

PROJECT: INSTITUTION  
DATE: Guangzhou, 2026


---

# ART AS SIMULATION, EXHIBITION AS PATH

VR pavilions, ALAN and the physical museum at Institution

艺术即模拟，展览即路径


TOM CEBIC · CO-FOUNDER, INSTITUTION



# Institution is a contemporary art institution working in Italy and New York, and now China, with five public systems

|                 |                                                                                                                           |
|-----------------|---------------------------------------------------------------------------------------------------------------------------|
| STAFF/FOUNDERS  | Fabio Cavallucci, Tom Cebic, Lucilla Grossi, Wang Lifu                                                                    |
| PROGRAMMES      | Decolonized History of Art: 24 lessons, 24 faculty, in Italian, English and Chinese · Culture Commons Forum, two editions |
| SYSTEMS         | vr.institution.art · curator.institution.art (ALAN) · atrium.institution.art · translate.institution.it · institution.tv  |
| PHYSICAL MUSEUM | Clouds A0                                                                                                                 |

## The system: one exhibition, two spaces, one curator, one operations layer




**THE EXHIBITION AS A SYSTEM OF PATHS · semantic, affective and physical**

**ATRIUM**  
plans, drafts, checks and records the institution's work

# The VR world runs today: 23 pavilions, 16 open, 37 works, three languages

|                |                                                                                        |
|----------------|----------------------------------------------------------------------------------------|
| ADDRESS        | vr.institution.art, browser (WebGL 2) and headset (WebXR)                              |
| PAVILIONS      | 23 in the database, 16 open, built from about 20 reusable templates                    |
| WORKS IN SLOTS | 37: 32 video, 4 article, 1 sculpture                                                   |
| TOPICS         | Institution Under Construction: Gaze & Memory, Body & Identity · AI and Sustainability |
| DESCRIPTIONS   | English, Italian and Chinese per pavilion                                              |
| NEXT           | multi-user, voice and avatars at scale                                                 |



HUB PAVILION, LIVE WORLD, 2026 SNAPSHOT

# The pavilion catalog defines nine pavilion types




CATALOG Pavilion Catalog, with the full catalog matrix revised in 2026.  
Clouds A0 (Masayuki Sono, Ostap Rudakevych)

TYPES FLOCK · ARTICLE · CINEMA · ART · LECTURE · PERFORMANCE  
(open, closed, chamber, sound, poetry, immersive) · VIRTUAL  
REALITY · BLACK ROOM · EDUCATION

SHEET FIELDS description · screen/art · users · control · lighting · materials ·  
remarks

FLOCK “transition from line to volumetric flock”: the constellation is a  
formation, not a map

IN UNITY Art A-F, Lecture 1S/5S, Sculpture, Performance A-E, Video,  
Article, Hub, Education



## The works in the pavilions




**MARINA ABRAMOVIĆ AND ULAY** · Light/Dark · single video




**JÓZEF ROBAKOWSKI** · Attention: Light! · film after a Paul Sharits sketch



**KAROLINA WOJTAS** · steals from net · restaged photographs



**ARTIOM CONSTANTINOV** · Pain, Suffering · audio-reactive video



**NEW-TERRITORIES (FRANÇOIS ROCHE) WITH EMANUELE COCCIA**  
· (n)Permanencies Tractatus · sculpture

### ALSO OPEN


**Suzanne Lacy** · By Your Own Hand · video, Quito bullring

**Katarzyna Kozyra** · The Rite of Spring · seven screens

**Mattia Cupelli** · SYNC/IDENTITY

**Analivia Cordeiro** · BodyWays

## Lectures, articles and the Education Pavilion



### LECTURE PAVILION

Thomas Seeley on honeybee democracy: a tiered amphitheatre, five screens

### ARTICLE PAVILIONS



Jeff Hawkins, A Thousand Brains · Sociocracy · "Institution: A Copernican Revolution for the Arts" (Fabio Cavallucci)

### EDUCATION PAVILION


a ring pavilion: lecture hall under one immersive screen for 900 people at tables of ten; workshop pods with switchable transparency

### IN THE WORLD

lecture and article pavilions · the Education Pavilion in Unity; its programme comes next




