A Cartography of Ambivalence: The Sublime and the Uncanny in the Generative Artificial Intelligence Iconosphere

Andrés Torres Carceller[1]*

[1]Universidad de Barcelona, Spain.

*Correspondence: andrestorres@ub.edu

RITA_24
December 2025
ISSN: 2340-9711
e-ISSN: 2386-7027

Received: 22-08-2025
Revised: 10-09-2025
Accepted: 12-11-2025
Published: 30-12-2025

Resumen

El artículo cartografía la ambivalencia de la iconosfera de la inteligencia artificial generative (IAGen), un ecosistema en el que la imagen deja de testimoniar para sintetizar, y establece un marco operativo para leer, discutir y someter a contraste sus producciones visuales. Para ello, se recurre a una reformulación teórico-conceptual de lo sublime (Burke, Kant) y lo siniestro (Freud) aplicada al paradigma de la IAGen, con el fin de derivar categorías analíticas y criterios de lectura. Como resultado, se definen el sublime algorítmico, asociado al asombro ante la escala de datos, la potencia de cómputo y la combinatoria abierta que desplazan el juicio hacia la infraestructura, y el siniestro sintético, ligado a la inquietud que producen verosimilitudes sin referente, activada por microindicios de fabricación y por una perfección aséptica que vuelve extraña la proximidad. Sobre esa base, se propone un protocolo de análisis en tres capas —tecno-material, semiótico-retórica e infraestructural y de circulación— orientado a distinguir la excelencia formal de la fuerza probatoria.

Palabras clave: Sublime algorítmico, siniestro sintético, inteligencia artificial generativa, cartografía visual, imagen

Abstract

The article maps the ambivalence of the iconosphere of generative artificial intelligence (GenAI), an ecosystem in which images no longer testify to events but synthesize plausible appearances. It provides an operational framework to interpret, discuss, and critically compare these visual productions. To do so, it reformulates the sublime (Burke, Kant) and the uncanny (Freud) for the GenAI paradigm in order to derive analytic categories and workable criteria for interpretation. It defines the algorithmic sublime as astonishment before massive datasets, computing power, and an open combinatorial space, a shift that reorients judgment toward the infrastructures that make such magnitude possible. It defines the synthetic uncanny as unease triggered by referentless verisimilitudes, prompted by micro-cues of fabrication and by an aseptic perfection that makes proximity feel strange. On this basis, the article proposes a three-layer analysis protocol, covering techno-material, semiotic-rhetorical, and infrastructural and circulation layers, designed to distinguish formal excellence from evidential force.

Keywords: Algorithmic sublime, synthetic uncanny, generative artificial intelligence, visual cartography, image

 DOI: https://doi.org/10.24192/2386-7027.2025.v11i24.05

INTRODUCTION

Borges[1] conceived of a library encompassing every possible book: not as a mere repository, but as a limitless horizon where the same volumes recur in varying sequences. This imagery conveys the experience of the unfathomable: proliferating corridors, shelves extending beyond the gaze, and permutations that constantly promise a further variant. It is not a matter of simple abundance, but rather the sensation of facing a magnitude that defies any attempt at mapping.

Figure 1: Anonymous face

Figure note. Image generated by the author on 9 October 2025 using *This Person Does Not Exist*; it does not depict a real person. Tool/source: Wang[2].

If books can be infinite, so too can faces. A face in an extreme close-up holds our gaze and activates an immediate pact of recognition. But what happens when that face belongs to no one? What if what we identify as human is a statistical composition? Once we learn that the image comes from an online generator—This Person Does Not Exist [2]—something shifts. The features, the light, and the framing remain, yet biography is absent. Certainty turns into a subtle unease, as if the gesture of trust were left without a recipient. What seemed like presence reveals itself as a data montage that mimics the human without leaving a trace. These two scenes point to two of the article's central concepts: the sublime and the uncanny. The former alludes to the shock produced by the disproportionate, when an image opens onto a horizon that exceeds all human scale. The latter highlights the unease that arises when what is closest—a face that returns our gaze—appears unmoored from the world. Aesthetic tradition has long theorized these experiences. Burke[3] described a sensibility oriented toward the obscure, the vast, and the incomprehensible. It is not the harmonious pleasure of the beautiful, but a mixture of attraction and awe before what exceeds our measure. Kant[4] distinguished between a "mathematical sublime," linked to the immeasurable, and a "dynamical sublime," associated with forces seemingly capable of annihilating us; in both cases, the tension between imagination and reason becomes a source of insight. Freud[5], for his part, located the unheimlich not in the absolutely strange, but in the familiar returning in an unsettling form. It is not the distant threat of the immense, but the distortion of the intimate.

We inhabit a visual oversaturation: ubiquitous screens, endless feeds, and real-time circulation turn images into a permanent backdrop. We do not merely see more; we see differently. Visibility is apportioned by metrics and recommendations, and attention oscillates between wonder and fatigue. In this context, the experience of looking no longer depends only on objects or styles, but also on the technical and logistical conditions that organize what appears, to whom, and at what pace. Consequently, the aim here is not to elaborate further on diagnoses of excess, but to propose criteria and a way of reading that ties these effects on attention to verifiable conditions of production and circulation.

The iconosphere, as formulated by Cohen-Séat[6], is the audiovisual milieu that surrounds us: not just a collection of images, but the network of devices, supports, exhibition circuits, and viewing habits that shape both what we see and how we see it. Speaking of an iconosphere therefore always implies device and context: how images are produced, scheduled, distributed, indexed, and read. Today, that iconosphere is increasingly algorithmic: massive repositories, generative models, interfaces, and attention metrics condition both what appears and how it is perceived. Historically, mimesis has functioned as a criterion for verisimilitude in images, not merely as resemblance but as an image's capacity to sustain a credible appearance [7]. In the modern iconosphere, this criterion is reconfigured with every technical shift. Likeness no longer depends solely on pictorial resources or manual skill, but is increasingly mediated by technologies of recording and reproduction. In the contemporary transition toward synthesis, a plausible appearance can be produced without reference to a recorded event. This shift intensifies two complementary experiences: astonishment at the scale of production and unease toward plausible appearances that lack verifiable grounding.

