



IMPLICIT AFFECTIVE STEERING

MoodCurator: A Design Probe for Generative Curatorial Storytelling

DIS 2026 · Singapore

Takeaway

Audiences steer AI curatorial stories through visible choices, not complex prompts.

palette selection

artwork curation

story output

Marci Chi Ma · Lei Shi · Hanwen Li · Jiacheng Cheng · David Kirk
Open Lab, School of Computing, Newcastle University, UK

WHY COLOUR?

Colour-emotion links have been discussed in art theory and psychology for more than a century.

Colour has long been used in art and psychology to express mood, atmosphere, and affect.

Colour carries feeling through properties such as hue, saturation, and brightness, but these associations are shaped by context and culture.

In  MoodCurator, colour is used as a visible, low-friction steering cue, not as a claim about the audience's true emotion.

MOTIVATION

Generative AI can write stories. The interaction still gets in the way.

Online art audience

I know the mood I want...

...but I may not have the art vocabulary
or prompt-writing skill.



Prompt box

“Translate your aesthetic
intention into precise text.”

**This is the *interaction*
bottleneck.**

DESIGN GAP

Steering should be low-friction, user-controlled, and legible.

Long text prompts

High semantic labour




Emotion self-report

Breaks the viewing flow



Passive sensing

Privacy + setup burden



Design opportunity:

Can aesthetic choices become steering controls?

Instead of asking audiences to explain intention as a prompt,  MoodCurator lets them compose intention through visible interface actions.

01 Palette

select colour as affective cue

02 Curation

choose artworks to ground the story

03 Voice

choose interpretive stance

The audience is not asked to write a prompt but steers the story through visible choices.

MoodCurator

01 Implicit Affective Steering

- Select a palette (affective cue)

02 The Affective-to-Narrative Pipeline

(a) Colour Analysis Module

Extract dominant colours and map to canonical colour bins

(b) Palette-to-Affect Mapping

Compute affect cue distribution

love trust surprise
sadness ... optimism
disgust anger fear

(15 emotions total)

(c) Diversity-aware Art Recommendations

Retrieve candidate artworks from a GAC dataset (with diversity constraints)

(10 artworks total)

(d) Grounded LLM Narrative Generator

Generates a short story grounded in the selected artworks and interpretive voice

03 Co-Created Interpretive Story

- Curate three artworks
- Select an interpretive voice
- Read the generated narrative
- Leave feedback
- Download story

BACKEND ARCHITECTURE

How visible choices become a grounded, steerable curatorial story.

Visible user choices:

Palette

Curation

Voice

Story

Runtime pipeline:

01 Colour Analysis

k-means (k=5) extracts 5 dominant colours; CIEDE2000 maps them to 12 colour bins.

palette → features

02 Affect Model

85-dim features → Gradient Boosting multi-label classifier → 15 affect cue scores.

self-trained

03 Recommendation Engine

Similarity over Google Arts & Culture metadata, then MMR diversity re-ranking.

curated slate

04 Prompt Builder

Selected cue + 3 artworks + metadata + interpretive voice form structured LLM input.

bounded prompt

05 LLM Story Generator

Claude 3.5 / GPT-5 mini generate a concise story grounded in selected artworks.

bounded prompt

Offline / stored components:

WikiArt-Emotions Dataset

Training labels for palette-to-affect mapping.

Self-trained Classifier

Gradient Boosting; micro-AUC 0.5906; cue not truth.

Google Arts & Culture Corpus

Google Arts & Culture-derived metadata for retrieval and grounding.

Story Controls

Read, export, regenerate, and leave feedback.

