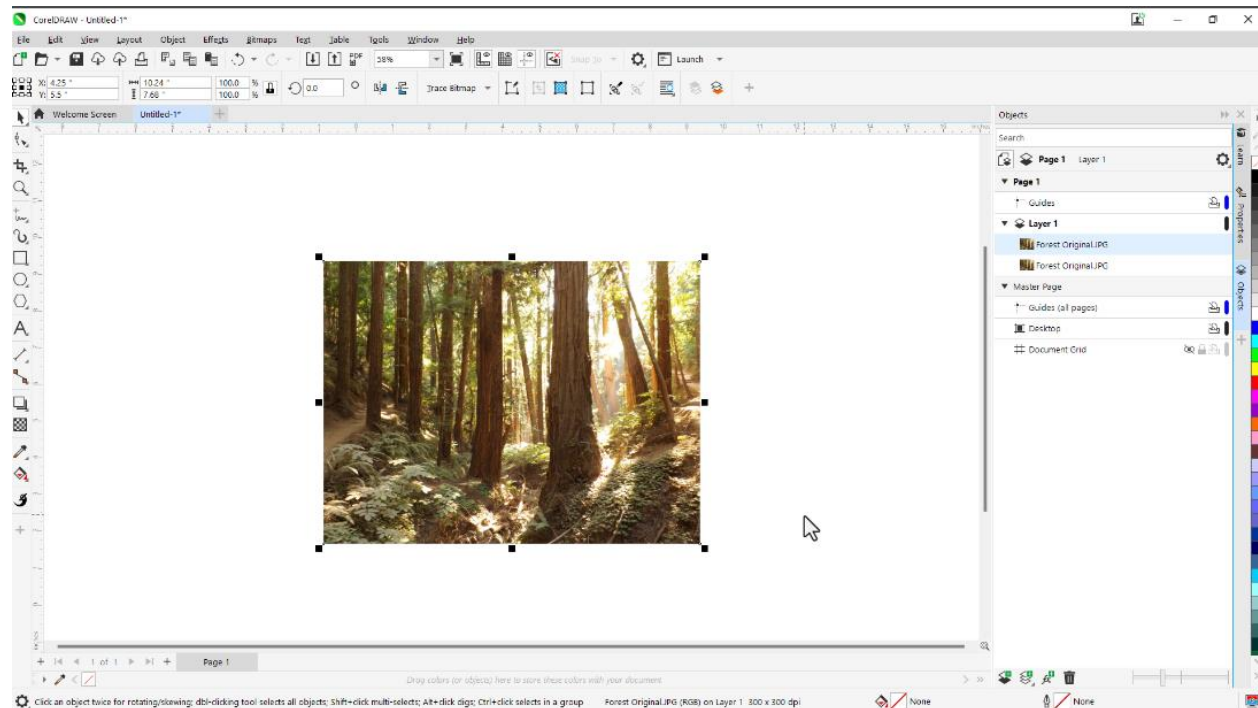
Because it runs directly on your local device, and not in the cloud, your creations are limitless, and your data remains private and secure.

Just provide a text prompt describing your idea, adjust a few simple settings, and the plugin generates several options to bring your ideas to life. Once your photo or image is transformed, and you choose your favorite result, bring it back into your host software program for final polishing or editing.

Optimize Vision FX performance

Vision FX uses an image-to-image AI algorithm, meaning that you need to start with a raster image, like this one I have imported into CorelDRAW. The Vision FX plugin can also be used with Corel PaintShop Pro, Corel PHOTO-PAINT and Corel Painter, as well as Adobe Photoshop and Affinity. Note that Vision FX doesn't work with vector graphics, only raster images such as JPEGs and PNGs.

To preserve my original image, I'll select it and press Ctrl + C and Ctrl + V, or Cmd C and V on the Mac. Now I have two JPGs in my Objects docker, or Objects inspector on the Mac.



I'll keep the copied layer selected, and choose Bitmaps / Plugins / Vision FX 2.0 / Vision FX 2.0.

Before you start using Vision FX, it's crucial to configure the program settings for optimal performance. Be sure to read over the Performance Tips that appear when you first open the program.

You can disable this pop-up if you don't want to see it every time you launch Vision FX. If you disable it, you can access the Performance Tips from the Help menu. We'll cover these tips as we go through the menus along the top, which are new in Vision FX 2.0.

There are two options in the Settings menu.

For AI Engine, the default is to run on your computer's CPU. But if your machine has an Intel or Nvidia graphics card, select the relevant GPU option to maximize performance. I'll show how to view your graphics settings in just a bit.

The Image Quality setting refers to the inference resolution. Choosing a higher resolution can help produce cleaner or more detailed images, but requires more processing time and power. Working with a lower resolution will speed up the processing time and still produce great results, just not as detailed or refined.

The Image Quality setting can be changed as needed. For example, you can use 512 when you're trying out different prompts and concepts, then increase to 1024 when you're ready to generate a final image at a higher quality. In addition, low resolution images can be enhanced with the Enhance Image Quality feature, which I'll show farther on.