Generative AI (GenAI) brings these experiences to the center of visual culture: not only do more images circulate, but plausible appearances are synthesized from models trained on massive datasets [8]. As in Borges's library, where the repetition of disorder generates an Order, these systems compress vast repositories into parameters and recombine prior materials—patterns, styles, correlations—to produce credible novelty. The prompt's ekphrastic phrasing situates the request within latent space, where a statistical procedure selects one among many possible solutions. Novelty no longer arises from a captured event, but from an effective permutation within an unfathomable space of possibilities, in which coherent light, texture, or composition emerges from statistical regularities rather than from a lived scene. This shift from capture to synthesis reconfigures time, authorship, and evidence. The image ceases to bear witness and begins to propose. If the iconosphere determines which images reach us and the conditions under which we perceive them, critical reading must incorporate these mediations—devices, interfaces, and attention metrics—as part of the method rather than treating them as mere context.

This article undertakes a cartography of that iconosphere through two analytical lenses: the algorithmic sublime—the experience of scale, density, and variation that arises when computational synthesis produces images and audiovisuals with a material power that is hard to encompass—and the synthetic uncanny—the specific unease of plausible appearances without referent, where credibility rests more on form than on trace. How does the pact of visual evidence change when what we see may not have occurred? What aesthetic grammars and infrastructural conditions sustain wonder and suspicion? With what tools can we read, verify, and teach these images in an environment of calculated verisimilitude?

THE SUBLIME AND THE UNCANNY: THEORETICAL BACKGROUND

The Sublime in the Classical Tradition

In Edmund Burke, the sublime is an effect on the body and on attention, not a stable property of objects. What threatens or overwhelms—darkness, vastness, emptiness, absolute silence, din, disproportionate magnitude—produces astonishment: a momentary suspension of the faculties and a mixture of fear and attraction. Unlike the beautiful, associated with softness, proportion, and pleasantness, the sublime generates a "negative pleasure" (delight) tied to self-preservation: enjoyment without harm when danger is contemplated at a distance. Burke enumerates efficient causes of the effect (darkness that prevents measuring, power that exceeds resistance, privation, infinity, magnificence) and stresses the phenomenon's somatic basis: muscular tension, heightened sensory alertness, quickened pulse. Aesthetics thus returns to the terrain of affect and embodied perception, against the mere rhetoric of grandiloquence [3,9].

Immanuel Kant reframes the problem as a relation between faculties: imagination (which organizes the sensible) and reason (which thinks totalities). He distinguishes the mathematical and the dynamical sublime [4]. In the mathematical sublime, imagination fails to present as a whole what is incommensurable (the endless series, the infinite); that failure gives way to reason, which can think the unlimited even if it cannot intuit it. In the dynamical sublime, forces of nature that seem capable of annihilating us arouse fear "without danger" for the spectator; again practical reason is affirmed: nature may overwhelm sensibility but not moral vocation. In both cases the affect is mixed: displeasure at imagination's inadequacy and a peculiar pleasure in reason's superiority.

Several Kantian points are operational for later analysis: (1) negative presentation—the infinite is not shown but suggested through imagination's insufficiency; (2) respect (Achtung)—the feeling that accompanies "elevation" is not aesthetic euphoria but a moral tension that makes the law sensible within us; (3) difference from the beautiful—whereas the beautiful harmonizes imagination and understanding (serene pleasure), the sublime dis-harmonizes and, by that route, educates the subject [10-12]; and (4) the safety condition—without the distance that prevents actual harm, terror does not become an aesthetic judgment [3,4]. These aspects provide reading criteria and guard against undue inferences that mistake aesthetic effects for evidence.

Thus Burke contributes the affective materiality of excess and Kant the cognitive architecture of disadjustment. Together they allow us to describe the sublime as a limit-experience that both shakes and educates: what exceeds sensible measure interrupts imagination and compels reason to put that excess in order.

Das Unheimliche in Freud

Before Jentsch[13], the uncanny was linked to an intellectual uncertainty: unease arises when we cannot decide whether something is alive or not (automatons, dolls, wax figures) or when we do not know the rule governing an ambiguous situation. The focus lies on cognitive indecision before liminal objects that disrupt pairs such as animated/inanimate or natural/artificial [13].

Freud[5] shifts the center of the problem. The uncanny (*unheimlich*) does not name what is wholly alien, but the familiar returning in an unsettling form. His philological reading of the pair *heimlich/unheimlich* shows a key ambivalence: *heimlich* means both "domestic, intimate" and "hidden, concealed," so that *un-heimlich* designates what "ought to have remained hidden" and comes to light. The core is no longer perceptual doubt but the return of the repressed: rejected motives, beliefs, or desires reappear in displaced form and produce disquiet.

Freud's discussion with Jentsch is evident in his reading of Hoffmann's "The Sandman": the disorientation before the automaton Olimpia is not explained by the question "Is she alive?" but by the reactivation of unconscious complexes (in particular, anxiety bound to the motif of the eyes and the sinister Coppelius). Freud also orders a repertoire of mechanisms: the double (*Doppelgänger*) as projection of the ego; the omnipotence of thoughts (animistic survivals in modernity); compulsion to repeat (encounters that recur "too much," insistent numbers); fate; and confusion between imagination and reality [5]. In every case, the uncanny does not reside in things themselves but in the affective and symbolic relation that renders them disturbing.

The distinction between life and art also matters. In lived experience certain motifs are immediately uncanny; in fiction, narrative conventions heighten or dampen the effect. The same device may cease to disturb if the work legitimizes it; conversely, "ordinary" devices can become unsettling depending on the framing and credibility of the device. Hence the uncanny is historically situated: it depends on frames of expectation, codes of realism, and prevailing criteria of credibility [5].

Recent contributions refine and extend this framework. For Royle[14], the uncanny operates as a regime of undecidability that disorders structural pairs (inside/outside, own/alien, living/nonliving, natural/artificial) and destabilizes even the language that tries to grasp it. Vidler[15] translates the *unheimlich* to space: architecture and the city become unhomely when the domestic reveals its reverse—endless corridors, repetitive modules, basements—showing that unease is also spatial and infrastructural. Masschelein[16] presents the uncanny as an operative "non-concept": its critical force lies not in fixing a single definition, but in activating tests of defamiliarization that illuminate diverse objects (film, art, theory, media) without closure. From aesthetic culture, Fisher[17] distinguishes weird (irruptions that introduce laws foreign to the world) and eerie (presence without agent or agency without presence), sharpening the affective map beyond individual psychodynamics.