## A visitor moves by walks, not menus: hub, topic map, flight, pavilion, portal



|          |                                                                      |
|----------|----------------------------------------------------------------------|
| ENTRY    | Google sign-in, then the Hub Pavilion                                |
| OVERVIEW | the topic map: topics and subtopics placed in space, zoom in and out |
| MOVEMENT | scripted flight along waypoints between pavilions, about 3 s         |
| INSIDE   | walk, look, read the work's text, ask ALAN                           |
| PORTALS  | 15 directed links; NEXT follows a curated path                       |
| TOGETHER | avatars, spatial voice, text chat, comments and likes                |

# The VR platform is a Unity WebGL client over Supabase, a dedicated multiplayer server and the ALAN backend

|             |                                                                                                                                                          |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| CLIENT      | Unity 2022.3 LTS, URP, WebGL 2 in the browser; WebXR for headsets; core download about 119 MB                                                            |
| DATA        | Supabase: Postgres (topics, subtopics, pavilions, portals, content, profiles, likes, playlists), Storage (35 pavilion bundles), Realtime, Google sign-in |
| MEDIA       | video streamed from CDNs onto in-world screens; images by URL                                                                                            |
| MULTIPLAYER | Netcode for GameObjects, dedicated server in AWS London, 30 ticks/s, WebSocket over TLS                                                                  |
| VOICE       | Photon Voice, spatial, lip-sync on avatars                                                                                                               |
| HOSTING     | S3 + CloudFront; staging on Amplify                                                                                                                      |
| ALAN        | HTTPS to curator.institution.art, token-gated; the client sends identity, location and the visitor's words only                                          |



DATA ARCHITECTURE ANALYSIS, EXTENDED REALITY (NILS GALLIST)

# A pavilion is data: a template, a row of coordinates, and works bound to named slots



|              |                                                                                                                                    |
|--------------|------------------------------------------------------------------------------------------------------------------------------------|
| PAVILION ROW | name · template · coordinates x;y;z · subtopic · visibility and scheduled opening · ambient sound · description in three languages |
| CONTENT ROW  | title · body text · type (video, article, image, 360 video, sculpture) · media URL · slot (Video;1, Article;2 ...) · creator       |
| BOOT         | download every visible pavilion, bind content to slots, place an exterior at each coordinate, land in the hub                      |
| CMS          | Pavilions · Topics · Positions (the walk graph) · Users · Bundle                                                                   |
| DESIGN RULE  | content changes without an application update; the database stays plain text, readable by AI                                       |

# ALAN is a system of six layers

|   |                      |                                                                                                   |
|---|----------------------|---------------------------------------------------------------------------------------------------|
| 6 | OUTPUT               | text with citations in chat; speech in VR, sources carried separately                             |
| 5 | MEMORY               | per-visitor summary rewritten every six turns; register inferred per turn, not stored             |
| 4 | VOICE AND DISCLOSURE | six disclosure tiers, six evidence roles, register filter, quote attribution check                |
| 3 | RETRIEVAL            | vector index with lexical fallback; scoped artwork → pavilion → reference                         |
| 2 | CORPUS               | Institution's own: authored notes, curator reference, course transcripts, cohort-reviewed answers |
| 1 | MODEL                | Anthropic Claude on Amazon Bedrock: Sonnet 4.6 at launch, Sonnet 5.5 now                          |



# The model layer: Claude on Amazon Bedrock,



|             |                                                                              |
|-------------|------------------------------------------------------------------------------|
| PROVIDER    | Anthropic Claude through Amazon Bedrock                                      |
| FALLBACK    | the previous Sonnet, then Haiku, on a provider error                         |
| TEMPERATURE | 0.35                                                                         |
| EMBEDDINGS  | OpenAI text-embedding-3                                                      |
| SPEECH      | browser speech synthesis in VR; server-side speech and transcription in chat |
| HOSTING     | FastAPI on one EC2 instance, us-east-1; Postgres; S3; Secrets Manager        |

# The corpus layer: Institution’s proprietary corpora, worldwide in their sources

CORPUS 1

Authored notes

CORPUS 2

Curator reference

CORPUS 3

Course transcripts

CORPUS 4

Cohort-reviewed answers

|              |                                                                                                                                                           |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| SIZE         | 1,519 documents, about 21,000 chunks; 45 named research paths                                                                                             |
| SOURCES      | Institution’s exhibitions, publications and recorded encounters; the Decolonized History of Art course; curatorial sources from outside the Western canon |
| MAINTENANCE  | nightly review of 10% of the corpus and 1% of cohort answers; nightly research on thinly covered artists                                                  |
| WEB FALLBACK | in public chat only, labelled, quarantined, never ingested automatically; none in VR                                                                      |
| POSTURE      | non-Western and under-covered voices; cohort thinking; conceptual mapping over movement labels                                                            |

