

# The Human Music Show: Crowd-Prompting as a Playable Social Interface for Generative Live Performance

MANASWI MISHRA, MIT Media Lab, USA  
ANTONIS CHRISTOU, MIT Media Lab, USA

Additional Key Words and Phrases: Generative Music, AI, Interactive Art, Live-Code Music, Community-driven Performance

## 1 Program Notes

*The Human Music Show* is a satirical, collective AI music performance, as well as commentary on the hyper-commercialization of musical artifacts featuring audience participation that melds the traditions of live-coding, DJing, and instrumental virtuosity. Prompting, which has emerged as the dominant control paradigm in commercial AI music tools, is reimagined here away from its usually solipsistic and solitary context, and instead as a way to integrate the collective will of the crowd and greater musical community into one cohesive musical whole. As part of our novel performance paradigm, we introduce *CrowdPrompting*, where audience feedback via software is turned live into generated AI samples, which are then chopped and sampled by the human performers and hosts of the show.

## 2 Project Description

How can the commercial, generative AI music paradigm of isolated song-making and soulless craft be inverted into something that an entire community can participate in? This is the aim of *The Human Music Show*, which reframes prompting as a playable social interface for generative live performance.



Fig. 1. Excerpts from a *Human Music Show* Performance.

Authors' Contact Information: Manaswi Mishra, MIT Media Lab, Cambridge, USA; Antonis Christou, MIT Media Lab, Cambridge, USA.



This work is licensed under a Creative Commons Attribution 4.0 International License.

NIME '26, June 23–26, 2026, London, UK

© 2026 Copyright held by the owner/author(s).

Blurring the line between satirical AI game show and participatory musical experience, the humans, who double as performers and actors, work in conjunction with the AI-generated "host" to guide the audience through the proceedings of this imaginary show that takes influence from the imagery of classic eighties media.

The piece was created with a fundamental desire to integrate the ethos and tradition of existing musical practices, particularly, those of live-coding and DJing, to better engage the audience and community even despite the use of often alienating commercial AI tools. Drawing inspiration from the newly emergent practice of "prompt jockeying" [3], we comment on the increasing trend to offload responsibility and control to unspecified and poorly understood AI systems.

In this case, *The Human Music Show* seems at first glance to relinquish decision making power even further to machines, featuring an almost entirely AI-generated backbone, complete with AI "host", visuals, and even theme song. In reality, however, the humans exhibit the ultimate power and control to shape the performance as a collective.

The overall structure features the audience participating in a series of three rounds where they, using their phones connected to a locally running server, "bet" on generative AI prompts that they would like to hear. The resulting weights are then fed into a satirical slot machine interface that "selects" the ultimate prompt. This role in past performances was embodied by the 2025 NIME and musical arcade machine, LIMITER [1], but will for this iteration be accomplished digitally. "Gambling" for prompts in this manner is used as an interaction paradigm to poke fun at the language of profit and efficiency that drives the consumer AI music industry, with all its devaluation of craft and process in favor of instantaneous dopamine hits and infinite content.

The resulting music is instantly fed into a Google Lyria / Magenta RT (open weights alternative) generative streaming model [5] and briefly played back to the audience, allowing them to experience the direct musical result of their choices. The human performers then assume control, taking the generated sample and chopping, splicing, and manipulating it into a new sonic form using the live-coding language Strudel [4], as well as a real-time prompt steering interface. This workflow is enabled by this new interface for generative music live-coding built by the authors as a custom fork of Strudel that incorporates generative music streaming models with a variety of interpretable controls, such as filtering, looping, temperature sampling, and more.