For what follows, two vectors are worth retaining. Disadjusted proximity: the troubling effect does not arise from what is wholly other, but from the near that twists—what resembles us, by excess resemblance or minimal torsion, ceases to be trustworthy. Displaced return: what stabilized experience (motifs, beliefs, reading habits) reappears under a profile that fractures its promise of security. Understood this way, the uncanny is not a stable essence but a relational dynamic with historical anchoring: a mode of reading attentive to when and how the known becomes disturbing as contexts, conventions, and devices shift.

The conceptual references reviewed here allow us to set an operative closure: on the one hand, the sublime is a limit experience that exceeds sensible measure and throws imagination off balance while reason seeks to order that excess; it also includes the thrust of forces that surpass the human and become aesthetic only under a safety distance. On the other hand, the uncanny names a proximity gone askew—the familiar that, through excess resemblance or minimal torsion, becomes unreliable—and the displaced return of motifs or habits that break their promise of stability. Both registers are historical and relational: the signs that activate them (scale, continuity, din or darkness; or else micro-indices, surfaces "too" seamless) shift with habits of sensibility and with criteria of credibility.

FROM THE DIGITAL SUBLIME TO THE ALGORITHMIC SUBLIME

We call the digital sublime the affective regime that emerges with the maturation of the internet as an architecture of flows [18] and that, philosophically, continues the displacement of wonder from nature to technology and its magnitudes [19]. In platform society, the prevailing techno-organizational form is not the isolated object but reprogrammable digital infrastructures that mediate connectivity—APIs, standards, and protocols—and expand through third-party integration and frequent updates [20]. The typical experience is the open-ended feed and real-time metrics: a continuous scroll layering notifications, windows, and tabs. Magnitude no longer resides in a single colossus (the mountain, the ocean) but in the combined density of flows: strata of information and stimuli competing for attention. Hence the mixture of awe and disproportionality: each gesture—click, like, share—barely modulates larger currents that exceed us. In this atmosphere, the sublime takes the form of permanent exposure and speed: an excess that is mobile and diffuse rather than monumental and static.

This shift was anticipated, philosophically, by Lyotard[21] with his diagnosis of inhuman temporalities and the opacity of techno-scientific systems: what exceeds us is no longer necessarily nature but processes and calculations that overflow phenomenological intuition. In the history and sociology of technology, Nye[22] described the technological sublime as a regime in which infrastructures and machines (railroads, skyscrapers, dams) produce grandeur, fear, and fascination comparable to natural landscape. Mosco[18] translated that logic to the digital sublime: an imaginary of transcendence and inevitable progress associated with networks and computing—from "cyberspace" to the "global village"—fueled by spectacular metaphors and myths of the machine. In parallel, Molinuevo Martínez de Bujo[23] outlines the everyday experience of hyperconnection: immediacy, automation, and simultaneity generate feelings of limit and overload relocated in the digital ecosystem. The promise of transparency coexists with pockets of algorithmic opacity, and this mixture of fascination and anxiety is further fueled by an attention economy in which platforms have incentives to design recommender systems oriented toward retention and engagement [24]. Hence its affinity with dynamics of shock [25] and iconophagy [26]: a visual metabolism that devours images while eroding our capacity to examine them critically.

This first horizon does not exhaust the contemporary condition. Alongside it we distinguish the algorithmic sublime, whose origin is not circulation but synthesis. Generative AI (GenAI) does not always point to a single, direct referent: it recombines patterns learned from large datasets and, as the case may be, conditions the output with text or prior images. The result is a synthetic approximation that, instead of recording the world, computes it. Excess no longer comes from the number of posts or the simultaneity of feeds, but from five traits that tend to coagulate into recognizable families:

1. Scale and density. Magnitude is anchored in training corpora of millions of examples and in resolutions that saturate the surface with hyper-detail. Perceptual "quality" shifts toward formal excellences—coherence of lighting, materials, textures—that work as shortcuts to verisimilitude. That "more" is no longer circulatory: it is materialized in each image or frame.

2. Potentially open combinatorics. Seeds, prompts, weights, and ensembles open an unbounded variation: series, grids, and genealogies in which each output invites another. The work ceases to be a fixed object and becomes a field of possibilities. Unlike the scroll—abundance of what already exists—here magnitude is born of a generative space that branches, and that branching is the source of awe.
3. Opacity and inhuman timescales of computation. Between prompt and result lies a statistical itinerary that the interface simplifies without explaining. Months of training condense into milliseconds of inference: a temporal asymmetry that unsettles understanding. This operational opacity is part of the effect: one senses that "something" is operating offstage and at another scale.
4. Distributed authorship. Data, architectures, loss functions, fine-tuning, platform governance, and user operations compose an authorship assemblage that exceeds classical criteria of creativity and intention. Algorithmic curation decides what appears, when, and to whom; the sense of power is attributed less to a person than to the system. That infrastructural prestige functions as symbolic capital in the field.
5. Programmed singularity. "Uniqueness" does not stem from a material here-and-now but from the system state—seed, noise, weights, temperature, ensemble—and from the circulation that positions the piece differentially. What appears "unrepeatable" is an effect of calculation and of its curation. The traditional aura does not return; an aura of performance emerges.

If the digital sublime is anchored in excess by circulation (quantity, speed, simultaneity), the algorithmic sublime is anchored in excess by generation (data scale, visible density, combinatorics, opacity, and distributed authorship). Both entail disproportionality and wonder, but with distinct grammars. In the first, the subject experiences smallness before the flow; in the second, before the power of synthesis that materializes possible worlds with a formal consistency that withdraws from direct verification. In attentional terms, the digital relies on shock and iconophagy; the algorithmic on formal finishes that impose themselves as aesthetic warrants where the referent is absent or irrelevant.

This turn relocates the authority of judgment. Faced with formal display and infrastructural prestige, spectators tend to suspend demands of provenance and verification, in line with critiques of systemic opacity and real-time living [27]. At the same time, it prolongs—under contemporary conditions—the shift that relocates natural grandeur in the technical [19]: where once there was mountain or storm, today there is model, dataset, and pipeline operating offstage. In the passage from the inexhaustible feed to the inexhaustible model —from "everything circulates" to "everything can be generated"—the algorithmic sublime takes shape as a distinctive affective regime of GenAI iconosphere. At the same time, it extends —under contemporary conditions— the shift of the sublime from natural grandeur to technical infrastructures operating at massive scale [28], so that where there were once mountains or storms, there are now models trained on billions of encoded images and a pipeline of data collection and filtering that remains outside the frame [29].