Let's now open the Help menu. FAQ links to the VisionFX website with some frequently asked questions. Performance tips opens the tips popup we saw earlier, and System Info opens a handy window with your computer specs. This is where I can see my graphics card, so I know to use Nvidia GPU in Model Settings. Check for Updates is how you can check periodically for enhancements and new models, and About displays version details for the software and AI engine.

Create an AI-generated image

The Vision FX UI is pretty straightforward, but some fine-tuning may be needed to get the results you want. A sample prompt is provided at the top right, with comma-separated terms specifying a detailed digital painting in the style of Van Gogh, with mystical colors and beautiful lighting. There is also a field for negative prompts, which we'll explore a bit farther on.

If you're not sure how to get started with your own prompts you can click the Get Inspired button. This will take you to a gallery where you can view artwork created with Vision FX, plus the prompts for each image, which you can copy and use.

The Strength slider controls how much the AI results will deviate from the original. As strength increases, less of the original image will appear in the results. I'll move this down a bit, in order to compare strength results.

Text Guidance controls how strictly the text prompt will be adhered to, which I'll leave as is.

Seed is a random number used to generate the first AI result, and the other four results will use consecutive seed values starting from this number. Seed values can be regenerated to get a different set of results.

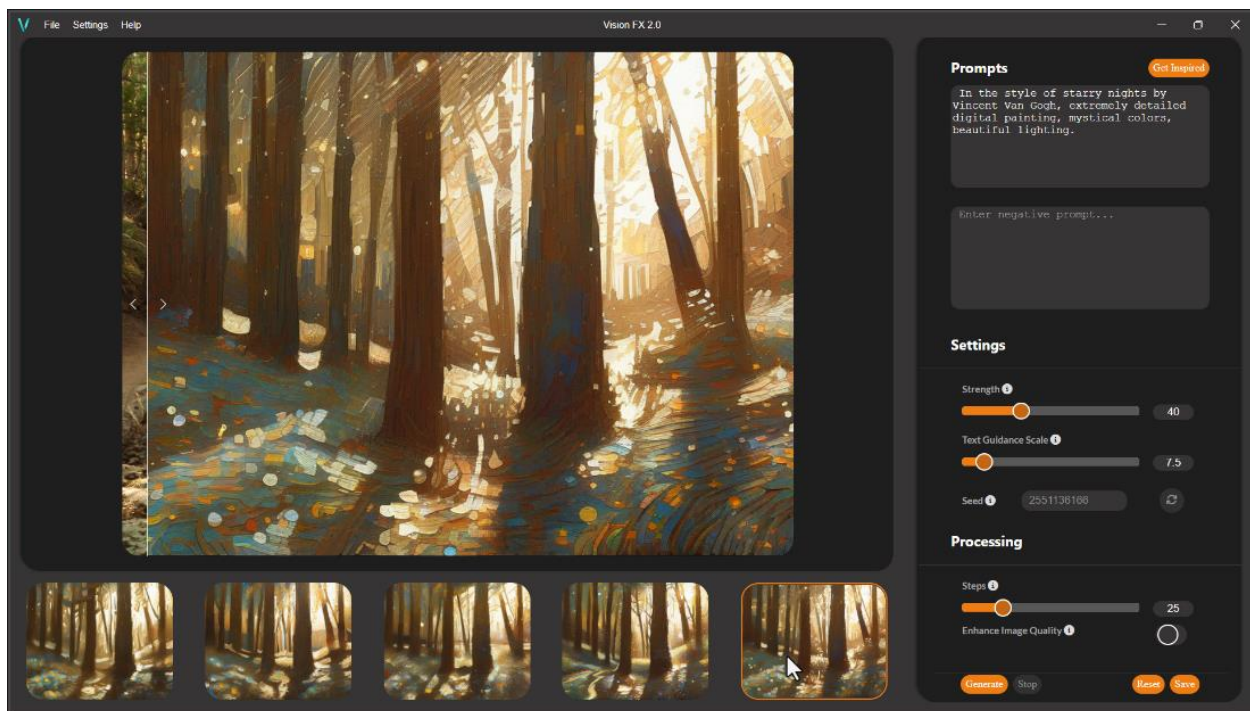
The Steps slider controls quality and speed. A high number of steps will produce slightly better results but will take longer.

When first generating options, leaving Enhance Image Quality disabled helps reduce processing time.