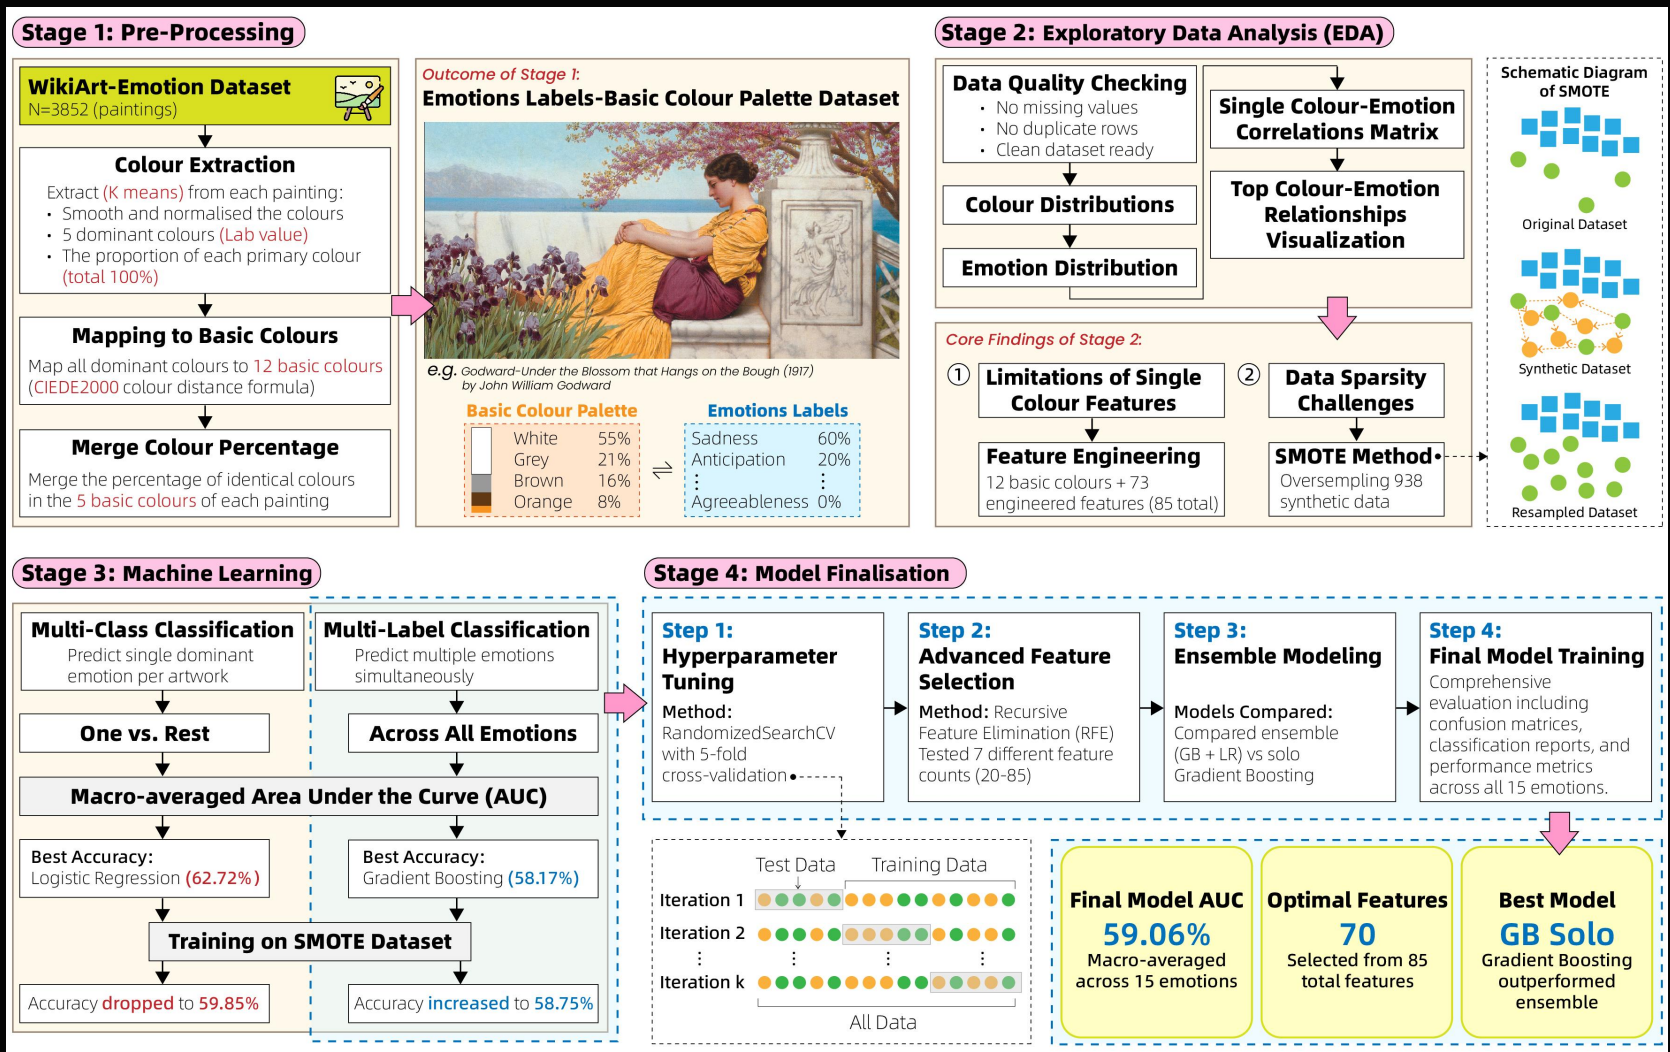
WikiArt-Emotions = affect model; GAC-derived corpus = diverse recommendations.

