

INTERLINKED, THE CUBE

Reading the Present

A measurement layer for the European primary market in emerging contemporary art.

ABSTRACT

The art market keeps a detailed record of its own past and almost none of its present. Auction results, press indices and career models describe transactions and reputations that have already formed, which is precisely the phase emerging artists have not yet entered. This paper sets out an alternative. A private viewing room, used as an ordinary working tool by galleries and art spaces, can also serve as a measurement instrument. Each session produces consented signals about how one person attended to a known set of works. Set against a structured description of those works, drawn from published research in empirical aesthetics, those signals separate into an account of taste that is useful to the gallery on the day it is produced. Interlinked connects such rooms across a European network of partner organisations, which lets emerging work circulate beyond its own scene and, in aggregate, produces a continuous reading of attention in the primary market before any price exists to be recorded. We describe the instrument, the model, the validation plan, the data architecture, and the limits we consider binding.

01 THE ARGUMENT

Emerging European art does not circulate, the tools that would help it circulate are absent at the small end of the trade, and everything the sector currently measures faces backwards. These three facts are connected, and one instrument addresses all of them.

- 1 **Circulation is the bottleneck, not talent.** The European art market concentrates in a handful of western hubs. An artist working in Bucharest, Riga or Palermo is unlikely to be seen by a collector in Vienna or Amsterdam, not because the work is weaker, but because no low cost path exists between the two.
- 2 **The standard instrument of the trade returns nothing.** Most galleries and art spaces still present available works as a PDF attached to an email. The sender cannot tell whether it was opened, which work held attention, or when the moment to follow up had arrived. The most important conversation in the primary market happens in the dark.
- 3 **Every existing art data product describes the past.** Auction records, press indices and career models are each genuinely useful, and each takes as input something that has already happened. Emerging artists are largely absent from all of them, because they are not yet in the record.
- 4 **A viewing room is already the right shape for an instrument.** It is a bounded session: one identified person, a known set of works, a moment of real intent. What is missing is not the technology to build such rooms, it is the decision to treat them as measurement rather than as presentation.
- 5 **The value appears at two levels, and only one of them can be assumed.** For a single organisation, a room produces a usable account of one collector. Across a consented network it produces something that does not exist anywhere: a reading of attention in the primary market as it forms. The second requires scale. So the first has to be worth using on its own, from the first day. The system is designed on that assumption.

02 WHAT THE SECTOR MEASURES TODAY

A partner or an investor should be able to place this project among the things that already exist. We have looked at that carefully, and the pattern is consistent enough to state as a table.

CATEGORY	REPRESENTATIVE PRODUCTS	WHAT IT RECORDS	TIME BASE	EMERGING ARTISTS
Auction record	Artnet, Artprice	Prices realised, at resale	Past	Largely absent
Press and attention indices	Articker	Published mentions and their velocity	Past, once written	Only after coverage
Career and valuation models	Wondeur, Limna	Exhibitions, endorsement, comparable sales	Past, projected forward	Thin record, weak signal
Gallery software and viewing rooms	Artsy, Artlogic, Arternal	Engagement inside one organisation	Present	Covered, not comparable
Virtual exhibition technology	Walter's Cube, Vortic, Kunstmatrix, Artland, Artsteps	The visit, wayfinding, dwell in a modelled space	Present	Institution centred

Categories are drawn as we understand the market in 2026. Products move, the structural point does not.

Two gaps run through the table, and they are structural rather than competitive.

The first is time base. The upper half of the table takes as input an event that has already concluded: a hammer price, an article, an institutional show. That is a reasonable design for the secondary market, where the object of study is a resale history. It is close to useless for an artist whose first solo show has not yet opened.

The second is the unit of observation. The lower half of the table does sit in the present, and it stops at the edge of one organisation. Engagement data inside a gallery's own system is real and it is private, and it can never become a reading of the market, because a single gallery's collector list is neither large enough nor diverse enough to say anything general.

Nobody measures the present. There is no account of how collectors and audiences are looking at emerging work right now, before a price exists to be recorded, because nobody sits inside the private flow where that looking happens.

03 THE INSTRUMENT

A room, not a replica

The Cube is a room. A partner uploads an exhibition, or a selection prepared for one collector, and receives a private link that opens like a small white room on any device, with no download and no account for the visitor. The visitor browses the works, marks what interests them, and can reach the organisation directly. The organisation sees which works drew attention, and for how long.

We are deliberately not building a three dimensional replica of a gallery. That decision is usually presented as a matter of cost. For us it is first a matter of measurement. Walking a visitor through a modelled space measures wayfinding as much as looking: dwell time in front of a work is confounded by how hard the work was to reach, how the camera behaved, and how comfortable the visitor is with the navigation. A flat, fast room removes most of that confound. It also keeps attention on the work itself, which serves the collector, and it can be produced in minutes rather than days, which is what makes it usable by a small organisation.