Table 1: Algorithmic Sublime: Operational Map

Axis	Operational definition	Infrastructural drivers	Analytical cues / mitigation
Scale & density	Massive training corpora; hyper-detailed outputs; formal coherence as shortcut to plausibility.	Model size/version; dataset curation; super-resolution; denoise / cleanup pipelines.	Document dataset / model / resolution; separate formal quality from reference; check for super-resolution/denoise footprints.
Combinatorics	Prompts/seeds/weights enable ramifying series; the work functions as a space of states rather than a single object.	Seed control; samplers; ensembles; LoRA; platform affordances for series.	Report seeds/iterations/branching; state selection criteria; archive representative variants.
Opacity & nonhuman time	Months of training condensed into millisecond inference; pipeline hidden by the interface.	Closed / proprietary models; opaque pipelines; API constraints; hardware acceleration.	Specify model/version; distinguish training vs. inference; state compute used and constraints.
Distributed authorship	Authorship as assemblage (data, architecture, tuning, operator, platform curation); reception shaped by recommendation.	Platform curation; ranking / metrics; A / B-tested presentation.	Credit pipeline actors; disclose platform's role; avoid single-author fallacies.
Programmed singularity	"Uniqueness" produced by system state (seed / weights / temperature) + circulation logic, not indexical capture.	Session hashes; content IDs; recommendation triggers; tokenization schemes.	Record system state; publish provenance; explain limits of repeatability/identity.

Note. Abbreviations: DOF = depth of field; LoRA = low-rank adaptation.

SYNTHETIC UNCANNY AND THE SOCIO-TECHNICAL UNCANNY VALLEY

We define the synthetic uncanny as the affective misalignment that arises when an image or audiovisual sequence appears plausible without referring to any real-world event. Denotation —what is in the image—loses its external grounding, while connotation—our cultural interpretation—becomes suffused with suspicion [30]. This drift aligns with the simulacrum: signs that no longer conceal or reveal an original but circulate as self-sufficient appearances [31]. Unease does not come from a visible trick, but from the success of verisimilitude. Unlike photography —where denotation could function as the trace of an event— in synthesis, form (the coherence of light, materials, continuity) acts as an implicit guarantor. Judgment shifts from referent to surface: we evaluate formal performance rather than provenance in the world.

To clarify this mechanism, it is useful to revisit the uncanny valley. In his formulation, Mori[32] described a non-linear relationship between the degree of human likeness and affective response: as appearance closely approaches the human without fully reaching it, affinity drops—with movement intensifying the fall—before recovering at the edge of indistinguishability. Mori, et al.[33] called the familiarity axis shinwakan and its negative pole bukimi (eeriness). The key point here is that strangeness is not a simple lack of liking, but a specific affect that emerges from minimal mismatches within a context of high verisimilitude. Reread today, that valley is a shifting sociotechnical threshold: it moves with model versions, datasets, interfaces, and viewing habits [34].

Subsequent literature has refined this intuition. Ho, and MacDorman[35] show that eeriness is not simply the reverse of warmth; they are partially independent dimensions. They also propose separate indices for humanness, eeriness, and warmth, and document that strangeness increases when minor imperfections appear near the human threshold, especially in motion. This framing reinforces that the valley is not a universal "symptom" of the object, but a zone of undecidability that depends on perceptual and technical context.

In still images, unease is activated along two complementary paths. Micro-cues: weak signals that break fluency once detected (reflections with no plausible source; incongruent depth of field —DOF— or bokeh; implausible continuity of hair or skin–textile edges; dental micro-asymmetries that "don't add up"). Polished perfection: the absence of irregularities proper to the living (excessively regular pores, skin without micro-imperfections, unworn textiles, dust-free surfaces). Without any flagrant error, uniformity becomes unnatural.

In audiovisual material, the valley widens because the illusion of life depends on fine temporal dynamics. Typical cues include minimal lip-sync offsets, overly regular prosody without micro-breath packages, blinks with atypical latencies, rigid saccades, or facial inertia between gestures that fails to "carry over" as a living face would. At the other extreme, extreme cleanliness —everything correct, nothing "off"— can itself be unsettling through over-correction.

This ambivalence finds recent empirical support. Nightingale, and Farid[36] show that GAN-synthesized faces can be indistinguishable from real ones—and even judged more trustworthy—by human observers. On the forensic side, Verdoliva[37] describes a co-evolution: as detection learns to exploit specific traces (upsampling patterns, compression signatures, geometric inconsistencies), new architectures attenuate them, shifting the acceptance threshold yet again. Audiences oscillate between habituation (normalizing the algorithmic look) and hypervigilance (hunting for microscopic cues), with risks of both false positives and false negatives. This pattern aligns with recent evidence on fluency and subjective validity, suggesting that ease of processing can function as a cue for truth or familiarity in the absence of additional support [38]. The effect does not reside solely in the image; it depends on the channel and the interface. Platform compression, denoising, frame interpolation, screen size, and zoom capabilities modulate which details emerge and how much attention they receive. Paratexts (headlines, captions, hashtags), perceived authorship, and recommendation norms predispose audiences toward trust or suspicion. The same clip may pass unnoticed in a high-trust feed yet trigger unease in a skeptical environment. Not every generative piece activates this register. Playful, parodic, or overtly fantastic works shift the reading, while unease arises primarily when plausibility is intended and the ecology of circulation demands evidence.

Operational criteria:

1. Baseline condition. Treat surface plausibility as non-probative: physical coherence (light, materials, continuity) can produce verisimilitude without reference.

2. Path A — Micro-cues. Inspect high-frequency areas, skin–textile edges, secondary reflections, and contact shadows; in video, fine lip-sync latencies, blink patterns, and prosodic micro-rhythm.
3. Path B — Polished perfection. Consider the absence of "biographical noise" a potential marker of synthesis (and check for default denoise/beauty).
4. Socio-technical threshold. Update criteria with each model release and platform habit; avoid essentialisms such as "you can always tell."

Reading protocol. First, evaluate under real viewing conditions; then re-examine under controlled conditions (higher resolution, pause/advance, no platform filters). Complete with a minimal reconstruction of fabrication (data, model, controls, pre/post-process) and with publication traceability (metadata and context). When available, integrate content credentials as a process trace, clarifying their scope.