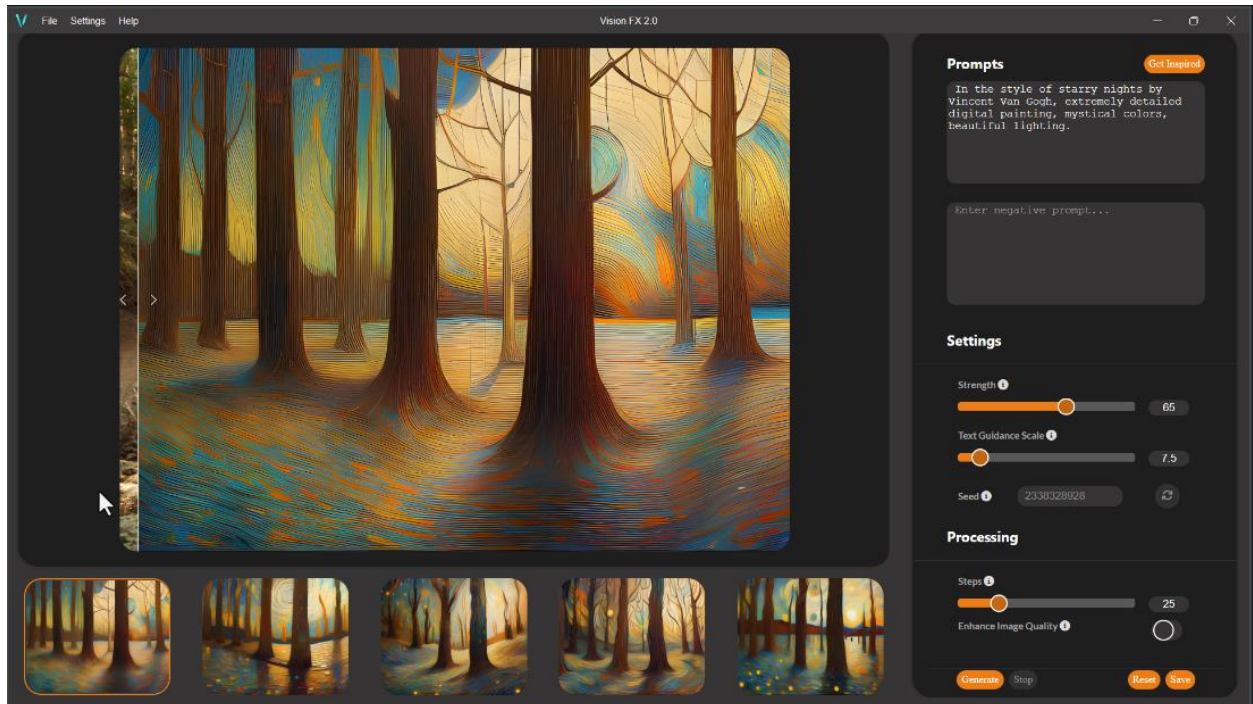
Clicking Generate invokes the AI engine to start producing the five options. This is a resource-heavy operation, so it might take some time to complete all five.

As options are generated, the slider on the image can be used to compare the AI version to the original, and I can click each option to see which I like best.

With these relatively low strength and text guidance values, the trees and overall colors and tone in each option stay pretty true to the original, while the AI engine has added the Van Gogh style specified in the prompt.