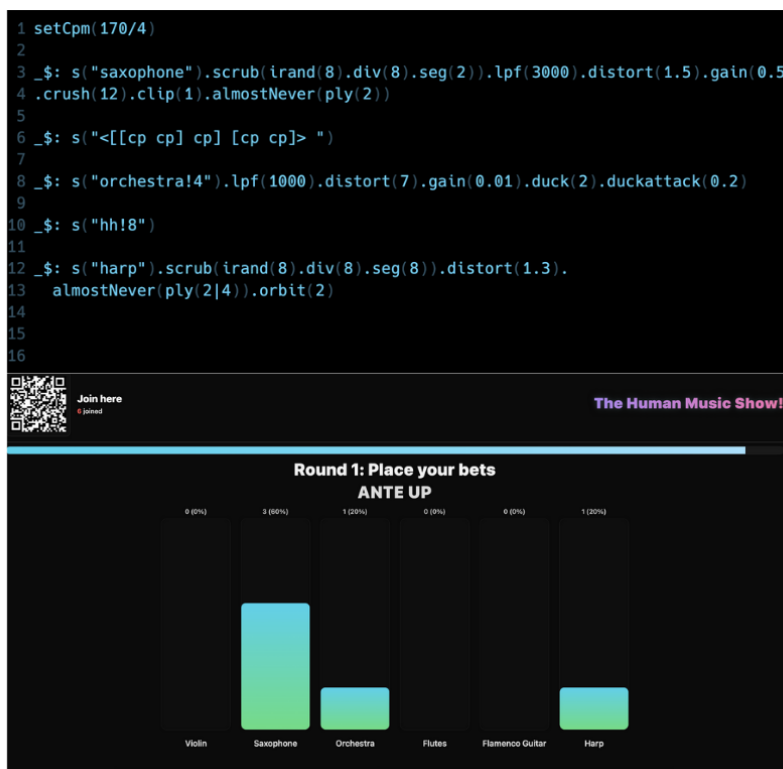


Fig. 2. The Strudel live coding and crowd-prompting "betting" interfaces

The rounds are structured according to increasing levels of musical abstraction going from "instruments" (e.g violin, saxophone, etc.) to "genres" (hyperpop, djent, classical) to "unusual sounds and prompts" that push the boundaries of the model (keyboard mashing, nonsense words, etc.)

The music gets increasingly frantic, up-tempo and unhinged, culminating in a final phase where the audience scans a QR code to contribute prompts that are then curated by the human performers and generated in the aforementioned custom generation pipeline that allows for the musicians to generate a final audience-inspired collective musical tapestry before slowly fading out. Our proposed performance paradigm introduces prompt based sample generation within a rich history of participatory code based music practice [2].

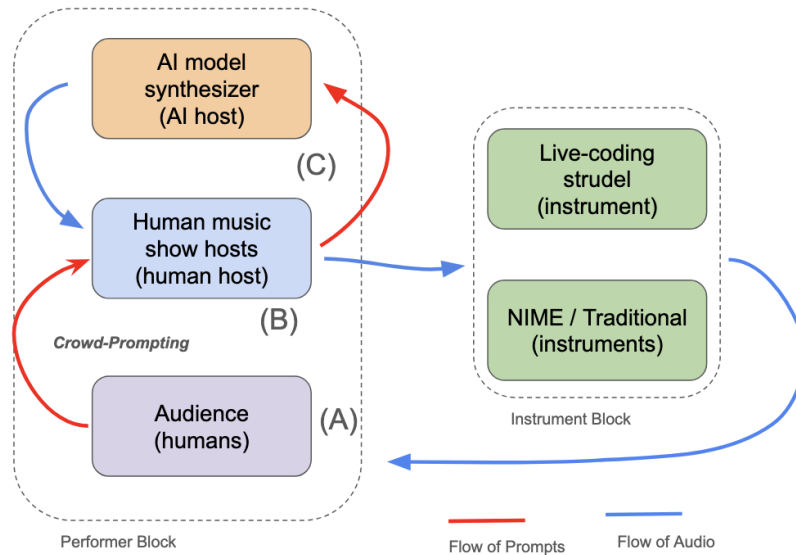


Fig. 3. Musical Agency Loops exhibited during performance

### 3 Technical Notes

The performance has a duration of approximately 15 minutes, and the essential requirements from the venue for performing *The Human Music Show* are the following:

- A large display, either projection or LED
- Sound system suitable for electronic music performance (i.e with a subwoofer, etc.)
- Relatively fast WiFi (100 Mbps or more) for communicating with audio streaming chunks over API and alternatively on-device Magenta RT open weights model with self-provided internet connection.
- An audience of around twenty individuals at a minimum that interacts with phones.
- A standing table for the performers
- A power strip and/or close proximity to at least two outlets for power
- Dim house lights to increase screen visibility, but preferably with a spotlight on the human performers

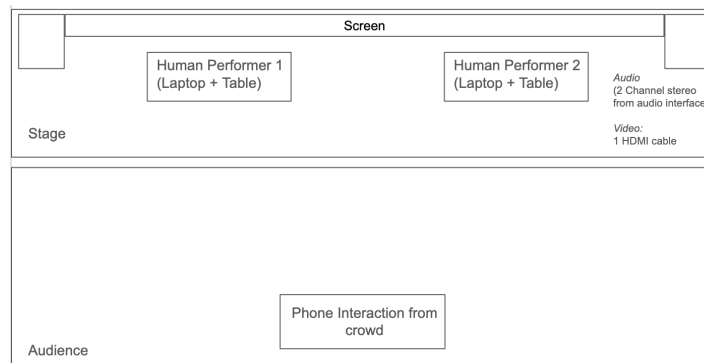


Fig. 4. Stage Diagram

The performers will bring:

- Costumes, for themselves as both performers and actors on the show
- One four input / four output audio interface
- Audio Cables, in particular two TRS for stereo audio output and one HDMI for sending visuals
- Two MacBook Pros, for running the various audio and AI systems

#### 4 Media Links

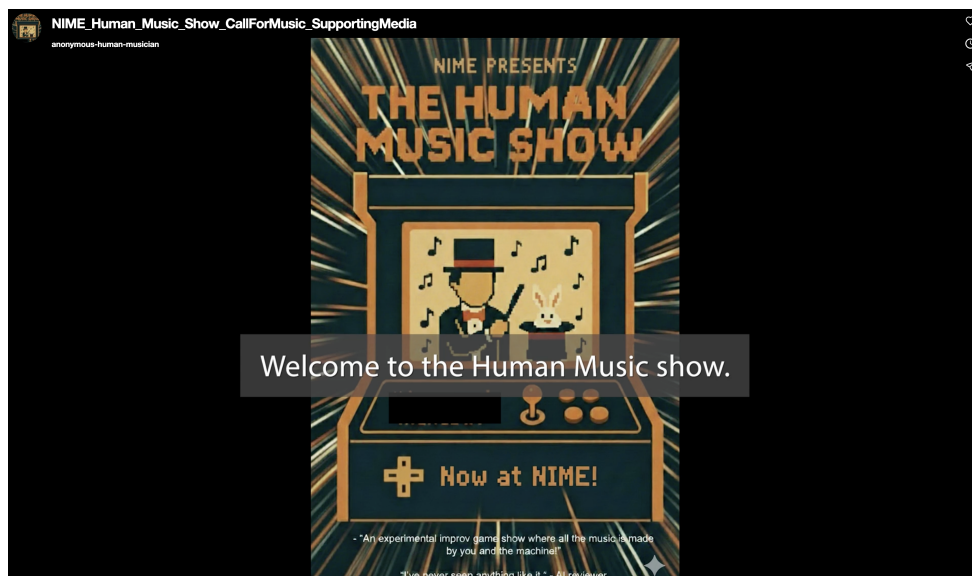


Fig. 5. Demo video on Vimeo (click the image): <https://vimeo.com/1164587954>

#### 5 Ethical Standards

This research does not involve any human or animal experimentation or research. No personal data from the audience was collected or stored. All AI models used are free and open source. The code used is built on Apache 2.0 license and audio samples generated from the Magenta RT Open weights are licensed under Creative Commons 4.0, but the authors acknowledge the environmental costs related to the training of these systems.

#### Acknowledgments

The authors would like to thank Tod Machover, as well as the entire MIT Media Lab Opera of the Future group for their constant support. This work was made possible by the MIT Media Lab 99Fridays event coordinators, who allowed for the original performance to occur.

#### References

- [1] Antonis Christou. 2025. LIMITER: A Gamified Interface for Harnessing Just Intonation Systems. In *Proceedings of the Proceedings of the International Conference on New Interfaces for Musical Expression*. <https://doi.org/10.5281/zenodo.15698896>
- [2] Michael Mulshine, Ge Wang, Chris Chafe, Jack Atherton, Terry Feng, and Celeste Betancur. 2023. WebChuck: Computer Music Programming on the Web.. In *NIME*.
- [3] Jordi Pons, Zack Zukowski, Julian D. Parker, CJ Carr, Josiah Taylor, and Zach Evans. 2025. Music and Artificial Intelligence: Artistic Trends. arXiv:2508.11694 [cs.CY] <https://arxiv.org/abs/2508.11694>
- [4] Felix Roos and Alex McLean. 2023. Strudel: Live Coding Patterns on the Web. In *Proceedings of the 7th International Conference on Live Coding*. <https://doi.org/10.5281/zenodo.7842142>
- [5] Lyria Team, Antoine Caillon, Brian McWilliams, Cassie Tarakajian, Ian Simon, Ilaria Manco, Jesse Engel, Noah Constant, Yunpeng Li, Timo I. Denk, Alberto Lalama, Andrea Agostinelli, Cheng-Zhi Anna Huang, Ethan Manilow, George Brower, Hakan Erdogan, Heidi Lei, Itai Rolnick, Ivan Grishchenko, Manu Orsini, Matej Kastelic, Mauricio Zuluaga, Mauro Verzetti, Michael Dooley, Ondrej Skopek, Rafael Ferrer, Savvas Petridis, Zalán Borsos, Aaron van den Oord, Douglas Eck, Eli Collins, Jason Baldrige, Tom Hume, Chris Donahue, Kehang Han, and Adam Roberts. 2025. Live Music Models. arXiv:2508.04651 [cs.SD] <https://arxiv.org/abs/2508.04651>