Table 2: Synthetic Uncanny: An Operational Map

Axis	Operational Definition	Sociotechnical factors shifting the valley	Reading / mitigation Cues
Referentless verisimilitude	Plausible appearance without a traceable world-event; denotation detached; connotation under suspicion.	Model upgrades; dataset curation; default beauty filters; platform compression/ interpolation. Render / compression pipeline.	Treat surface plausibility as non-evidentiary; anchor with provenance/context; test "signs of reality" against production conditions.
Micro-cues	Weak signals that break the spell once noticed.	New model versions; beautification defaults; UIs that hide metadata/ recompression.	Inspect high-frequency zones, skin–textile borders, secondary reflections, contact shadows.
Polished perfection	Absence of "biographical noise" (imperfections and wear).	Platform aesthetics; aggressive denoise; voice-cloning with heavy post.	Reintroduce friction/context (making-of, declared limits, synthesis disclosure).
Context & interface	Channel modulates credibility and scrutiny.	Recommendation policies; captions/ hashtags; audience culture; player defaults (autoplay/interpolation)	Situated evaluation: who publishes, where, for whom; avoid decontextualized inspection.
Plastic threshold	The "valley" moves with models, habits, and pedagogy.	New model baselines; viewing habits; platform tutorials/defaults.	Update criteria; avoid essentialisms ("you can always tell").

Note. Abbreviations: DOF = depth of field; UI = user interface. "Valley" refers to the uncanny valley understood as a sociotechnical threshold.

Two operational lenses guide our reading of the IAGen iconosphere. The algorithmic sublime names the vertigo of generative excess —scale, density, variation, opacity, and distributed authorship— which shifts judgment toward formal performance and the prestige of infrastructure. The synthetic uncanny designates the unease of verisimilitude without a referent, activated by micro-cues or by polished perfection, within a mobile socio-technical uncanny valley (models, datasets, interfaces, platform habits). Methodologically, it is advisable to proceed in two passes: (1) examine the piece under real viewing conditions; (2) re-examine it under controlled conditions (higher resolution, pause/advance, without platform filters), and complete the assessment with a minimal reconstruction of the pipeline and publication traceability. The aim is not to unmask by default, but to make explicit what guarantees each layer (techno-material, semiotic-rhetorical, infrastructural) actually provides and where their limits lie; these lenses do not replace verification criteria—they offer a reading framework for distinguishing effects of scale/ combination from plausible appearances without a referent.

POST-PHOTOGRAPHIC GENEALOGY AND THE POST-SIMULACRUM: FROM REPRODUCIBILITY TO SYNTHESIS

The recent history of the technical image can be read as a gradual shift from the logic of copying to the logic of synthesis. Along the axis of reproducibility, Benjamin[39] argued that technical serialization weakened —and at the same time reshaped— the uniqueness of the work. Value tied to singular presence gave way to regimes of multiplication and circulation, opening new forms of reception and legitimation. This shift did not amount to a complete disappearance, but to a transition toward "cold auras" sustained by mediations, contexts, and display devices that restore singularity through circulation [40]. In the same turn, a technological sublime takes hold. The grand no longer anchors itself in nature but in machines and infrastructures; what exceeds is no longer the landscape, but the technical power that replaces or duplicates it [41].

Digital culture consolidates postphotography as a regime in which the photographic index loses centrality and the image is defined by its place within flows, archives, and operations—less trace than operation, less capture than

management [42]. In its most recent iteration, this arc culminates in the post-simulacrum. It is no longer only a matter of images replacing the real, but of images competing with it from an autonomous universe of visual production governed by platforms and synthesis models [43]. The move from reproducibility to synthesis does not cancel photography; it subsumes it as a repertoire from which appearances are computed.

A key element in this genealogy is the database image. Rubinstein, and Sluis[44] describe how photography becomes data that is taggable, aggregatable, retrievable, and recombinable, so that its identity depends less on the shutter release than on its position within collections, metadata, and queries. From this perspective, aesthetics is no longer only a matter of framing and light; it also involves archive architecture and a grammar of querying that shapes which criteria classify, how taxonomies group, and which styles emerge from statistical patterns [45]. GenAI inherits and intensifies this principle because it is trained on massive corpora of image–text pairs and, in the process, absorbs biases, regularities, and representational gaps present in the dataset itself [46]. Synthesis tends to reproduce those statistical regularities, producing plausible appearances and, in some cases, even retrieving material from the training data.

This arc illuminates the two registers developed above. On the one hand, the algorithmic sublime feeds on the magnitude and combinatorics of the massive archive: training sets, model scalability, and the branching of prompts and seeds turn the "life of the archive" into a feeling of excess. The work ceases to be an object that multiplies and becomes a *state space* to be explored. Disproportion stems both from visible density (hyper-detail, global coherence of lighting and material) and from the awareness that there is *always another variant* one computational step away. On the other hand, the synthetic uncanny emerges when that grammar —learned as statistical verisimilitude— produces plausible presences without a worldly trace: denotation is emptied and connotation is traversed by a structural suspicion. Where reproducibility still referred, however faintly, to an original, synthesis offers appearances without antecedent.

The notion of post-simulacrum sharpens the difference with the classical semantics of the simulacrum. It is not mere concealment or substitution of the real, but competitive convergence: synthetic images and empirical records share platforms, interfaces, and metrics; they compete for attention with the same quality grammar (sharpness, coherence, resolution) under regimes of visibility governed by algorithmic curation. Photography loses its ontological privilege; "proof" migrates toward formal performance and contexts of circulation.

A methodological consequence follows: authenticity becomes documentary and procedural. If indexicality is no longer sufficient, an image's grounding shifts toward traces of making (metadata, edit logs), its archival ecology (versioning, publication trails), and platform norms (tagging policies, provenance seals). Hence efforts to standardize provenance and traceability. These do not guarantee truth, but they restore context and delimit what can be warranted; assurance shifts from the supposed purity of the medium toward process documentation and conditions of circulation [47]. This consequence sets the stage for the reading and verification protocol presented in the following section.

GenAI both extends and disrupts the history of technical images. It continues the automation, circulation, and editing associated with photography and cinema, yet it breaks their indexical principle by no longer recording an event but generating it statistically. The algorithmic sublime and the synthetic uncanny do not emerge from nowhere; they draw on this history. The former turns the magnitude of the archive and the model into an experience of excess, while the latter turns calculated plausibility into unease through a proximity without a referent. In this framework, analyzing GenAI images requires fewer inventories of effects and more maps of procedures, vocabularies for reading archives, and protocols that return to reception what technical finish tends to conceal.

