



SIGGRAPH 2026
Los Angeles 19–23 JUL

The Premier Conference & Exhibition on
Computer Graphics & Interactive Techniques

LatentScore

Sketching Soundscapes with
LLM-Distilled Retrieval for
Procedural Synthesis

PRABAL GUPTA

prabal.ca/siggraph · [paper](#) · [github](#)



2026



ACT I · MOTIVATION

**What if your life
had a soundtrack?**









ACT I · MOTIVATION


What if emotions
could play a musical
instrument?

ACT I · MOTIVATION



... and 28 more knobs ...

ACT I · MOTIVATION



ACT II · DATA BOOTSTRAP

How do we generate
a library of musical
recipes?

Emotional Descriptions

“warm jazz café”

“sunken coral reef”

“fireside reading nook”


Musical Recipe

mode : **dorian**
tempo : **slow**
echo : **subtle**

pad : **ambient_drift**
space : **large**
depth : **true**

texture : **vinyl_crackle**
grain : **warm**
space : **small**

We mine emotion descriptions from ordinary text.



Emotional Descriptions

“warm jazz café”

“sunken coral reef”

“fireside reading nook”

Musical Recipe

mode : **dorian**
tempo : **slow**
echo : **subtle**


pad : **ambient_drift**
space : **large**
depth : **true**

texture : **vinyl_crackle**
grain : **warm**
space : **small**

**How do we
get the
recipes?!**

language models can turn emotional descriptions to recipes

*“fireside reading
nook”*



gemi-3-flash




texture : vinyl_crackle
grain : warm
space : small



what if we use a teacher to grade recipes?

“beautiful sunrise”