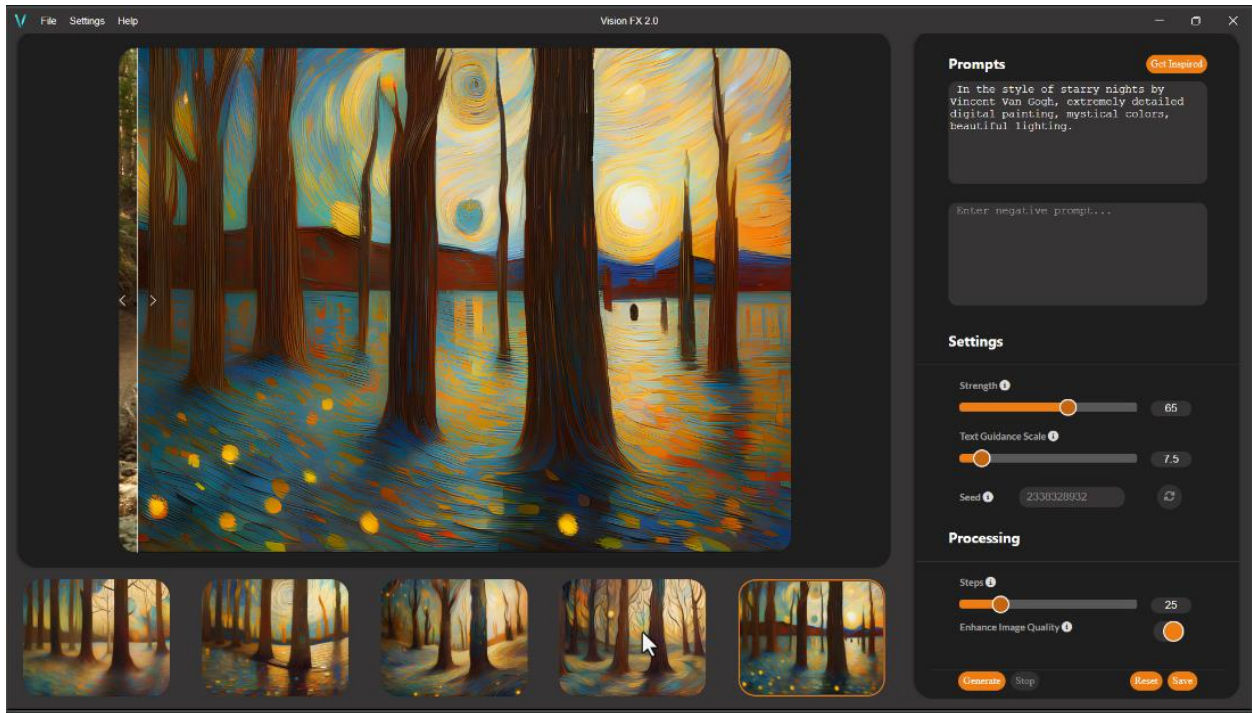


If I don't love any of the results, I can reset the values, which also resets the seed number, and this time I'll keep the default strength value. Clicking Generate again will produce five new options based on the default settings.

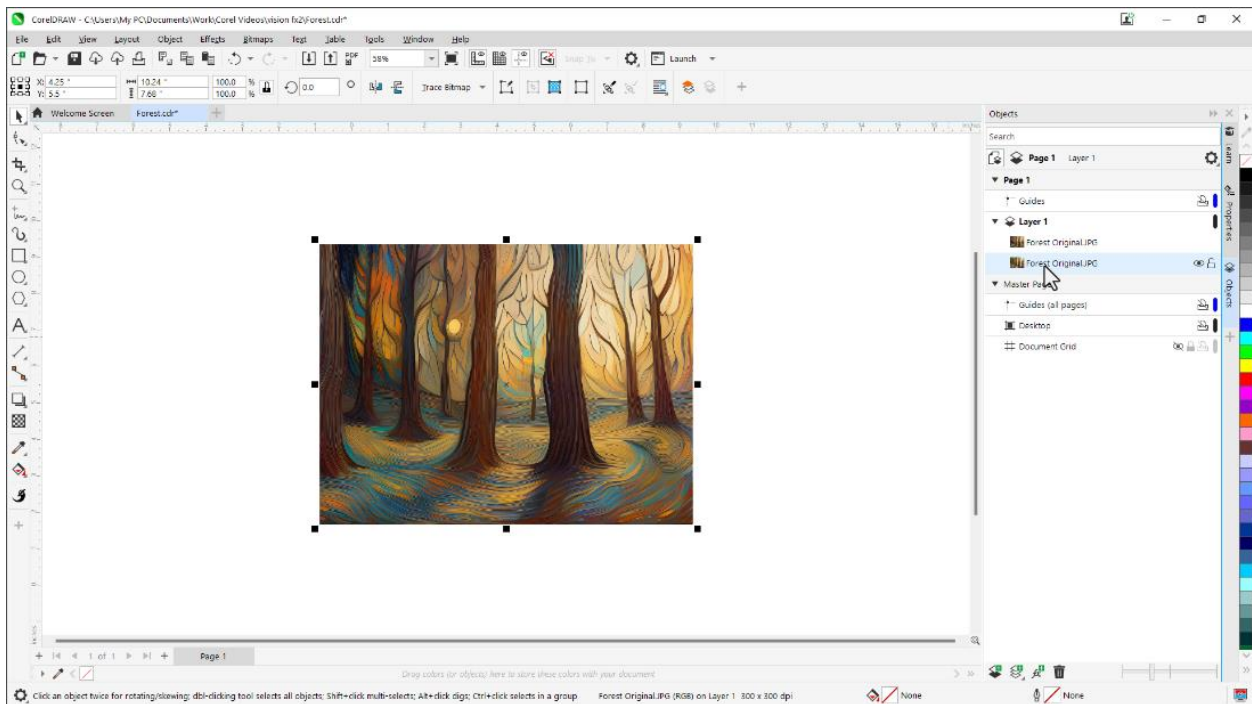


I can only choose only one of these to save, but if I wanted to get the same set of five results again, I would first note these exact settings, including the seed value, to run again later.


Let's say I like result #4, and want to see an enhanced version. With that option selected, I'll turn on Enhance Image Quality and generate again. This recreates this option and any options that come after, though I could click Stop after I get the one I want. The enhanced results are a bit crisper and smoother than their originals.



To bring this result back into CorelDRAW, I'll select it and click Save. This closes the plugin and brings me back to CorelDRAW, where the duplicated bitmap has been replaced with the AI result I chose.

