

Cochise College Center for Lifelong Learning

Media Generation with AI

Creating Images, Music, and Video with AI Tools

George Self

Fall 2026

Email me at: contact@georgeteaches.com

Visit: georgeteaches.com

Visit: cochiseai.com

Introduction

Welcome to **Media Generation with AI**.¹ Not long ago, creating a polished image, a piece of music, or a short video required years of training or an expensive professional. Today, AI tools let anyone do all three with nothing more than a web browser and a clear description of what they want.

That does not mean the results are always perfect, or that these tools replace human creativity. What it means is that the gap between having an idea and being able to express it visually or musically has never been smaller. In this class, you will close that gap for yourself.

What You Need

- A device with a web browser and internet access
- A Google account (Gmail) — you will sign in during class
- Go to <https://gemini.google.com> for image and video work
- Go to <https://labs.google/fx/tools/music-fx> for music (no separate login needed)

Learning Goals

By the end of this session, you will:

- Generate original images from text descriptions using Google Gemini
- Write effective prompts that produce images matching your intent
- Create AI-generated background music for a real project using MusicFX
- Describe the current state of AI video generation and its practical limits
- Make informed decisions about copyright and disclosure when using AI-generated media

Big Idea

AI media tools do not replace your creativity — they give you a faster way to express it. The skill that matters most is not technical. It is learning to describe what you want clearly enough that the AI can follow your lead.

¹Written by George Self with assistance from Claude. It was formatted with $\text{\LaTeX}$.

Module 1: AI Image Generation

Think about the last time you needed an image for something: a church bulletin, a club newsletter, a flyer for a neighborhood event, a greeting card for a friend. You probably searched online and settled for something close enough, or went without. AI image generation changes that. You describe the image you want in plain English, and the AI creates it from scratch in seconds.

We will use Google Gemini for image generation today. Gemini is built into your Google account, which means no new signup and no new password. Open your browser, go to <https://gemini.google.com>, and sign in with your Google account.

How Image Generation Works

Gemini has been trained on an enormous number of images paired with descriptions. When you type a prompt, it uses that training to generate a new image that matches your description. It is not searching the internet for an existing image — it is creating something new.

That means two important things. First, the image belongs to you to use (with some limits we will discuss in Module 3). Second, if the first result is not quite right, you can describe the change you want and Gemini will generate a new version.

Writing a Good Image Prompt

The quality of your image depends almost entirely on how clearly you describe it. A vague prompt produces a generic result. A specific prompt produces something you can actually use.

A useful image prompt answers four questions:

- **What is the subject?** What is in the image?
- **What is the setting or background?** Where does the scene take place?
- **What is the style?** Realistic photograph, watercolor, illustration, oil painting?
- **What is the mood or tone?** Warm, professional, whimsical, dramatic?

You do not need to answer all four every time, but the more detail you give, the more useful the result.

Activity 1.1: Your First Image

Let's create an image for a real purpose. Think of something you actually need: a photo for a newsletter, an illustration for a class, a landscape for a desktop background, or an image to accompany a story.

1. Open <https://gemini.google.com> and sign in

2. Click the image generation option or type your prompt directly in the chat
3. Describe your image using the four-question framework above
4. Review the result. If it is not quite right, type a follow-up description of what to change
5. Save an image you are happy with by right-clicking and choosing **Save image as**

Discussion: What worked well in your prompt? What did you have to adjust?

Activity 1.2: Refining With Follow-Up Prompts

One of the most useful skills in AI image generation is the follow-up prompt: using the conversation to refine a result rather than starting over from scratch.

1. Look at the image you generated in Activity 1.1
2. Identify one thing you would change: the lighting, the style, the background, a detail in the subject
3. Type a follow-up prompt describing only the change you want. For example: Make the lighting warmer and more golden OR Change the background to a desert landscape
4. Compare the new result to the original

Discussion: Is it easier to refine an existing image or start with a better prompt?

Module 1 Checklist

1. Signed into Gemini at gemini.google.com
2. Generated at least one original image from a text prompt
3. Used a follow-up prompt to refine the result
4. Can describe the four elements of an effective image prompt

Module 2: AI Music Generation

Images are the most visible use of AI media tools, but music generation may be the most surprising. You can describe the kind of music you want — the mood, the

instruments, the tempo — and have a completely original piece ready in about thirty seconds. No musical training required.

We will use Google MusicFX for this module. Unlike Gemini, MusicFX does not require a separate login. Open a new browser tab and go to <https://labs.google/fx/tools/music-fx>.

What MusicFX Does

MusicFX generates short instrumental tracks from text descriptions. You describe the style, mood, and instruments you want, and it produces a loopable audio clip you can download and use. Common uses include:

- Background music for a slideshow or home video
- Ambient music for a podcast or presentation
- A custom ringtone or notification sound
- Background music for a social media post or video

Writing a Music Prompt

Music prompts work differently from image prompts. Rather than describing a scene, you are describing a feeling and a sound. Useful elements to include:

- **Genre or style:** jazz, classical, lo-fi, ambient, folk, blues
- **Instruments:** piano, acoustic guitar, strings, light percussion
- **Mood:** calm, energetic, melancholy, uplifting, professional
- **Tempo:** slow, moderate, upbeat
- **What to avoid:** specify `no lyrics` or `no drums` if needed

Activity 2.1: Background Music for a Project

Think of a slideshow, video, or presentation where background music would help. This could be a family photo slideshow, a church presentation, a club meeting intro, or anything else you have in mind.