# The grounding layer: three-pass retrieval, disclosure tiers and a pre-publish check



RETRIEVAL

8 chunks per pass, public tier only in VR

TIERS

public direct · public paraphrase · public link only · public silent · operator only · private archive — enforced before any prompt is built

ROLES

primary source · supporting evidence · bridge · voice calibration · prosodic template · counterposition

CHECK

forbidden verbs · quote attribution · five-word match against calibration text

LENGTH

400 characters first answer · 900 on a follow-up · 1,500 on a deeper one

RESULT

answers come from documents; when none support an answer, ALAN says so

# Fine-tuning work: QLoRA adapters on an open 32B model, and a video-language model trained on artworks

## VOICE ADAPTERS


|          |                                                                                                                                                                             |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BASE     | DeepSeek-R1-Distill-Qwen-32B                                                                                                                                                |
| METHOD   | supervised fine-tuning with QLoRA: 4-bit NF4, LoRA rank 64 then 128, 4,096-token sequences                                                                                  |
| DATA     | Institution's own writing                                                                                                                                                   |
| VERSIONS | seven adapters; one preference-optimisation run tried and reverted                                                                                                          |
| HARDWARE | H100, then L40S on AWS, then A100-class                                                                                                                                     |
| FINDING  | early training text had been drafted by a language model; v7, retrained on 489 checked essays, dropped the model's punctuation habit (em-dashes 14.0 → 2.1 per 1,000 words) |

## ARTWORK MODEL

|               |                                                                           |
|---------------|---------------------------------------------------------------------------|
| BASE          | VideoLLaMA-3                                                              |
| DATA          | five artworks with curator-written labels for style, colour and emotion   |
| TRAINING      | fine-tuned and sharded across GPUs by the research team                   |
| DEPLOYED      | Hugging Face, with a basic front end                                      |
| IN PRODUCTION | ALAN answers through Claude with retrieval; these two are research models |

## In a pavilion, ALAN knows where the visitor stands and answers from that work's documents first

|                |                                                                                                       |
|----------------|-------------------------------------------------------------------------------------------------------|
| CLIENT SENDS   | session identity · pavilion arrived or left · artwork in focus · the visitor's words                  |
| SERVER RETURNS | a spoken answer (400 / 900 / 1,500 characters) · its sources, separately · the updated memory summary |
| PRESENCE       | a floating "Ask ALAN" sphere that follows at 3.7 m; a chat panel in headset mode                      |
| MEMORY         | a short per-visitor summary across sessions; the next visit resumes where the last stopped            |
| LANGUAGES      | English, Italian and Chinese                                                                          |



# Semantic and affective readings are separate layers of the same work



|        |                                                                                                                             |
|--------|-----------------------------------------------------------------------------------------------------------------------------|
| DESIGN | two attention pathways; multimodal encoders for image, audio and text; expert curators as ground truth for cultural context |
| RULE   | a bird is mourning in one culture and fortune in another; the affective reading is cultural, so the training data must be   |
| WHERE  | the semantic layer is the corpus; the affective layer is the data programme on the next sheets                              |

# The exhibit already produces data: position, dwell, questions, register

ONE EVENT, AS RECORDED TODAY

|         |          |      |          |         |          |          |
|---------|----------|------|----------|---------|----------|----------|
| visitor | pavilion | work | t arrive | t leave | question | register |
|---------|----------|------|----------|---------|----------|----------|

|           |                                                                                                                        |
|-----------|------------------------------------------------------------------------------------------------------------------------|
| POSITION  | which pavilion, which work in focus, timestamped                                                                       |
| DWELL     | time between arrival and departure, per pavilion and per work · computed from the events above                         |
| QUESTIONS | what visitors ask at a given work, and in which language; stored per session                                           |
| REGISTER  | factual · operational · scholarly · contemplative · lyric; inferred per turn to set pacing; not stored across sessions |
| REACTIONS | like · save to playlist · comment                                                                                      |
| USE TODAY | pacing of answers; curators reading what was asked; corpus research on thinly covered works                            |