READING AND VERIFICATION PROTOCOLS: TECHNO-SEMIOTICS AND INFRASTRUCTURE

This section turns the conceptual framework into a procedure for analyzing images and audiovisual pieces generated with IAGen. It rests on two premises: no reading is neutral with respect to the technical support and its ecology of circulation, and formal plausibility does not amount to reference. The aim is to reintroduce interpretive friction: to make processes and limits visible so as not to confuse a polished finish with evidence.

A) Techno-material layer: how it is made and why it has effect

Document data, model, parameters, and stages. For data: size and diversity of the set, curation criteria, and licenses. For the model: type and version, whether open or closed, available controls. For parameters: prompts (including negatives), seed, temperature, number of iterations/branches, ensembles and/or LoRA adapters. Distinguish training (prolonged) from inference (fast) to situate the temporal asymmetry that often feeds opacity. Record pre- and post-processing/postproduction (upscalers, denoise, "beauty" filters, color grading, frame interpolation) and the distributed authorship of the pipeline (data team, architecture design, fine-tuning, operator, platform curation). Finally, explain programmed singularity: how the system state (weights, seed, combinations) and the publication circuit generate the appearance of irrepeatability.

What it guarantees—and what it does not. This layer clarifies causalities (why there is hyper-detail, why a voice sounds "clean," why the series branches), but it does not accredit that the event occurred. It helps explain the algorithmic sublime (scale, density, combinatorics, opacity) and also why certain "over-polished" finishes appear (post-processing, normalizations).

B) Semiotic-rhetorical layer: which signs of reality it mobilizes and how it persuades

Analyze the surface and its rhetorics. First, physical signs of coherence: lighting and contact shadows, materials, spatial/temporal continuity, depth of field/bokeh. Make their scope explicit: they guarantee internal consistency, not the existence of a referent. Second, weak cues that destabilize reading (implausible secondary reflections, improbable skin-to-textile edges, "impossible" micro-asymmetries; in audiovisual, fine lip-sync latencies, blink patterns, facial inertia, prosodic microrhythms). Third, the reverse: smooth perfection (absence of "biographical noise": overly regular pores, unworn textiles, voices without micro-breath) as a possible marker of synthesis. Identify rhetorics of verisimilitude (sharpness, global coherence, a learned "look" from a genre/style) and rhetorics of excess (hyper-detail, seriated morphogenesis), avoiding the move from formal performance to proof. What it guarantees and what it does not. This layer delimits what form allows one to conclude and where external context becomes necessary. It connects directly with the synthetic uncanny (micro-indicia and smooth perfection) and, by contrast, with the formal excellence typical of the algorithmic sublime.

C) Infrastructural layer: where it circulates, who prioritizes it, and with what traces

Situate the piece within its ecology of circulation. Describe channel and format (platform, feed, player, compression, possibility of zoom), recommendation/curation policies, visible social signals (metrics, trending), and audience expectations (skeptical forum vs. trusted environment). Where present, extract content credentials [48]: which production/editing steps are signed, by whom, when, and with which tools—clarifying their scope (they provide process traceability; they are not, by themselves, factual truth). Note how the environment amplifies certain effects (monumental formats, autoplay, dashboards and counters intensify excess) or softens cues (cosmetic filters, denoise, interpolation).

Guarantees-and-limits report (recommended closure for each analysis). Conclude with a concise report that states, first, what is known (provenance and edit chain; technical aspects of the process—data where available, model/version, relevant controls and parameters, pre- and post-processing; observable physico-semiotic consistencies; and context of circulation); second, what is not known (opaque stretches of the dataset or fine-tuning, missing traces across platforms, and limits of the expertise itself); third, what each layer does and does not provide (the techno-material layer explains procedures without validating occurrence; the semiotic-rhetorical layer establishes formal plausibility without proving reference; the infrastructural layer accredits authorship/editing and chain of trust where traces exist, without in itself constituting factual proof); and, finally, the residual risks (possible false positives/negatives, habituation or hypervigilance biases, and channel effects on perception). The logic is to reintroduce interpretive friction without cancelling wonder: to make how and where visible in order to situate the what. The three layers complement one another and allow one to apply, where appropriate, the two reading lenses without confusing formal excellence with evidence or suspicion with certainty. This closing step matters because, in an environment where surface verisimilitude can come apart from worldly anchoring and platform conditions change how we see, the report makes judgments comparable, auditable, and accountable: it makes assumptions explicit, delimits the scope of conclusions, and reduces both uncritical enthusiasm and indiscriminate suspicion.

CONCLUSIONS

Rather than mere labels, the algorithmic sublime and the synthetic uncanny are modes of attention. They help us move beyond the dichotomies of "works vs. deceives" and "real vs. fake," restoring to GenAI images their technical, historical, and sensory density. If this trajectory reveals anything, it is that the current iconosphere cannot be understood solely by what it shows, but by how it orients us, whether to surrender to generative magnitude or to distrust impeccable proximity. These lenses do not replace verification criteria; they offer a descriptive framework for distinguishing effects of scale and combinatorics from plausible appearances without a referent, and for reintroducing process and context into reading.

In the algorithmic sublime, we can recognize a politics of scale that shifts judgment toward formal performance and the prestige of infrastructures. Its philosophical contribution is not merely aesthetic. It sketches an economy of the possible, in which the unfathomable no longer refers to nature but to a state space that expands with computation. When this vertigo becomes habitual, it reshapes the notion of measure. It teaches us to calculate differently, yet it can also lead us to cede authority to what presents itself as unreachable.

The synthetic uncanny, by contrast, does not denounce isolated "tricks." It points to a fragility in the pact of evidence that is activated precisely when everything seems in order. Its value is not to moralize suspicion, but to refine it. It trains us to detect when form demands belief, and what the gaze loses when it mistakes verisimilitude for proof. Here philosophy does not offer a verdict, but an art of hesitation: holding doubt a few seconds longer in order to ask about provenance, the scene of circulation, and the conditions of possibility of seeing.

Together, these lenses suggest an ethics for our shared visual space that is less defensive and more curatorial. The point is not to close off the inventive power of synthesis, which would impoverish the world, but to govern its ambivalence by arranging contexts, making limits explicit, and fostering habits of reading that turn wonder into a criterion and suspicion into care. We might call this a politics of attention, training the gaze to oscillate without cynicism between openness to the new and a demand for context and provenance. If contemporary visual culture rewards what dazzles while normalizing what unnerves, reading in a double key is not a hermeneutic luxury; it is a prudent way of inhabiting an ecosystem where evidence is not a given. Between an excess that expands and a proximity that unsettles, what is at stake is not only how we see, but what forms of shared interpretation we can sustain through our collective gaze.