1. Go to <https://labs.google/fx/tools/music-fx>
2. Write a prompt describing the music you want. Include genre, instruments, mood, and tempo
3. Generate the track and listen to the result
4. If the result is not right, adjust one element of your prompt and generate again

5. Download a track you like using the download button

Example prompt to try: Soft instrumental background music, light jazz with piano and acoustic bass, slow tempo, calm and professional, no lyrics

Discussion: How does changing the genre or instrument list change the feel of the result?

Module 2 Checklist

1. Opened MusicFX at labs.google without a separate login
2. Generated at least one music track from a text prompt
3. Adjusted the prompt to change the result
4. Downloaded a track for potential use in a real project

Module 3: AI Video and Responsible Use

AI Video Generation

AI video generation is the newest and least mature of the three media types we are covering today. The technology is improving rapidly, and what is possible now is likely to look quite limited by the time you need to use it in a real project.

Currently, Google Gemini can generate short video clips (a few seconds to about a minute) from text descriptions using a model called Veo. The results can be impressive for simple scenes but struggle with complex motion, human faces, and anything requiring precise physical accuracy.

A Practical Reality Check

AI video is worth knowing about and worth experimenting with, but for most practical purposes — club newsletters, family slideshows, community announcements — AI image generation combined with a simple slideshow tool will serve you better today. Check back in a year. The gap is closing quickly.

Activity 3.1: Exploring AI Video

1. In Gemini, type a prompt asking for a short video. For example: Create a short video of a sunset over a desert landscape with saguaro

cactus

2. Review the result. Notice what it does well and where it falls short
3. Try a simpler scene if the first result is not what you expected

Discussion: How does this compare to what you expected? What would you need from a video tool that this does not yet provide?

Using AI Media Responsibly

Creating media with AI is easy. Using it responsibly takes a little more thought. Three areas deserve your attention.

Copyright and ownership. Images, music, and video you generate with AI are generally yours to use for personal and non-commercial purposes. The rules for commercial use — selling a product, using it in advertising, publishing it in a paid publication — vary by tool and are still evolving legally. When in doubt, check the terms of service for the tool you used.

Disclosure. If you use AI-generated media in a public context, say so. A note like “Image created with AI” in a newsletter or presentation is good practice and increasingly expected. It is also a conversation starter: many people will be curious about how you did it.

Accuracy and representation. AI image generators can produce images of real places, people, or events that never happened. Be careful about generating images that could be mistaken for photographs of real situations, and never use AI-generated images to misrepresent facts or people.

Module 3 Checklist

1. Generated or viewed an example of AI video output
2. Can describe the current practical limits of AI video generation
3. Can state the three responsible use principles: copyright, disclosure, and accuracy

Tools to Explore on Your Own

The tools we used today are a starting point. Here are others worth trying when you are ready to go further.

Image Generation

- **Adobe Firefly** (<https://firefly.adobe.com>) — Adobe’s image generator, integrated with Photoshop and Canva. Strong for professional design work.
- **Microsoft Copilot** (<https://copilot.microsoft.com>) — Includes image generation powered by DALL-E. Works with your Microsoft account if you have one.

- **Canva** (<https://canva.com>) — Combines AI image generation with a full design tool. Excellent for newsletters, flyers, and social media graphics.

Music Generation

- **Suno** (<https://suno.com>) — Generates full songs with lyrics and vocals from a text prompt. Free tier available with account signup.
- **Udio** (<https://udio.com>) — Similar to Suno, with strong results for a variety of musical styles.

Video Generation

- **Runway** (<https://runwayml.com>) — One of the leading AI video tools for creative work. Limited free tier.
- **Pika** (<https://pika.art>) — Generates short video clips from images or text. Free tier available.

Prompt Templates

Use these as starting points and adapt them to your own needs.

Image: Flyer or announcement illustration

A warm, welcoming illustration for a [type of event] flyer. [Setting or location]. Watercolor style, soft colors, no text in the image.

Image: Portrait or group scene

A friendly, realistic illustration of [subject]. [Setting]. Natural lighting, professional but approachable mood.

Image: Landscape or place

A [style: photograph / painting / illustration] of [location or scene]. [Time of day]. [Mood: peaceful / dramatic / golden hour].

Music: Slideshow background

Soft instrumental background music for a [occasion] slideshow. [Genre] with [instruments]. Slow to moderate tempo, no lyrics, [mood].

Music: Upbeat background

Upbeat, cheerful instrumental music. [Genre]. [Instruments]. Moderate tempo, no lyrics, suitable for a [context: community event / family video / club presentation].