What one session yields

- Time attended to each work, and the sequence in which works were opened
- Returns to the same image, which behave differently from first views
- Marks of interest, saves, and requests for detail
- Enquiries sent, and what the visitor asked about
- Where the session ended, which is often as informative as where it lingered

A BINDING DESIGN CONSTRAINT

The room has to be worth using for its own sake. A gallery adopts it because the follow up is better and the work looks right on a phone, not because it contributes to a research programme. Any design that trades usefulness for data quality is rejected, because a room nobody sends produces no data at all.

04 FROM SIGNALS TO TASTE

Two descriptions meet inside a viewing room, and neither is very interesting alone.

The first is the work, which can be described in structured terms: palette, contrast, chroma and composition at one level, and at another its subject, its degree of abstraction, its scale, its emotional register. The second is the visitor's attention, which the room records. Set the second against the first, and the weights begin to separate. One collector's attention loads onto palette and tonal range. Another's loads onto subject and figuration. The two may have spent the same twenty minutes in the same exhibition and meant entirely different things by it.

The architecture, from neural mapping

Over the past decade the Connectomics team at Google Research, working with Jeff Lichtman's laboratory at Harvard, reconstructed neural tissue at synaptic resolution, including a cubic millimetre of human cerebral cortex published in *Science* in 2024 and released openly. What we took from that work is not biology but shape: a model of perception that assembles meaning in stages, through levels that feed one another, rather than reading an image off in a single pass. Our system is built in that shape. It does not treat a painting as a flat bag of tags.

The feature model, from empirical aesthetics

Kiyohito Iigaya and colleagues at Caltech, publishing in *Nature Human Behaviour* in 2021, showed that aesthetic preference for art can be predicted from a mixture of low level visual features such as hue, contrast and brightness, and high level features such as concreteness, dynamism and emotional tone. Their predictions held across many individuals, and the same features emerged unprompted inside a neural network trained for nothing more than recognising objects. Their follow up in *Nature Communications* in 2023 used functional imaging to show the brain assembling aesthetic value hierarchically across those levels, weighting them and integrating them.

That gave us two things: a vocabulary for describing a work, and a reason to believe that individual taste lives in the weights rather than in the works themselves.

LAYER 3 · REPORT

What this person responds to, and which works they are most likely to respond to next, including works held by other galleries in the network.

LAYER 2 · WEIGHTS

How much this individual's attention loads onto each level of description. This is the part that is personal, and the part that is learned.

LAYER 1 · HIGH LEVEL FEATURES

Subject, degree of abstraction, scale, dynamism, emotional register.

LAYER 0 · LOW LEVEL FEATURES

Hue, chroma, contrast, brightness, composition.

The training data is the part no one else holds

The architecture and the feature vocabulary come from published research and are available to anyone. The signals are not. They are produced only inside private viewing rooms, only with consent, and only by organisations that have chosen to grant aggregation rights. Those rights, granted by many organisations over several years, are the part of this that a later arrival cannot assemble quickly.

What the gallerist receives

The output is a report written in language a gallerist already thinks in. It describes what a person responds to and proposes what to show them next, including work held by other galleries in the network that this collector has never seen. Two people who viewed the same show should not receive the same follow up, and the report says why. No questionnaire is involved. Nothing is asked of the collector that the collector did not already do.

05 DOES ATTENTION SURVIVE A SCREEN

This is the fair question, and the honest answer is partly. A 2021 study in *Scientific Reports* compared how people looked at an installation artwork in its physical form and in a virtual reproduction, and found the gaze patterns broadly similar. Attention to art does not simply evaporate when the work moves online, which is what makes a viewing room a usable instrument rather than a curiosity.

What a screen does compress is scale, surface and material presence. A large canvas and a small one are the same size in a browser until the interface says otherwise, and texture is nearly absent. We treat this as a known distortion rather than a solved problem: the room states physical dimensions prominently, and the model carries scale as a declared property of the work rather than inferring it from the image.

06 FROM ONE ROOM TO AN INDEX

Inside one gallery, this is a useful convenience. Across a consented network in several countries it becomes something else: a reading of demand for emerging art as it forms, week by week, and eventually a barometer of present attention on the primary market which, as far as we can establish, does not yet exist anywhere.

Such an index would describe where attention is gathering by medium, subject, scale and register, and how that differs between scenes and audiences, at the stage before prices exist. It is useful to galleries deciding what to bring to a fair, to curators reading a generation, to funders assessing whether public support actually moves work across borders, and to the artists themselves.

THE RISK, STATED PLAINLY

An index of this kind becomes valuable only once the network reaches a certain scale, and networks fail more often than they form. We do not treat scale as a plan. We treat it as an outcome that may or may not arrive, which is why the room is designed to be worth paying for on its own from the first day, and why the aggregation layer is opt in rather than assumed.

07 CONSENT AND DATA ARCHITECTURE

A system that reads attention has to be designed for European data protection from the first line of code, not adapted to it afterwards. What follows are commitments we intend to be held to, by partners and by a research collaborator.

