THE CREATIVE MEMORY ENGINE: REALIZATION-ORIENTED AI FOR INDEPENDENT MUSICIANS

Todd Anstis
StudioOps / Hewn Ventures, Fort Worth, TX, USA
todd.anstis@hewnventures.com

ABSTRACT

Independent musicians accumulate more voice memos, demos, lyric fragments, and references than they can keep available for finishing. Generative systems reduce the cost of producing new material, but they do not restore connection to work an artist has already begun. This late-breaking demo presents the Creative Memory Engine, a working prototype of realization-oriented AI: systems designed to increase the share of a musician's own unfinished possibilities that can be perceived, judged, and completed, while preserving human authorship. StudioOps, the host system, ingests heterogeneous personal artifacts into an explorable semantic constellation. The representation surfaces thematic proximity, recurrence, and cross-modal relationships; the artist inspects, accepts, or rejects suggested links; and selected context can be carried toward a digital audio workstation. The live session walks through ingestion, constellation navigation, bloom-style clusters, and an authorship-preserving handoff. This is a design case and prototype, not a completed user study. The live demo is organized around six affordances: semantic perception, longitudinal memory, judgment support, prioritization, synthesis, and execution continuity. This extended abstract is not an ISMIR 2026 proceedings paper.

1. INTRODUCTION

Most independent songwriting practice is not a shortage of ideas. It is a surplus of fragments. Voice memos, abandoned demos, lyric drafts, and reference recordings pile up across phones, folders, and DAW sessions. Capture tools externalize memory, which is valuable, but they often turn an active idea into a stored file. Reuse then depends on remembering that the file exists, finding it, and reconstructing why it mattered: idea → capture → archive → forgetting → retrieval → reconstruction → development.

Generative music systems enlarge that possibility space while increasing the burden of comparison, organization, and commitment [1, 2]. Studies of human–AI songwriting show that musicians already struggle to steer generative models, retain authorship, and keep generated snippets in a coherent personal process [3–5]. Those systems answer a different question than the one many working artists face: not *what else could be generated*, but *what of mine is still unfinished, related, and worth returning to*.

This demo starts from a practical thesis: most artists do not have an idea problem; they have a connection problem. It instantiates a complementary objective that we call realization-oriented AI. The aim is closer to augmenting human intellect [6] and to creativity support tools that help more people create more often [7, 8] than to end-to-end song generation. StudioOps is a design case: a working prototype in which MIR-style retrieval and representation serve memory, judgment, and continuity rather than autonomous composition. An earlier talk on this direction was given at Berklee's AI Music Summit (AIMS) in June 2026.

2. REALIZATION-ORIENTED AI

We define realization-oriented AI as systems designed to increase the proportion of creator-valued possibilities that become meaningful completed outcomes. "High potential" is not an objective label assigned by a model. It is a situated judgment grounded in the artist's values and commitments. The system can supply perspective; the person retains authority over taste, purpose, and completion.

A useful construct is *creative surface area*: the ideas, memories, references, relationships, and developmental pathways meaningfully available to a creator for perception and action at a given time. Availability is situated. Material must be relevant enough to notice, interpretable enough to judge, and actionable enough to influence a decision. The claim is not that the system "knows" the artist. It is that well-designed representations can change what the artist can notice and do, in the spirit of sensemaking with external representations [9] and of computers as partners in parts of a creative process [10, 11]. Table 1 lists six affordances exercised in the demo.

These affordances complement guidelines for human–AI interaction: make capabilities inspectable, support correction, and keep the human in control of consequential decisions [12]. Agency is not a final approval button. It requires rejectable suggestions, visible provenance, and the ability to withhold sensitive material. The system should also leave room for wandering, incubation, and deliberate non-completion.

© T. Anstis. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** T. Anstis, "The Creative Memory Engine: Realization-Oriented AI for Independent Musicians", in *Extended Abstracts for the Late-Breaking Demo Session of the 27th Int. Society for Music Information Retrieval Conf.,* Abu Dhabi, UAE, 2026.