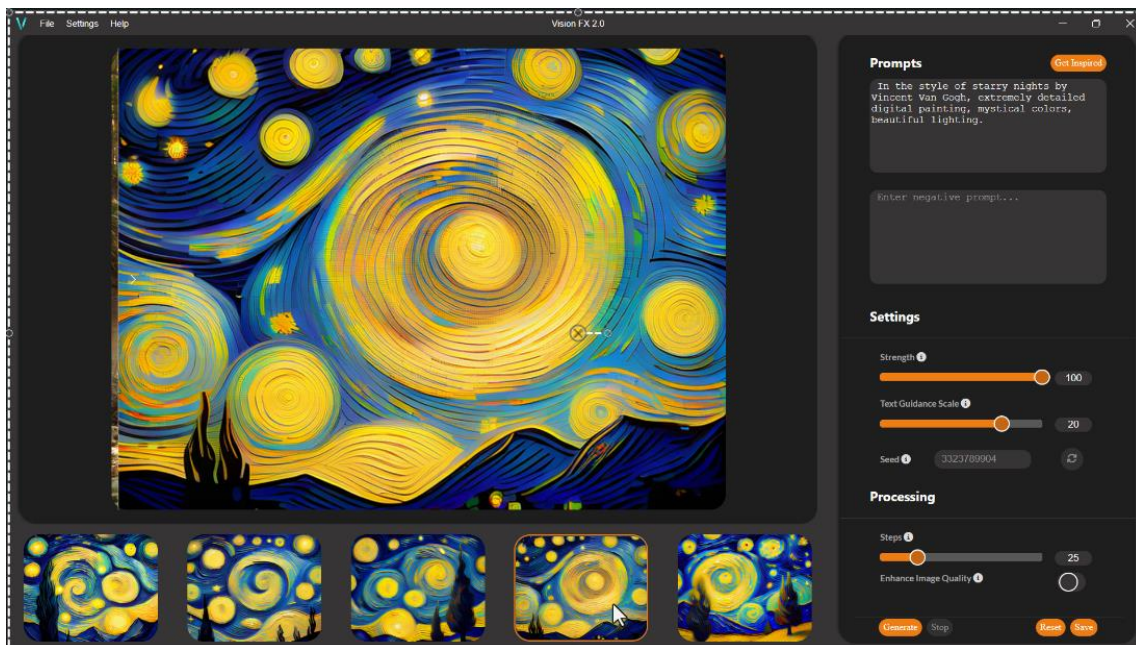


To see how the results would look with different values, I'll make another copy of the original image, then activate the plugin again. To get these results, I disabled Enhanced Image Quality to save time, and tried higher strength and text guidance values. While the results are still based somewhat on the original forest image, each option looks less like the original and more like a Van Gogh painting.

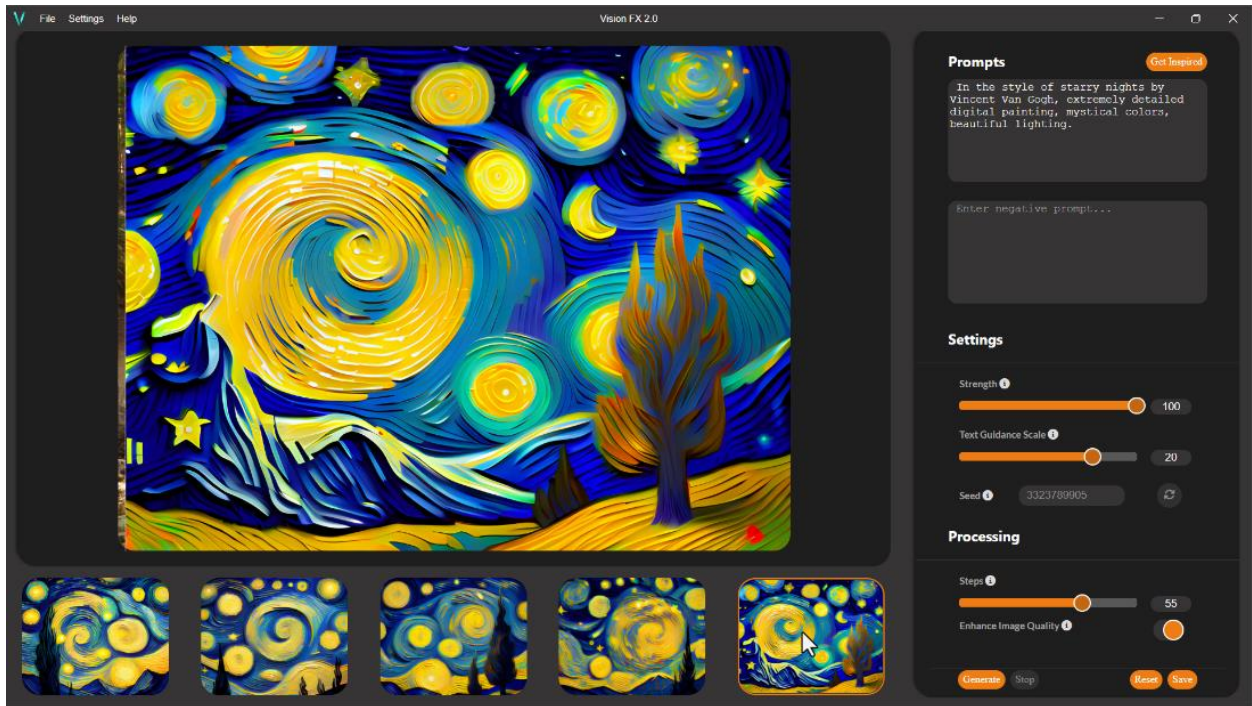


I'll save this result and return to CorelDRAW again.