# Sensing: eye tracking, gaze, pulse and heart rate, in the headset and in the building

AI GUIDE PROGRAMME

|                      |                                                      |                                                  |                                     |                                                     |
|----------------------|------------------------------------------------------|--------------------------------------------------|-------------------------------------|-----------------------------------------------------|
| PHASE 1<br>MVP guide | PHASE 2<br>Enhanced guide: image recognition, speech | PHASE 3<br>Mobile guide and physical institution | PHASE 4<br>Emotional-impact sensing | PHASE 5<br>Artist-specific models; affective corpus |
|----------------------|------------------------------------------------------|--------------------------------------------------|-------------------------------------|-----------------------------------------------------|

|                 |                                                                                                                                                                                 |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IN THE HEADSET  | head pose; eye tracking and gaze target on headsets that provide it                                                                                                             |
| ON THE BODY     | pulse and heart rate from a wearable, opt-in                                                                                                                                    |
| IN THE BUILDING | image recognition so the guide follows the visitor from work to work; NFC or QR at the work; room sensors for emotional impact, with artists making installations that use them |
| TODAY           | no biometric data is collected                                                                                                                                                  |



# One work, two sensors: the same event schema in VR and in the building

ONE EVENT, THE SAME IN BOTH SPACES

| visitor | work | time | position | gaze target | dwell | question | affect · opt-in |
|---------|------|------|----------|-------------|-------|----------|-----------------|
|---------|------|------|----------|-------------|-------|----------|-----------------|

|                                                                                                                                                                  |                                                                                                                                                                                                                                  |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <div>IN VR</div> <div>position from the engine</div> <div>gaze from the headset</div> <div>dwell from the events</div> <div>affect from a wearable, opt-in</div> | <div>IN THE BUILDING</div> <div>position from image recognition or a tag at the work</div> <div>gaze from eye-tracking glasses or a headset</div> <div>dwell from the same events</div> <div>affect from the same wearable</div> |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|         |                                                                                                                                                                                                      |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PAIRING | every pavilion has a physical counterpart; a work has one identity in both spaces                                                                                                                    |
| PULSE   | the wearable that records a visitor's heartbeat can drive the gravity pulse of the rule sheet: the room breathes with the body in it, affective data becomes physics, and the physics becomes a path |
| FEEDS   | ALAN's pacing · curators' reading of a show · installations that respond · an affective corpus with cultural labels                                                                                  |

# Consent and data rules


|            |                                                                                                          |
|------------|----------------------------------------------------------------------------------------------------------|
| OPT-IN     | any body signal is collected only when the visitor turns it on; off by default                           |
| OWNERSHIP  | affective data belongs to the visitor; sharing with an institution or other visitors is a visible choice |
| NOT STORED | the per-turn register; raw conversation beyond the session summary                                       |
| NOT USED   | visitor data is never used to train a model; one deployment per institution, no shared corpus            |
| SEPARATION | public and operator indices are separate systems; the private archive is outside every index             |
| PROCESSING | processor agreements and GDPR processing records exist for the VR platform                               |
| BUILDING   | cameras raise questions a headset does not; tags and visitor-carried devices come first                  |

# Atrium is the operations layer: a typed engine, eight role rooms, checkpoints, cost caps



|             |                                                                                                               |           |                                                                                                                            |
|-------------|---------------------------------------------------------------------------------------------------------------|-----------|----------------------------------------------------------------------------------------------------------------------------|
| ENGINE      | LangGraph; agents: planner, critic, researcher, drafter, verifier, aggregator                                 | COSTS     | tokens × tool calls × wall-clock per model tier (Haiku, Sonnet, Opus); default cap \$25 per institution per day            |
| ROLE ROOMS  | curatorial · grants · development · education · press · support · artist relations · archive · daily briefing | ISOLATION | Postgres row-level security per institution                                                                                |
| CHECKPOINTS | budget · plan coherence · drift · value · voice · faithfulness · citation integrity · human approval          | KNOWLEDGE | Bedrock knowledge base over eight shared drives; 286 of 316 documents indexed                                              |
| AUTONOMY    | L0 drafts only · L1 confidence-gated side effects · L2 background loop with daily caps                        | TOOLS     | Drive, Docs, Sheets, Slides, Calendar; Gmail send with audit log and revert; email and Slack access with verified identity |