Affordance	Role in the demo
Semantic perception	Make recurrences inspectable beyond unaided recall.
Longitudinal memory	Preserve intent, provenance, and idea evolution.
Judgment support	Compare and contextualize without claiming taste.
Prioritization	Surface recurrence or readiness; the artist chooses.
Synthesis	Link fragments across media and show why.
Execution continuity	Carry selected context toward making.

Table 1. Six design affordances of realization-oriented AI.

3. THE CREATIVE MEMORY ENGINE

The Creative Memory Engine is the part of StudioOps demonstrated here: ingest an artist's own fragments, represent them in an explorable semantic field, help the artist choose what to finish, and hand selected context toward a DAW (Figure 1). The system does not generate the song.

The prototype ingests heterogeneous artifacts from a personal corpus—voice memos, lyric drafts, demo recordings, images, quotations, and notes. Embedding models construct semantic representations so text, audio-derived descriptions, and other media can occupy a shared field of meaning. Clustering and retrieval organize those artifacts by conceptual relationship rather than by folder or capture date alone. Every item remains linked to its source so a suggested relationship can be checked, rejected, or reinterpreted. Similarity is a hypothesis about the artist's own material, not a claim that two songs are the same, and not a ranking of what is "best."

The primary interface is a navigable constellation. Distance is a visual proxy for semantic and thematic proximity. A *bloom* is an emergent cluster around a selected fragment: the originating sketch plus related lyrics, memos, references, and unfinished takes. A guitar idea recorded late at night may sit beside a lyric written months later and a forgotten reference, because they share an emotional or thematic center. Named surfaces in the current prototype include Signal (the constellation), Bloom (the selected cluster), and Workbench (where lineage stays intact while a thread is developed). The point is not that AI found "similar files." The point is that relationships previously distributed across notebooks, drives, apps, and memory become perceptible together.

When a thread looks worth finishing, the artist—not the model—commits. A companion integration extends selected context into Ableton Live, the environment in which many independent artists actually develop a song. The handoff is an execution-continuity prototype: it aims to reduce the archive break between recognition and making. It is not claimed as a finished commercial plugin, and it does not write parts into the arrangement. AI contributes memory, comparison, synthesis, and continuity. The artist contributes taste, intention, commitment, and meaning.

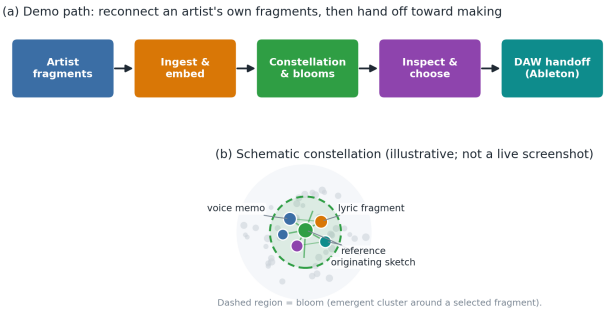


Figure 1. Live demo path from personal fragments to an authorship-preserving DAW handoff. Diagram, not a product screenshot.

4. LIVE DEMO

The session is a live walkthrough with a short personal corpus prepared in advance. Attendees will see ingestion of the author's own fragments (voice memo, lyric, guitar demo, reference); constellation update and unexpected neighbors; a bloom around one originating sketch, with inspectable provenance; accept/reject of suggested links; and handoff of selected context toward Ableton Live. The five-minute format emphasizes interaction over narration: what becomes newly available to attention, and what remains the artist's decision. Prototype limits will be stated in the room.

5. STATUS AND LIMITATIONS

StudioOps is a working prototype and a design framework, not a validated product claim. We do not report user-study statistics, finish rates, or retrieval metrics here because those studies have not been completed. Candidate future measures include forgotten material recovered and judged relevant by the creator; effort to resume interrupted work; usefulness of surfaced connections; and changes in perceived authorship over longer use. Short lab tasks will miss the phenomenon; consented longitudinal traces are more appropriate.

