

# Google DeepMind Lyria 3: Music Generation Model Outputs 48kHz Stereo with Real-Time Steering and SynthID Watermarking

Kabui, Charles

2026-03-03

---

[Read at ToKnow.ai](#)

---

The image is a screenshot of a dark-themed announcement page for Google DeepMind Lyria 3. At the top left, the title "Google DeepMind Lyria 3" is displayed in white and light blue. Below it, the subtitle "AI music generation with real-time steering and SynthID watermarking" is shown in a smaller white font. The main content is organized into three vertical columns, each with a colored header and descriptive text below it. The first column has a red header "48kHz" and text "Studio-quality stereo" and "Full arrangements from scratch". The second column has a yellow header "<2s" and text "Real-time latency" and "Steer music as it plays". The third column has a blue header "SynthID" and text "Inaudible watermark" and "Survives compression". At the bottom left, there is a line of text: "Text, image, or audio → full multi-instrumental tracks with vocals". At the bottom right, the "ToKnow.ai" logo is visible. In the top right corner of the page, there is a faint, stylized graphic of a house or a similar structure.

**Google DeepMind Lyria 3**

AI music generation with real-time steering and SynthID watermarking

- 48kHz**  
Studio-quality stereo  
Full arrangements from scratch
- <2s**  
Real-time latency  
Steer music as it plays
- SynthID**  
Inaudible watermark  
Survives compression

Text, image, or audio → full multi-instrumental tracks with vocals

ToKnow.ai

Figure 1: Google DeepMind Lyria 3

Google DeepMind released Lyria 3, its most advanced music generation model. The model generates 30-second tracks in studio-quality stereo (48kHz) from text, image, or audio inputs. It does not piece together loops. It generates full musical arrangements from scratch, including vocals and multi-instrumental tracks. Lyria 3 is now integrated into the Gemini app. Describe a mood, genre, or set of instruments, and it outputs a finished audio file. Music generation is hard because it is continuous and multi-layered: melody, harmony, rhythm, and timbre all need to stay coherent from the first second to the last.

The real shift is the Lyria RealTime API. Traditional models work like a jukebox: input a prompt, wait for a file. Lyria RealTime generates audio in 2-second chunks as a live stream. Users can steer style, tempo, and instrumentation while the track plays, with under 2 seconds of delay. The Music AI Sandbox builds on this, transforming a hum into a full orchestral arrangement or turning chord progressions into a vocal choir. Every track includes a SynthID watermark embedded in the audio waveform, inaudible to humans but detectable by software even after MP3 compression or analog re-recording. As AI-generated music floods streaming platforms, permanent attribution solves a real trust problem for creators and listeners. Lyria 3 competes directly with Suno and Udio but brings multimodal input and Google's infrastructure advantage.

Sources:

- [Google DeepMind Lyria Technical Details](#)
- [MarkTechPost: Google DeepMind Lyria 3](#)

---

*Disclaimer: For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation. **Read more:** [/terms-of-service](#)*