# Atrium services in use: the staff gateway and the live interpretation service



## STAFF GATEWAY

alan@institution.art and Slack; identity verified first (DKIM and DMARC alignment; signed requests and a workspace allowlist); answers in seconds with real tools; an email is sent only after a one-word confirmation; every reply copies the operator

## STUDIES

agent research pinned as signed, dated institutional positions

## INTERPRETATION

translate.institution.it: 13 listener languages from an AI cascade, transcribe → translate → synthesise, 2–5 s; at the forum, three human interpreters in browser consoles, AI engine installed and off

## HUB

institution.tv: two days, four audio channels, eight streams; a direct room feed where YouTube is blocked

## COST

a few dollars per language-hour, against \$60–120 per hour for commercial interpretation

## A pavilion is a simulation; its rules go on the sheet next to lighting and materials

|          |                                                                                                                                                                                                      |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TODAY    | every pavilion is already a simulation: a Unity scene runs gravity, collision, light and sound for the works in it. The rules are Earth's, and fixed                                                 |
| THE STEP | write the rules on the pavilion's sheet, next to LIGHTING and MATERIALS, so an artist can author them as part of the work                                                                            |
| RULES    | eighteen fields: gravity and its direction · tilt · pulse · grip · bounce · viscosity · spin · lift per theme · cohesion · touch and touch memory · softness · ripple · time bubble · time scale     |
| BODIES   | works, each with a theme vector (water, stone, light, body, machine, memory) · visitors · a ball · a wheeled robot                                                                                   |
| SURFACES | floor and ceiling as $72 \times 72$ pressure grids; a dent forms where load accumulates                                                                                                              |
| PULSE    | the wearable that records a visitor's heartbeat can drive the gravity pulse of the rule sheet: the room breathes with the body in it, affective data becomes physics, and the physics becomes a path |
| PROGRAM  | the sheet is executed: 120 Hz in a browser today; the same rules can run inside a Unity pavilion                                                                                                     |



$$a_i = f_i g(t) - v v_i - 2\Omega \times v_i - \Omega \times (\Omega \times x_i) + \text{slope} + \text{cohesion} + \text{query} + \text{policy}$$

## ALAN writes the rules: a sentence becomes a change to the sheet, and the simulator checks ALAN's own prediction

|              |                                                                                                                                             |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| WHO          | ALAN, the same curator that answers from the corpus. Here it reads a request and edits the rule sheet                                       |
| INPUT        | a sentence: "turn the room upside down and make the floor remember where people stood"                                                      |
| ALAN RETURNS | a change set over the eighteen fields · a one-line summary · predictions: metric, direction, expected value, reason                         |
| VALIDATOR    | clamps every value to its range; reports ignored fields; never executes code                                                                |
| CHECK        | the benchmark runs before and after; each prediction is scored on direction and on error                                                    |
| WHY          | language models slide back toward Earth values when asked about a world that is not Earth; here that is a number on the screen, per request |



# Benchmark: one task, eleven rule sheets; the robot’s path is the measurement

The task never changes: a wheeled robot has to push a ball to a goal while a visitor walks the room. Only the rule sheet changes. Each row is 12 seeded runs of 25 s at 60 Hz, measured October 2026. The last number is how far a forecast that assumes Earth physics is wrong about where the ball will be 2 s later.

| RULE SET                 | DELIVERED, OF 12 | MEDIAN TIME, S | CLOSE CALLS, S / TRIAL | EARTH-FORECAST ERROR, M | WHAT HAPPENED                                             |
|--------------------------|------------------|----------------|------------------------|-------------------------|-----------------------------------------------------------|
| Earth                    | 11               | 6.6            | 0.16                   | 0.00                    | the baseline                                              |
| Moon                     | 7                | 13.3           | 0.17                   | 0.97                    | every push floats; slow                                   |
| Underwater               | 2                | 21.4           | 0.32                   | 1.12                    | drag stops the ball short                                 |
| Ice floor                | 3                | 15.8           | 0.39                   | 1.19                    | no grip; the robot slides                                 |
| Zero gravity             | 0                | —              | 0.00                   | 0.23                    | wheels have nothing to push against                       |
| Upside down              | 10               | 8.5            | 0.16                   | 6.36                    | same task on the ceiling; the Earth forecast is 6 m wrong |
| Tilted floor             | 6                | 8.1            | 0.30                   | 0.38                    | the ball rolls downhill, away from the goal               |
| Spinning room            | 3                | 8.7            | 0.49                   | 0.36                    | Coriolis bends every push; more close calls               |
| Memory-foam floor        | 12               | 6.4            | 0.10                   | 0.04                    | dents hold the ball on course                             |
| Light rises, stone sinks | 11               | 6.6            | 0.16                   | 0.00                    | works lift; the ball does not                             |
| Time bubble              | 7                | 6.6            | 0.24                   | 0.11                    | slow near the visitor; close calls rise                   |