We do not seek to dissolve this ambivalence, but to map it so that it does not govern us blindly. If the GenAI iconosphere draws a line between "everything can be generated" and "everything can appear so," these two lenses do not hand down verdicts; they open operational questions and restore the role of analysis. In an environment where attention is captured by what astonishes and trust is won by what does not jar, looking in a double key is no theoretical ornament. It is the minimum rigor required to inhabit images that no longer ask for faith, but for reading.

REFERENCES

1. BORGES, Jorge Luis. "La biblioteca de Babel." In *Ficciones*, Debolsillo, 2011. https://books.google.com.pk/books?id=bx7DygAACAAJ

2. WANG, Philip. "*This person does not exist*." https://this-person-does-not-exist.com/en

3. BURKE, Edmund. *Indagación filosófica sobre el origen de nuestras ideas acerca de lo sublime y de lo bello*. Alianza editorial, 2023. https://www.alianzaeditorial.es/libro/filosofia/indagacion-filosofica-sobre-el-origen-de-nuestras-ideas-acerca-de-lo-sublime-y-de-lo-bello-edmund-burke-9788411483636/

4. KANT, Immanuel. *Crítica del juicio*. Alianza, 2003. https://www.filosoficas.unam.mx/docs/556/files/kant_-_critica_del_juicio%20(extractos).pdf

5. FREUD, Sigmund. "Lo ominoso." In *Obras completas*, vol. XVII, Biblioteca Nueva, 2006. https://www.psicopsi.com/wp-content/uploads/2021/05/Freud-Amorrortu-17.pdf

6. COHEN-SÉAT, Gilbert. *Essai sur les principes d'une philosophie du cinéma*. Presses universitaires de France, 1946. https://books.google.com.pk/books?id=lDxNxwEACAAJ

7. ARAZI, Alejandro Cohen. "¿ Qué es el Valle Inquietante? De Aristóteles a la Inteligencia Artificial." *Reflexión Académica*, 2025, no. 58, pp. 143–146. https://dspace.palermo.edu/ojs/index.php/rad/article/view/12844

8. SOMAINI, Antonio. "Algorithmic images: Artificial intelligence and visual culture." *Grey Room*, 2023, no. 93, pp. 74-115. DOI: https://doi.org/10.1162/grey_a_00383

9. SHUSTERMAN, Richard. "Somaesthetics and Burke's Sublime." *The British Journal of Aesthetics*, 2005, vol. 45, no. 4, pp. 323–341. DOI: https://doi.org/10.1093/aesthj/ayi047

10. CLEWIS, Robert R. *The Kantian sublime and the revelation of freedom*. Cambridge University Press, 2009. https://books.google.com.pk/books?id=PCvT_ZivMOAC

11. GUYER, Paul. *Kant and the experience of freedom: Essays on aesthetics and morality*. Cambridge University Press, 1993. https://books.google.com.pk/books?id=De1uZ3DF5m4C

12. SHAW, Philip. *The Sublime*. Routledge, 2006. https://books.google.com.pk/books?id=3HnOLEMSjPMC

13. JENTSCH, Ernst Anton "Zur Psychologie des Unheimlichen." *Psychiatrisch-Neurologische Wochenschrift*, 1906, vol. 8, no. 22, pp. 195–198. https://publikationen.ub.uni-frankfurt.de/frontdoor/index/index/docId/19552

14. ROYLE, Nicholas. *The uncanny*. Manchester University Press, 2003. https://books.google.com.pk/books?id=XkvSWxjrMN8C

15. VIDLER, Anthony. *The architectural uncanny: Essays in the modern unhomely*. MIT press, 1994. https://books.google.com.pk/books?id=kVKxBIB-En4C

16. MASSCHELEIN, Anneleen. *The unconcept: The Freudian uncanny in late-twentieth-century theory*. State University of New York Press, 2012. https://books.google.com.pk/books?id=9XgohiN3vOwC

17. FISHER, Mark. *The Weird and the Eerie*. Repeater Books, 2016. https://books.google.com.pk/books?id=XeNYDQAAQBAJ

18. MOSCO, Vincent. *The digital sublime: Myth, power, and cyberspace*. MIT Press, 2005. https://books.google.com.pk/books?id=cenuCwAAQBAJ

19. SERRANO, Inmaculada Murcia. "Lo sublime de Edmund Burke y la estetización postmoderna de la tecnología." *Revista de Estética y Teoría de las Artes*, 2009, vol. 8. https://d1wqtxts1xzle7.cloudfront.net/58596347/murcia-libre.pdf

20. CHIMENTI, Gianluca; HAGBERG, Johan and ARAUJO, Luis. "Platforms, infrastructures and the Futures of market society." *Journal of Business Research*, 2025, vol. 189, p. 115167. DOI: https://doi.org/10.1016/j.jbusres.2024.115167

21. LYOTARD, Jean-François. *Lo inhumano. Charlas sobre el tiempo*. Manantial, 1988. https://books.google.com.pk/books?id=3ie9swEACAAJ

22. NYE, David E. *American technological sublime*. MIT Press, 1996. https://books.google.com.pk/books?id=LrdbOJFxWIoC

23. MOLINUEVO MARTÍNEZ DE BUJO, José Luis. "El sentimiento de los límites y lo sublime tecnológico." In *Filosofía y razón: Kant, 200 años*, edited by VÁZQUEZ, M. E and DE LA CALLE, R., PUV, 2005, pp. 75-91. https://dialnet.unirioja.es/servlet/articulo?codigo=5063137

24. SÖDERLUND, Kasia; ENGSTRÖM, Emma; HARESAMUDRAM, Kashyap; LARSSON, Stefan and STRIMLING, Pontus. "Regulating high-reach AI: On transparency directions in the Digital Services Act." *Internet policy review*, 2024, vol. 13, no. 1, pp. 1-31. DOI: https://doi.org/10.14763/2024.1.1746

25. KLEIN, Naomi. *La doctrina del shock: el auge del capitalismo del desastre*. vol. 1, paidós 2007. https://revistainvi.uchile.cl/index.php/INVI/article/download/62340/65999