Framing note:

Colour is not claimed as true emotion detection. It is a transparent, probabilistic steering cue that audiences can accept, reject, or redirect.


COLOUR-TO-AFFECT MODEL

The self-trained pipeline that maps palette features into candidate affect cues before narrative generation.



WHY THESE PALETTES?


Extending basic colour categories into co-designed affective palette options.




The palette set extends the 12 basic colour categories used in Google Arts & Culture, and was refined through expert co-design with colourists, graphic designers, artists, and other visual practitioners.

15 CANDIDATE AFFECT CUES

The affect cue space used by the colour-to-affect model.



Trust



Surprise




Sadness



Pessimism




Optimism



Love




Humility




Happiness



Gratitude



Fear



Disgust



Disagreeable



Arrogance



Anticipation




Anger

These 15 labels span **Positive**, **Negative**, and **Neutral** affect categories, and come from the WikiArt-Emotions dataset used to train our colour-to-affect mapping.

FIVE INTERPRETIVE VOICES


DIS 2026

Five narrative stances that let audiences shape how the same artworks are interpreted.



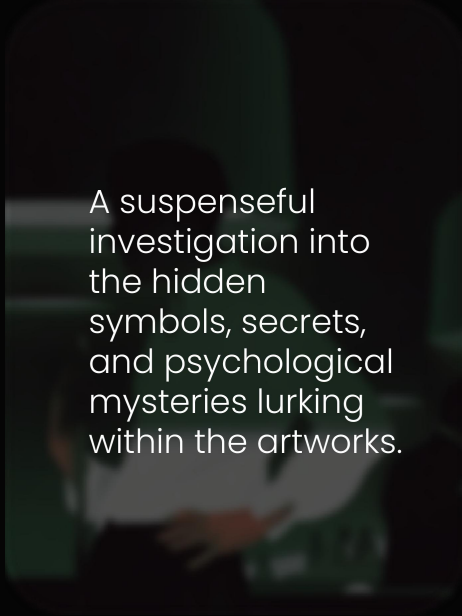
A journey through time, uncovering the stories behind the art, the artist's life, and the world they lived in.

Historian's Chronicle



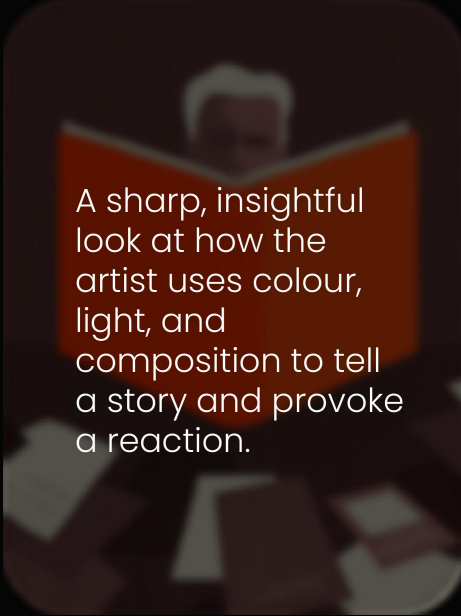
A lyrical and imaginative story that explores the beauty, emotion, and sensory details within each painting.