## Four commissions written as sheets: the rule is the authored object

### MEANING PULLS

constellation · rule: cohesion and query

#### THE VISITOR

says “water and memory”, standing anywhere in the show

#### THE WORK

the works that carry that meaning come to the visitor: the constellation re-forms around them, light and sound follow, the rest of the show goes quiet

#### IN VR

the pavilions drift together; the visitor is the centre of a new constellation

#### IN THE BUILDING

nothing heavy moves: the drones of the void gather overhead, each work’s light and sound rise or fall by the same potential, and the floor lights the path to the nearest one

#### ARTIST

writer, with a light and sound artist

### FLOOR MEMORY

gallery floor · rule: touch memory, softness



#### THE VISITOR

walks the show, for weeks

#### THE WORK

the floor records pressure; the dent field shows underfoot and in VR; at the end it is milled in marble

#### IN VR

the floor yields and remembers

#### IN THE BUILDING

a pressure floor, projection, then stone

#### ARTIST

sculptor

### TURNING ROOM

sunken arena · rule: spin

#### THE VISITOR

throws and catches

#### THE WORK

the room rotates; throws curve; the Earth forecast is projected and is wrong in public

#### IN VR

the room spins

#### IN THE BUILDING

a turntable at 2–3 rpm; the forecast on the screen

#### ARTIST

sound artist

### SLOW NEAR ME

performance pavilion · rule: time bubble

#### THE VISITOR

approaches a dancer or a machine

#### THE WORK

everything near the visitor slows by one function; machines and dancers obey it together

#### IN VR

exact

#### IN THE BUILDING




machines exact; dancers by ear

#### ARTIST

choreographer


## The building concept: a vertical atrium wrapped in a translucent screen

|           |                                                                                                                                                     |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| CONCEPT   | Physical Space Interiors, Clouds A0                                                                                                                 |
| ATRIUM    | a vertical void wrapped in a translucent screen; content visible from many viewpoints                                                               |
| GATHERING | benches at the foot of the void; projections visible from the street                                                                                |
| GALLERIES | organised around the void, digital content on each floor                                                                                            |
| BRIEF     | a hall of 1,000–3,000 m <sup>2</sup> with a 6–10 m ceiling; a stage that rises and lowers; retractable seating; an underground level for dark works |
| PAIRING   | every VR pavilion has a physical counterpart in the building                                                                                        |



# The building as mechanism: a hall that reconfigures, a void that holds the constellation, a dark level for the rules a body cannot feel

|                |                                                                                                                                                                                                                                                                                        |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PRECEDENT      | buildings that move: a shell that telescopes out over its plaza; theatres whose walls, floors and seating reset into dozens of configurations (The Shed; the Perelman Performing Arts Center). Ours moves by the exhibition's rule sheet                                               |
| THE HALL       | 1,000–3,000 m <sup>2</sup> , 6–10 m high. The rising stage is the tilt and pulse platform; retractable seating turns lecture into arena into open floor; a surround screen shows the VR state at the scale of the large pavilion, about 900 people                                     |
| THE VOID       | the constellation, for real: the FLOCK of the catalog flown as tens of drones standing for thousands of VR points, each one a cluster of works, its brightness the cluster's mass; it re-forms as visitors speak (k-means, Hungarian assignment, minimum-snap paths, person-exclusion) |
| THE DARK LEVEL | black-room works; inverted gravity and time dilation seen enacted, not felt                                                                                                                                                                                                            |
| GALLERIES      | a pressure floor and trolleys; every work has its VR twin                                                                                                                                                                                                                              |
| THREE TIERS    | matter does it (tilt, turntable, foam, pressure floor) · a machine does it (drones, trolleys, winches; $a_{phys} = \Gamma_{safe}(a_R)$ ) · sim-only, for human bodies                                                                                                                  |
| NUMBERS        | tilt 6° gives 1.03 m/s <sup>2</sup> lateral; 0.3 rad/s (2.9 rpm) gives 0.6 m/s <sup>2</sup> Coriolis on a body at 1 m/s. Mechanics, not clearances                                                                                                                                     |