26. BAITELLO, Norval. *A era da iconofagia: ensaios de comunicacao e cultura*. Hacker Editores, 2005. https://books.google.com.pk/books?id=yhioPgAACAAJ

27. MOLINUEVO MARTÍNEZ DE BUJO, José Luis. *La vida en tiempo real: la crisis de las utopías digitales*. Biblioteca Nueva, 2006. https://portalcienciaytecnologia.jcyl.es/documentos/63cf36c80eea6d35507ca46a

28. DEŽMAN, Dominik Vrabič. "Interrogating the deep blue sublime: images of artificial intelligence in public media." In *From hype to reality: Artificial intelligence in the study of art and culture*, edited by CETINIĆ, Eva; Negueruela del Castillo, Darío, vol. 3, Bibliotheca Hertziana – Max Planck Institute for Art History, 2024, pp. 57–66. DOI: https://doi.org/10.48431/hsah.0307

29. SCHUHMANN, Christoph; BEAUMONT, Romain; VENCU, Richard; GORDON, Cade; WIGHTMAN, Ross; CHERTI, Mehdi et al.JITSEV, Jenia. "LAION-5B: An open large-scale dataset for training next generation image-text models." *Advances in Neural Information Processing Systems, 35*, vol. 35, 2022, pp. 25278-25294. DOI: https://doi.org/10.48550/arXiv.2210.08402.

30. ECO, Umberto. *Semiótica y filosofía del lenguaje. Lumen*, 1990. https://books.google.com.pk/books?id=YvORSAAACAAJ

31. BAUDRILLARD, Jean. *Cultura y simulacro*. kairós, 1993. https://www.editorialkairos.com/catalogo/p/cultura-y-simulacro

32. MORI, Masahiro. "Bukimi no tani [The uncanny valley]." *Energy*, 1970, vol. 7, p. 33. https://cir.nii.ac.jp/crid/1370013168736887425

33. MORI, Masahiro; MACDORMAN, Karl F and KAGEKI, Norri. "The uncanny valley [From the field]." *IEEE robotics & automation magazine*, 2012, vol. 19, no. 2, pp. 98-100. DOI: https://doi.org/10.1109/MRA.2012.2192811

34. MACDORMAN, Karl F and ISHIGURO, Hiroshi. "The uncanny advantage of using androids in cognitive and social science research." *Interaction studies*, 2006, vol. 7, no. 3, pp. 297-337. DOI: https://doi.org/10.1075/is.7.3.03mac

35. HO, Chin-Chang and MACDORMAN, Karl F. "Revisiting the uncanny valley theory: Developing and validating an alternative to the Godspeed indices." *Computers in Human Behavior*, 2010, vol. 26, no. 6, pp. 1508-1518. DOI: https://doi.org/10.1016/j.chb.2010.05.015

36. NIGHTINGALE, Sophie J and FARID, Hany. "AI-synthesized faces are indistinguishable from real faces and more trustworthy." *Proceedings of the National Academy of Sciences*, 2022, vol. 119, no. 8, p. e2120481119. DOI: https://doi.org/10.1073/pnas.2120481119

37. VERDOLIVA, Luisa. "Media forensics and deepfakes: an overview." *IEEE journal of selected topics in signal processing*, 2020, vol. 14, no. 5, pp. 910-932. DOI: https://doi.org/10.1109/JSTSP.2020.3002101

38. VOGEL, Tobias; SILVA, Rita R; THOMAS, Aurelia and WÄNKE, Michaela. "Truth is in the mind, but beauty is in the eye: Fluency effects are moderated by a match between fluency source and judgment dimension." *Journal of Experimental Psychology: General*, 2020, vol. 149, no. 8, pp. 1587-1596. DOI: https://doi.org/10.1037/xge0000731

39. BENJAMIN, Walter. *La obra de arte en la época de su reproductibilidad técnica*. Editorial Ítaca, 2003. https://editorialitaca.com/libro/la-obra-de-arte-en-la-epoca-de-su-reproductibilidad-tecnica

40. BREA, José Luis. *Las auras frías*. Anagrama 1991. https://d1wqtxts1xzle7.cloudfront.net/38025420/Brea__Jose_Luis_-_Las_auras_frias-_el_culto_a_la_obra_de_arte_en_la_...-libre.pdf?1435610086

41. RONCORONI-OSIO, Umberto. "Benjamin y la teoría de lo sublime tecnológico." *Contratexto*, 2007, no. 15, pp. 133-140. DOI: https://doi.org/10.26439/contratexto2007.n015.777

42. SONESSON, Göran. "Post-photography and beyond." *From Mechanical Reproduction to Digital Production, Visio*, 1999, vol. 4, no. 1, pp. 11-36. https://d1wqtxts1xzle7.cloudfront.net/32553963/VisioPhoto-libre.pdf

43. FONTCUBERTA, Joan. *Desbordar el espejo*. Galaxia Gutenberg, 2024. https://sgfm.elcorteingles.es/SGFM/dctm/MEDIA03/202403/25/00106510646224____2_.pdf

44. RUBINSTEIN, Daniel and SLUIS, Katrina. "A life more photographic: Mapping the networked image." *Photographies*, 2008, vol. 1, no. 1, pp. 9-28. DOI: https://doi.org/10.1080/17540760701785842

45. CARCELLER, Andrés Torres. "The ARTificial revolution: Challenges for redefining art education in the paradigm of generative artificial intelligence." *Digital Education Review*, 2024, vol. 45, pp. 84-90. DOI: https://doi.org/10.1344/der.2024.45.84-90

46. BIRHANE, Abeba; PRABHU, vinay; HAN, Sanghyun; BODDETI, Vishnu and LUCCIONI, Sasha. "Into the LAION's Den: Investigating hate in multimodal datasets." *Advances in Neural Information Processing Systems 36 (NeurIPS 2023), Datasets and Benchmarks Track*, vol. 36, 2023, pp. 21268-21284. https://arxiv.org/abs/2311.03449.

47. COALITION FOR CONTENT PROVENANCE AND AUTHENTICITY (C2PA). *C2PA Technical Specification* 2024. https://spec.c2pa.org/specifications/specifications/1.4/specs/C2PA_Specification.htmlhttps://spec.c2pa.org/specifications/specifications/1.4/specs/C2PA_Specification.html

48. CONTENT AUTHENTICITY INITIATIVE (CAI). *Content Credentials: Technical Overview*. 2022. https://contentauthenticity.orghttps://contentauthenticity.org