- **Consent at the point of looking.** The visitor is told in plain language, before the room opens, what is recorded and why. Consent is granular and withdrawable.
- **Aggregate by default.** Signals leaving a partner organisation are aggregated. Individual profiles are never sold, and never leave the organisation that holds the relationship.
- **The gallery owns its client and its data.** We are private infrastructure, not a marketplace. We take no commission, we do not stand between a gallery and its collector, and data is exportable.
- **Purpose limitation.** Signals are used to serve the gallery that produced them and, where rights are granted, to build the aggregate index. Nothing else.
- **Aggregation is opt in, per partner.** A partner may use the room and decline the network layer entirely.
- **Retention is bounded** and stated, with EU hosting and a data protection impact assessment maintained as a live document rather than a launch artefact.

08 HOW WE INTEND TO VALIDATE THIS

Claims about reading taste are easy to make and rarely tested in public. We intend to develop this strand with a university laboratory in empirical aesthetics or vision science as a full partner, so that the method is designed, reviewed and published to research standards rather than asserted by a company about its own product. Four questions, in order of difficulty.

H1 · Within session prediction

Do attention signals over part of a room predict which works a visitor marks or enquires about in the remainder, better than a baseline of overall work popularity? This is the weakest claim and the necessary one.

H2 · Stability of the weights

Do feature weights estimated for a person in one session predict their attention in a later session, with different works? If taste lives in the weights, weights should travel.

H3 · Transfer across organisations

Do weights learned at one partner predict attention at another partner's room? This is the claim the network layer depends on, and it cannot be tested inside a single gallery.

H4 · Usefulness, judged blind

Do gallerists, shown a report and a plausible decoy, identify the report that matches the collector they know? A model can be statistically real and practically worthless, and this is the test that separates the two.

Evaluation is on held out data, hypotheses are registered before the pilot rather than after it, and results are published whichever way they fall. We would want the limits stated plainly in that work as well.

09 LIMITS

A model of this kind describes tendencies across many sessions. It does not read minds, and it should not be sold as though it does. Beyond that general statement, five specific limits shape what we will claim.

- **Cold start.** A single short session supports very little. Early reports will be closer to a summary of what happened in the room than to a model of a person, and we would rather say so than dress it up.
- **Selection bias.** Whoever receives a viewing room link is already a filtered population, chosen by a gallery. Aggregate readings describe the attention of people galleries chose to write to, which is not the same as a market.

- **Small numbers per individual.** The strongest claims live at the aggregate level. The individual report is a working hypothesis handed to a professional who knows the person, not a verdict.
- **Attention is not intent.** Long looking correlates with interest, and interest is not a purchase. We report attention, and we do not translate it into predicted revenue.
- **Response, not quality.** The model describes what a person responds to. It says nothing about whether a work is good, and we will resist every commercial incentive to blur that line.

10 WHERE THE PROJECT STANDS

Interlinked, the Cube, is being developed by Solaas Tech S.R.L., based in Bucharest and founded by Andreea Stoian and Alexandru Cojanu. We are preparing an application to the Creative Europe programme, within the Culture strand, for a European Cooperation Project, and we are inviting a small circle of cultural organisations across Europe to shape it with us.

The project answers both objectives of the European Cooperation Projects call. It strengthens the transnational creation and circulation of European works and artists, by moving emerging art across borders through shared rooms and exchanges between partners. It enhances the capacity of the sector to innovate and prosper, by equipping smaller organisations with digital tools, data and new ways of working. A digital preview reduces the physical transport of works, which serves environmental sustainability. The selection brings forward overlooked scenes and artists, which serves inclusion.

For a first application we would work at the small scale of the programme, which asks for three organisations from three countries, and grow toward the medium scale, which asks for five countries, as the partnership matures. The call opens once a year and we are preparing for the 2027 round, expected in the first quarter of that year. The months ahead are for shaping the consortium, agreeing roles and preparing a pilot. After the funded period the platform is built to sustain itself through a light subscription for galleries, so that a network created with public support continues to serve the sector.

AN INVITATION

We are looking for a coordinating organisation with programme experience, partners rooted in emerging scenes, a digital or media art partner, a university laboratory to lead and publish the perception research, and a partner closer to established markets to open the way toward collectors. At this stage we are not asking for a commitment, only for a conversation. Write to stoiana00@gmail.com.

REFERENCES

- ligaya, K., Yi, S., Wahle, I. A., Tanwisuth, K. and O'Doherty, J. P. (2021). Aesthetic preference for art can be predicted from a mixture of low and high level visual features. *Nature Human Behaviour* 5, 743 to 755.
- ligaya, K. et al. (2023). Neural mechanisms underlying the hierarchical construction of perceived aesthetic value. *Nature Communications* 14.
- Reitstätter, L. et al. (2021). Similarity of gaze patterns across physical and virtual versions of an installation artwork. *Scientific Reports* 11.
- Shapson-Coe, A. et al. (2024). A petavoxel fragment of human cerebral cortex reconstructed at nanoscale resolution. *Science* 384.

· · ·