architecture

Article

Design in the Age of Predictive Architecture: From Digital Models to Parametric Code to Latent Space

José Carlos López Cervantes * and Cintya Eva Sánchez Morales

Department of Graphic Expression in Architecture and Engineering, School of Architecture, University of Granada, 18071 Granada, Spain; cysamo@ugr.es
* Correspondence: joseclc@ugr.es

Abstract

Over the last three decades, architecture has undergone a sustained digital transformation that has progressively displaced the ontology of the geometric generator, understood here as the primary artefact through which form is produced, controlled, and legitimized. This paper argues that, within one extended digital epoch, three successive regimes have reconfigured architectural agency. First, a digital model regime, in which computer-generated 3D models become the main generators of geometry. Second, a parametric code regime, in which scripted relations and numerical parameters supersede the individual model as the core design object, defining a space of possibilities rather than a single instance. Third, an emerging latent regime, in which diffusion and transformer systems produce high plausibility synthetic images as image-first generators and subsequently impose a post hoc image-to-geometry translation requirement. To make this shifting paradigm comparable across time, the paper uses the blob as a stable morphological reference and develops a comparative reading of four blobs, Kiesler's Endless House, Greg Lynn's Embryological House, Marc Fornes' Vaulted Willow, and an author-generated GenAI blob curated from a traceable AI image archive, to show how the geometric generator migrates from object, to model, to code, to latent image-space. As a pre-digital hinge case, Kiesler is selected not only for anticipating blob-like continuity, but for clarifying a recurrent disciplinary tension, " form first generators" that precede tectonic and programmatic rationalization. The central hypothesis is that GenAI introduces an ontological shift not primarily at the level of style, but at the level of architectural judgement and evidentiary legitimacy. The project can begin with a predictive image that is visually convincing yet tectonically underdetermined. To name this condition, the paper proposes the plausibility gap, the mismatch between visual plausibility and tectonic intelligibility, as an operational criterion for evaluating image-first workflows, and for specifying the verification tasks required to stabilize them as architecture. Selection establishes evidentiary legitimacy, while a friction map and Gap Index externalize the translation pressure required to turn predictive imagery into accountable geometry, making the plausibility gap operational rather than merely asserted. The paper concludes by outlining implications for authorship, pedagogy, and disciplinary judgement in emerging multi-agent design ecologies.

Keywords: digital design; parametric design; generative artificial intelligence; diffusion models; latent space; geometric generator; architectural ontology; plausibility gap; image based design

Academic Editors: Tilanka Chandrasekera and Mallikarachchige Nishan Rasanga Wijetunge

Received: 15 December 2025
Revised: 26 January 2026
Accepted: 2 February 2026
Published: 10 February 2026

1. Introduction

This article evaluates the impact of generative artificial intelligence on what has been called digital architecture since the 1990s. The goal is not to decide whether AI constitutes a separate “revolution” within the digital revolution, but to determine whether it introduces a paradigm shift in architectural design, specifically in its operative germ: what is here termed the geometric generator, the artefact or instance from which form is derived and project agency is organized. The underlying question is therefore double: does GenAI represent an inflexion point that inverts the design process or is it the logical culmination, intensification, and automation of procedures already present in the mature phases of digital architecture? [1,2]

In recent historiography, authors such as Mario Carpo have described the evolution of digital architecture through a sequence of “turns”, distinguishing successive transformations in tools, methods, and notions of authorship [2]. From this perspective, a common contemporary diagnosis holds that AI does not inaugurate a new epoch, but automates imitative transfer, learning stylistic correlations from massive corpora and scaling logics already implicit in mass customisation and serial variation [1,2]. Read this way, the move toward GenAI would be a radically accelerated continuity: a spectacular intensification rather than a rupture.

This article adopts a critical position toward that strong continuity thesis. Without denying genealogy, it argues that we are facing an ontological shift, because the location of the geometric generator changes, and with it the conditions under which architectural form becomes evidence, can be justified, and can be judged as legitimate. To name these phases, the paper uses the term regimes, not as a stylistic label, but as a comparative framework describing relatively stable constellations of media, artefacts, agency, and disciplinary legitimation. The term becomes particularly operative in continuity with Michael Young’s account of image regimes, where architecture aligns itself with dominant visual technologies and economies more than with classical categories of object, type, or construction [3].

The central thesis is that, within a prolonged digital condition, the geometric generator has migrated in both substrate and intelligibility: from digital model to code, and from code to a latent regime. The latter introduces an inversion that until recently was taken for granted. In digital architecture, the hegemonic workflow has been model-to-image: the architect produces explicit geometry in a readable environment and then generates visualizations. By contrast, in many emergent uses of GenAI, the process begins with a synthetic image—a visually persuasive prediction produced from a prompt—and only then does the disciplinary challenge arise: translating that image into coherent geometry. Here the architect operates as a kind of digital archeologist, inferring what model could plausibly produce what the image suggests [4]. This inversion is diagrammed in Figure 1.