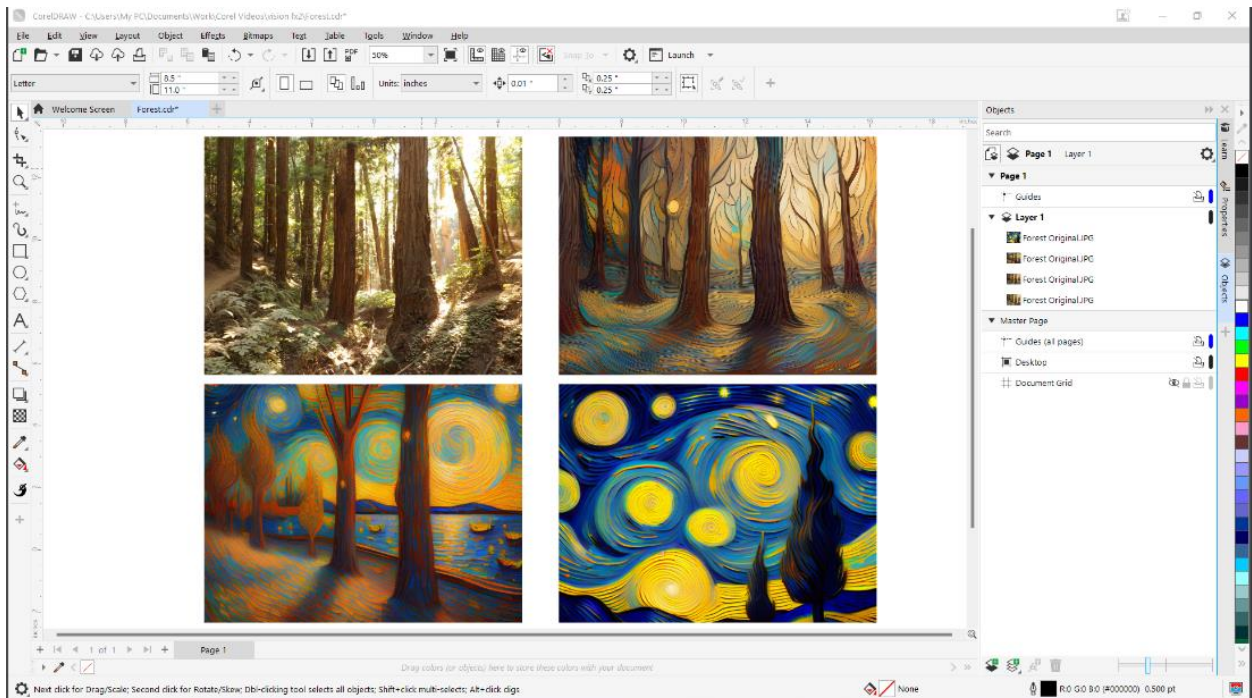
Trying again on a third copied image, I get these results, using maximum strength and higher text guidance. These images have less connection to the forest, and are more like Starry Night. High strength is also a great way to generate an AI image from scratch, as we'll see again a bit farther on.



To generate a couple fresh results at a higher resolution, I'll change this in the settings, increase steps, select the fourth result, and generate again. These two new results took a bit longer, but have tighter strokes and finer detail.



After saving the result I want, here are the three Vision FX-produced images in CorelDRAW, spaced apart in order to compare with the original.



Negative prompts

Now let's look at negative prompts, which can be used to reinforce what you're trying to achieve with your prompt terms, while also defining any objects, elements, styles, or moods that you don't want to see in your results.


I have this beach sunset image open in Paintshop Pro.

To preserve the original Background layer, I'll right click and choose Duplicate.

The plugin can be found in Effects / Plugins / VisionFX 2 – VisionFX 2.