Poet's Dream



A suspenseful investigation into the hidden symbols, secrets, and psychological mysteries lurking within the artworks.

Detective's Case



A sharp, insightful look at how the artist uses colour, light, and composition to tell a story and provoke a reaction.

Critic's Analysis



A playful and imaginative tale that connects the historical artworks to our modern world in surprising and humorous ways.

Time Traveller's Report

PROMPT SCAFFOLD

The prompt constrains structure, grounding, and stance.



01 Inputs



Interpretive voice

Selected interpretive stance



Affect cue + intensity

Emotion label with intensity



Character name

Narrative focal character



Three selected artworks

In fixed sequence



Artwork metadata / visible details

Titles, artists, dates, materials, visible elements



02 Output Structure



One cohesive micro-story

A continuous narrative with a clear arc



~300 words total

Concise and focused



Exact 3-section sequence

One section per artwork in order

Painting 1



Painting 2



Painting 3



03 Grounding Rules



Anchor the story in the provided images and metadata

Ground every section in what is given



Use only visible details

Describe what can be seen in the artworks




Do not invent unsupported provenance, dates, locations, or artist intent

Avoid assumptions and additions beyond the provided information

DESIGN ITERATION

From early critique to the final iteration pattern.

01 Initial Prototype




- Low-fidelity screens
- Core flow: palette → artworks → voice
- Early assumptions on layout and controls

02 Pilot Study




- N=5, 60-minute in-person
- Tested task flow and comprehension
- Observed navigation behaviour
- Collected feedback on clarity, pace, and preferences

03 Expert Co-design Workshop



- N=5 experts, 120-minute remote
- Critiqued legibility, contrast, and information hierarchy
- Discussed controllability and trust
- Generated design revisions and new interaction ideas

04 Final System



- Refined layout and visual hierarchy
- Clearer controls and feedback
- Stronger link between palette, artworks, and narrative voice

EVALUATION: STUDY 1 FINDINGS


Positive scores, but one response exposed a deeper framing problem.

Study 1: Does the loop work?

- Remote mixed-methods deployment
- N=64 survey

What we measured?

- 13-item post-experience questionnaire
5-point Likert scale
- Three evaluation constructs
Usability · Agency & Personalisation · Narrative Quality & Engagement
- Interaction logs + open responses
- Descriptive indicators



Unexpected tension:

One participant with strong art engagement noted:

The generated story is too western-based, totally tearing apart my emotional experience.



Motivation for study 2:

Positive scores showed the loop could work, but mismatch cases showed where it broke. study 2 zooms in on why some stories felt forced or culturally skewed.

EVALUATION: STUDY 2 FINDINGS

Two recurring mismatch patterns in the think-aloud follow-up.

Study 2: Where does it break?

- Think-aloud follow-up
- N=10 purposive sample
- Focus: interpretive skew, cultural framing, model defaults



Pattern 1: Recurring Western Narrative Tropes



Terrace in the Garden near the Wannsee towards Northwest
by Max Liebermann, 1916

- The model sometimes defaulted to conflict-driven or suspenseful arcs.
- A tranquil garden scene was reframed as a psychological-thriller-like setting.

Pattern 2: Individualised Interpretation



Rainy Season in the Tropics
by Frederic Edwin Church, 1866

- Shared or cultural meanings were reduced to a private emotional story.
- A landscape read as communal heritage was reframed as a personal mood swing.

01 Make framing visible

Show how palette, artwork, and voice constrain generation.



02 Make uncertainty discussable

Present affect outputs as candidate cues, not answers.



03 Add post-generation control

Let audiences revise the interpretive lens after reading.



04 Test across language contexts

English deployment surfaced mismatch; multilingual studies are needed.



01 Legible steering paradigm

Palette + curation + voice as visible controls for AI-supported cultural interpretation.

02 Empirical evidence

A remote deployment showing low-friction agency and narrative resonance patterns.

03 Contestable framing

Interpretive mismatch points toward more explicit, disputable AI framing controls.

Final takeaway:

Generative AI for cultural interpretation should be personalised, yes, but also steerable, legible, and contestable.