# World models in October 2026: pixels learn the world, engines hold it; a pavilion is an engine world model with art in it


|                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>PIXEL-NATIVE</p> <p><b>text or image → navigable video</b></p> <p>looks real; physics learned and unreliable; a requested physical parameter barely changes what is generated</p> | <p>ENGINE + RENDERER · THE PAVILION IS ONE</p> <p><b>rules in a simulator, a neural model for appearance</b></p> <p>state persists by construction; language becomes rules registered in the engine. A pavilion is this kind of world model, with art in it: the rules are on the sheet, ALAN edits them, a verifier checks</p> |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|                   |                                                                                                                                                                                                              |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| WHAT ART ADDS     | works with meaning as a force; visitors as boundary conditions; a curated path through the room; a public that comes back                                                                                    |
| MEASURED FAILURE  | video models fall under sub-Earth gravity; language models revert to Earth ratios in counterfactual worlds (0 of 15 in a “decay world”); requested-to-realised slope near zero in open image-to-video models |
| ROBOTICS STANDARD | a world model is judged by how well its predicted robot success correlates with real success; $\rho$ and $r$ between 0.87 and 0.99 in 2026                                                                   |
| GAP               | robustness benchmarks perturb visuals, layout, language and objects; none perturbs the physics itself                                                                                                        |
| ADAPTATION        | a frozen world model adapted at test time to unseen friction, mass and human-defined rules: prediction error 15.2 → 7.0, success 33.7% → 73.0%                                                               |

# Exhibition as path: a visitor, a robot and a world model walk the same room, and the path is the measurement

|                  |                                                                                                                                                                        |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| THE PATH         | for a visitor, the sequence of positions, dwells and gaze targets from sheet 20; for a robot, its trajectory; for a world model, its rollout. One schema for all three |
| SUBJECTS         | a human visitor · a robot policy · a world model, in the same pavilion under the same edited rules                                                                     |
| CURVE            | $e_k = e_\infty + (e_0 - e_\infty) e^{(-k/\kappa)}$ ; fit $\kappa$ per subject and compare how fast each adapts                                                        |
| EARTH INDEX      | the gravity a predictor behaves as if it believes: 0 follows the edit, 1 stays on Earth                                                                                |
| POLICY $\rho, r$ | predicted against real robot success, per rule sheet, which no benchmark reports under shifted physics                                                                 |
| SAFETY           | minimum separation · time to collision · safety-filter interventions; a proxy after ISO/TS 15066                                                                       |
| PROTOCOL         | fixed step and seeds; pre-registered splits; confidence intervals                                                                                                      |
| DATA SOURCE      | a public venue: many visitors, repeated exposure, curated content, with consent                                                                                        |

ADAPTATION CURVES · ILLUSTRATIVE, NOT DATA



# What the future projects need, and where the work stands

## NEEDS

|              |                                                                                                                                                                                                                                                     |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CAPABILITIES | server-authoritative physics in the VR world · eye tracking in headsets · one event schema across VR and building · Chinese-language corpus and voice for ALAN · fine-tuning and world-model evaluation on local GPUs                               |
| EQUIPMENT    | eye-tracking headsets · optical motion capture · a pressure-sensing floor section · a netted flight box with micro-drones · omni-wheel trolleys · a small motion platform · projection · opt-in wearables · a GPU cluster                           |
| ROLES        | engine and Unity engineer · robotics and controls engineer · machine-learning engineer · affective-computing researcher · curator-researchers for non-Western and Chinese-language corpora · artists in residence · producer · data-protection lead |
| RESEARCH     | E1 spec match · E2 prediction calibration · E3 adaptation and safety · twin fidelity between VR and room · legibility studies · the affective corpus with cultural labels                                                                           |

## NEXT

### VR WORLD

multiplayer and voice at scale · ALAN on the 2D map · the Education Pavilion programme

### SENSING

eye tracking · wearables · the building side of the event schema · the heartbeat pulse

### BUILDING

the physical museum with Clouds AO

### ART AND RESEARCH

the commissions · the constellation flown · the test-bed experiments