Text & Audio
Embedding Model




83%

retry multiple times to do better

*“fireside reading
nook”*



gemini-3-flash



pad : ambient_drift
space : large
depth : true



texture : vinyl_crackle
grain : warm
space : small



mode : dorian
tempo : slow
echo : subtle



retry multiple times to do better


pad : ambient_drift
Space : large
depth : true



texture : vinyl_crackle
Grain : warm
Space : small



Mode : dorian
tempo : slow
echo : subtle




Text & Audio
Embedding Model

55% ✗

62% ✗

83% ✓

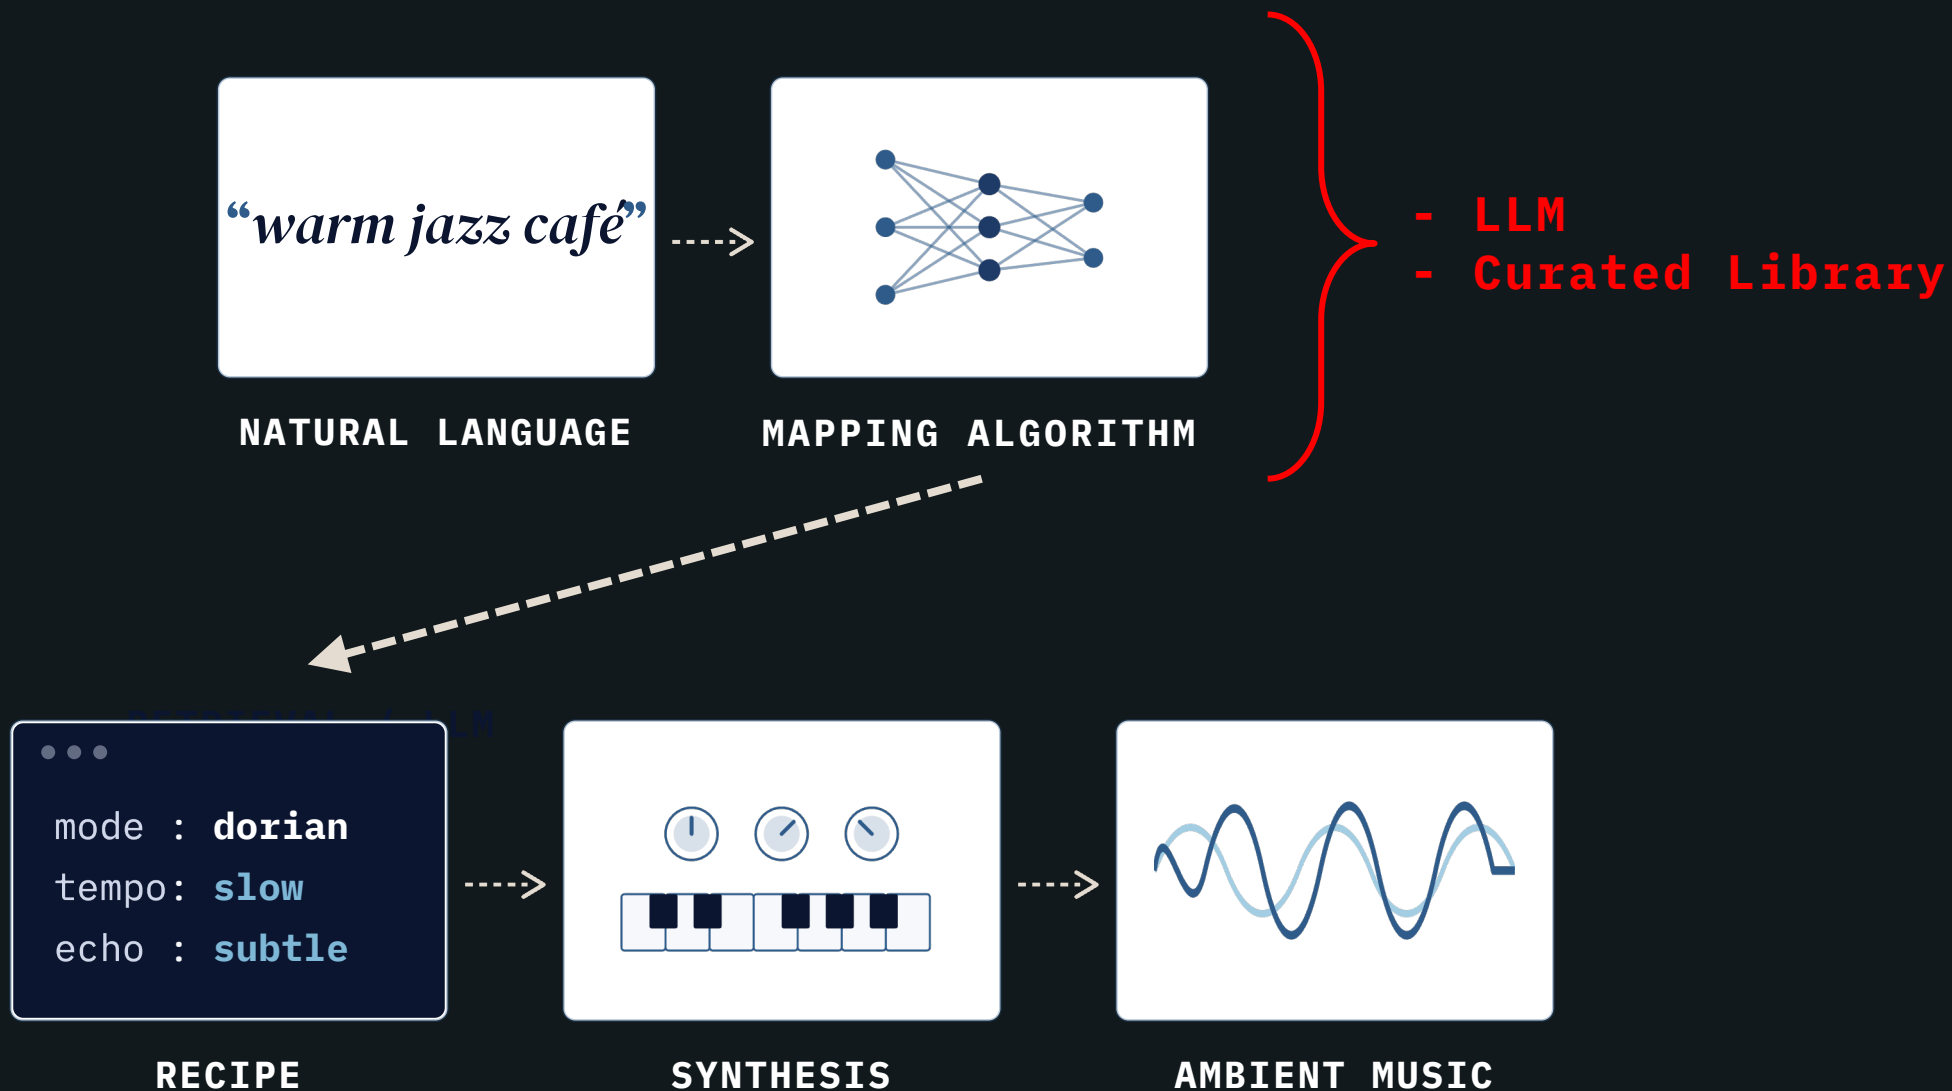
we now have a library of emotional descriptions and musical recipes



ACT III · SHOWTIME

How do we play
music now?

Use LLMs or recipe library to turn natural language to recipes



search for a similar emotional description in the library



generate and steer generated music in (near) real-time, on CPU

DESCRIBE

SHAPE

STEER

tempo: **very_slow**



tempo: **very_fast**


ACT III · SHOWTIME

Built for ambient use.

BUILT FOR

non-vocal procedural
music


ambient · steerable · runs on CPU



NOT DESIGNED FOR

vocals, lyrics, or
recorded samples

exact genres · authored melody · timbre specificity



try now at latentscore.com

• OPEN SOURCE

LatentScore

A Python library that generates ambient music from text descriptions. No GPU required. Turn vibes into sound with a single line of code.

► Live Demo

 GitHub

`pip install latentscore`





SIGGRAPH 2026
Los Angeles 19–23 JUL

The Premier Conference & Exhibition on
Computer Graphics & Interactive Techniques

LatentScore

Editable Music from
Natural Language

PRABAL GUPTA



prabal.ca/siggraph
[paper](#) · [github](#)



2026