Semantic representations can privilege what is easily digitized, misread ambiguity, and reproduce model or corpus bias. Resurfacing can distract, overwhelm, or expose material the artist intended to leave dormant. Personal archives also create privacy and ownership risks; the public demo uses the author's own material or explicitly consented excerpts. The system should support withholding and forgetting, not only remembering.

The intended contribution to ISMIR is a working alternative to generation-centered music AI: retrieval, representation, and interaction that help independent musicians remain in relationship with more of what they have already begun. If generation has become abundant, realization—finishing work that still sounds like the artist—is the scarcer MIR problem.

6. AI USAGE STATEMENT

A large language model was used to help draft and organize this extended abstract from the author's existing manuscripts, talk notes, and design documents, including selection and wording of the related-work discussion. All technical claims were checked against those sources. No user-study results, model names, or citations were invented. The author is responsible for the entire content of this abstract, including text, figure, and references.

7. REFERENCES

[1] A. R. Doshi and O. P. Hauser, "Generative AI enhances individual creativity but reduces the collective diversity of novel content," *Science Advances*, vol. 10, no. 28, p. eadn5290, 2024.

[2] Z. Epstein, A. Hertzmann, and Investigators of Human Creativity, "Art and the science of generative AI," *Science*, vol. 380, no. 6650, pp. 1110–1111, 2023.

[3] C.-Z. A. Huang, H. V. Koops, E. Newton-Rex, M. Dinculescu, and C. J. Cai, "AI song contest: Human-AI co-creation in songwriting," in *Proc. of the 21st Int. Society for Music Information Retrieval Conf.*, Montréal, Canada, 2020.

[4] R. Louie, A. Coenen, C. Z. Huang, M. Terry, and C. J. Cai, "Novice-AI music co-creation via AI-steering tools for deep generative models," in *Proc. of the CHI Conf. on Human Factors in Computing Systems*, 2020, pp. 1–13.

[5] B. L. Sturm, O. Ben-Tal, Ú. Monaghan, N. Collins, D. Herremans, E. Chew, G. Hadjeres, E. Deruty, and F. Pachet, "Machine learning research that matters for music creation: A case study," *J. New Music Research*, vol. 48, no. 1, pp. 36–55, 2019.

[6] D. C. Engelbart, "Augmenting human intellect: A conceptual framework," Stanford Research Institute, Menlo Park, CA, USA, Tech. Rep., 1962.

[7] B. Shneiderman, "Creativity support tools: Accelerating discovery and innovation," *Commun. ACM*, vol. 50, no. 12, pp. 20–32, 2007.

[8] J. Frich, L. MacDonald Vermeulen, C. Remy, M. M. Biskjaer, and P. Dalsgaard, "Mapping the landscape of creativity support tools in HCI," in *Proc. of the CHI Conf. on Human Factors in Computing Systems*, 2019, pp. 1–18.

[9] D. M. Russell, M. J. Stefik, P. Pirolli, and S. K. Card, "The cost structure of sensemaking," in *Proc. of the INTERCHI Conf. on Human Factors in Computing Systems*, 1993, pp. 269–276.

[10] T. Lubart, "How can computers be partners in the creative process: Classification and commentary on the special issue," *Int. J. Human-Computer Studies*, vol. 63, no. 4–5, pp. 365–369, 2005.

[11] N. Davis, C.-P. Hsiao, Y. Popova, and B. Magerko, "An enactive model of creativity for computational collaboration and co-creation," in *Creativity in the Digital Age*. Springer, 2015, pp. 109–133.

[12] S. Amershi, D. Weld, M. Vorvoreanu, A. Fourney, B. Nushi, P. Collisson, J. Suh, S. Iqbal, P. N. Bennett, K. Inkpen, J. Teevan, R. Kikin-Gil, and E. Horvitz, "Guidelines for human-AI interaction," in *Proc. of the CHI Conf. on Human Factors in Computing Systems*, 2019, pp. 1–13.