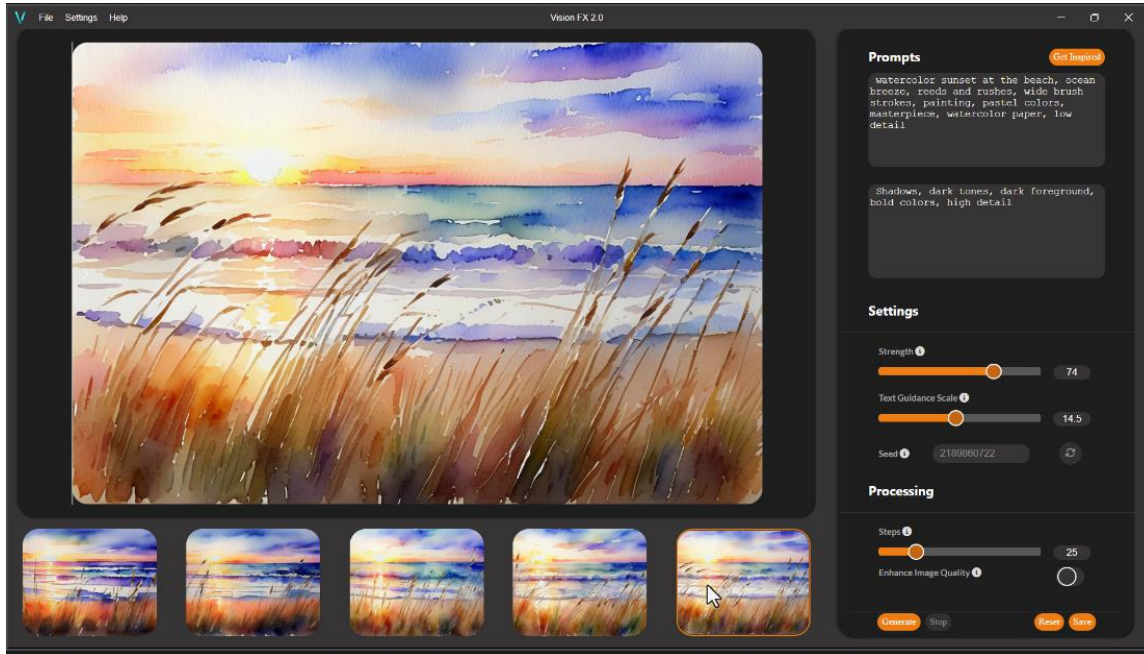
To generate some watercolor results, I entered a prompt about a watercolor sunset painting with low detail and pastel colors. The best practice is to be ultra-specific with your prompts, and start with a few terms that summarize the end result you're looking for. Subsequent terms should mainly contain adjectives that describe or enhance the style description.

Bumping up the strength and text guidance produced these results.

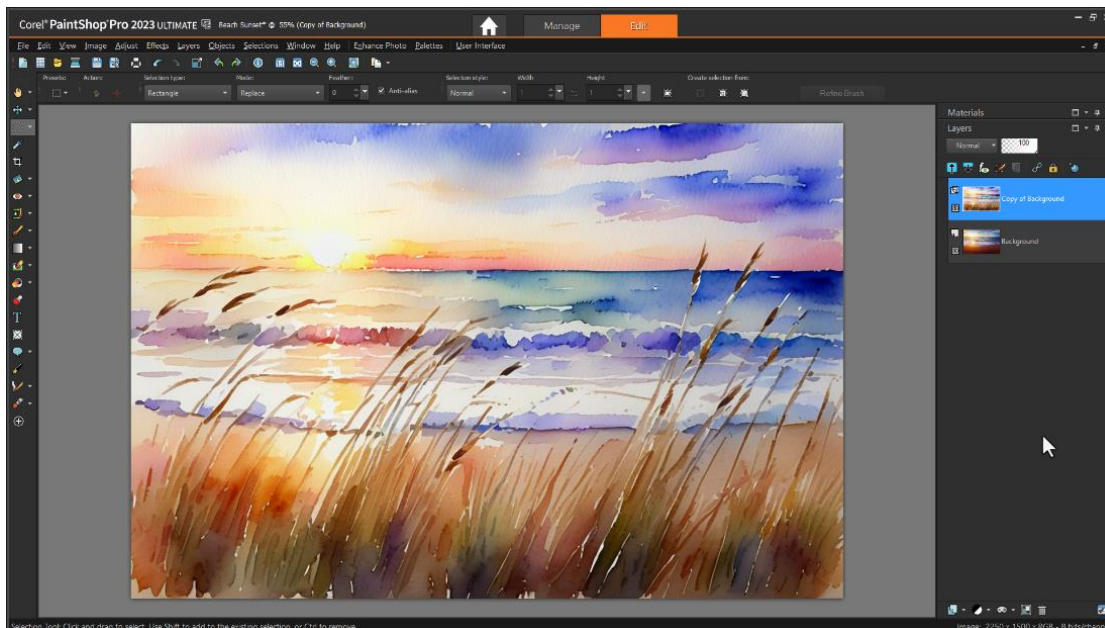


I can make some minor changes by adding negative prompts, specifying elements that should not be included. To try and brighten things, I'll choose the third option to regenerate the last few results. For negative prompts, I'll add:

I'll generate again, keeping the other settings as they are, to produce brighter and more uniformly-toned results.




With this option selected I'll save, which brings me back to PaintShop Pro. The duplicated layer is replaced with the AI image, which I can hide to see that the original photo is still there.



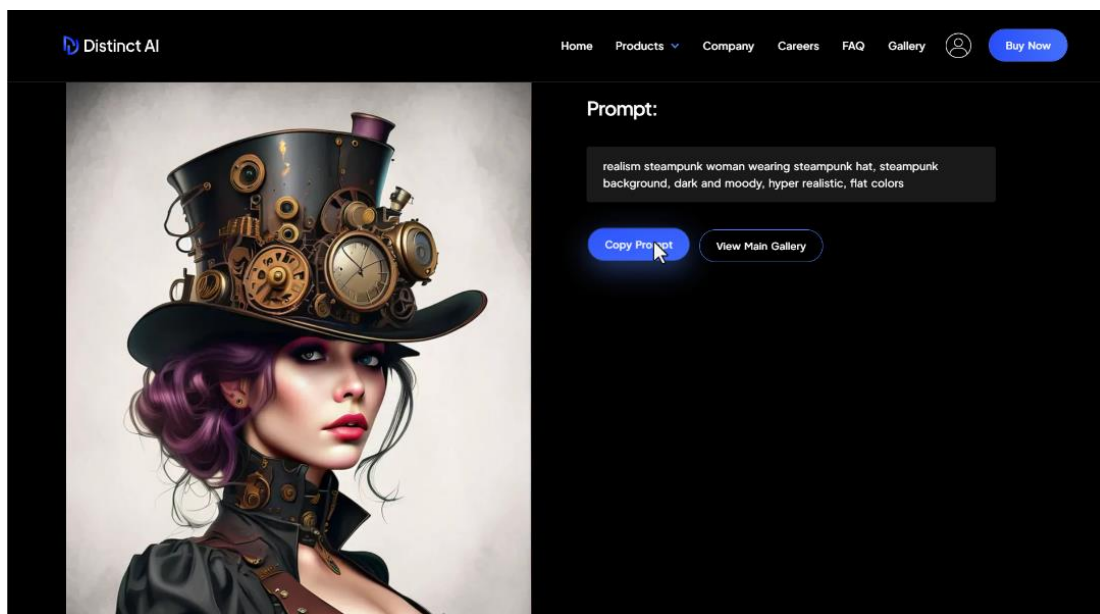
Create a character portrait

To create a portrait, I could start with a photo as we've seen, but to generate ideas from scratch, I can start with an empty canvas.

I'm using PaintShop Pro again, and my starting image is a simple gradient fill. And I can now access Vision FX directly from the Effects menu.

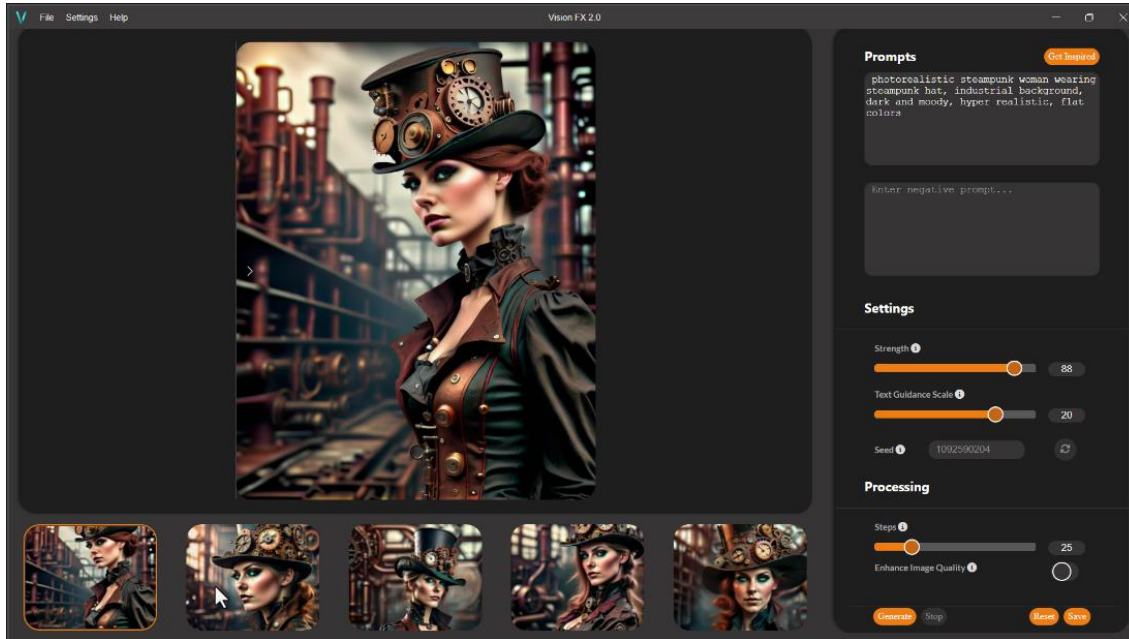


I want to get something similar to this example from the Vision FX gallery, so I'll copy its prompt, and paste it into Vision FX.



I'll also make a few changes to the prompt, including adding photorealistic and changing the background style.

Strength is high, but low enough to incorporate the gradient in the results.



You can even bring your pets into the fun. Here's my dog at home, as a cute Pixar 3D animated character with atmospheric lighting.