This shift relocates the geometric generator to a new and strange place: the model’s latent space, an opaque statistical domain trained on millions of images whose internal logic is not directly accessible. The difference is qualitative. A project developed through modelling or parametric code begins from a legible artefact—a mesh, a surface, a parameter space, a script—something that can be read, audited, and debated as a human intellectual construction. Even when the system is complex, control remains syntactic and explicit: the designer manipulates NURBS (Non-Uniform Rational B-Splines), a standard mathematical representation for smooth curves and surfaces in CAD modelling, or writes relations and understands which operations generate which effects [2,5–8]. In the latent regime, by contrast, control becomes semantic and indirect, exercised through natural language and iterative negotiation with a black box of billions of parameters [1]. This mutation redefines

https://doi.org/10.3390/architecture6010025

authorship: the architect is no longer only a direct producer of geometry but becomes a critical editor, curator, and stabilizer of variants, selecting, filtering, and translating [4].

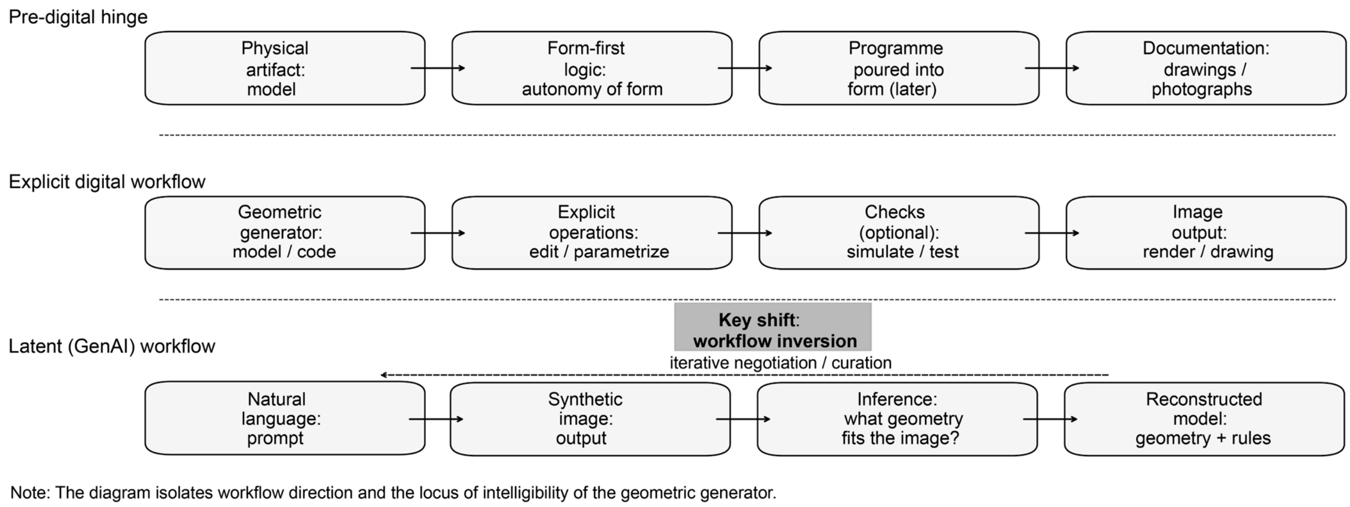

Figure 1. Workflow inversion across digital regimes: in model- and code-based workflows, geometry precedes the image (model-to-image); in the latent regime, a synthetic image is produced first and geometry must be reconstructed post hoc (image-to-model requirement).

At the same time, the intuition of a "form-first" design, in which an abstract form precedes its functional rationalization, is not new. For this reason, the article introduces Frederick Kiesler's Endless House as a pre-digital antecedent and disciplinary hinge: a mid-twentieth-century investigation that begins from a physical model, a continuous sculptural blob conceived as an autonomous object before domestic programme is retrofitted within it [9–11]. GenAI may replace the physical object with a synthetic image, but the avant-garde principle of formal autonomy and subsequent architectural translation was already firmly established. The novelty does not lie in form preceding function, but in the change in the generator's nature: from a tangible artefact or readable construct to an opaque statistical space.

To sustain this argument, the paper defines three regimes of digital architecture—the digital-model regime, the parametric code regime, and the latent regime—while using the pre-digital case as a disciplinary hinge to clarify contrast. The transition is traced through a consistent formal family, the blob, precisely because it allows regimes to be compared without changing morphological intuition, isolating what matters: not style, but the status of the geometric generator. The reading draws on Carpo's historiography of the digital turn [2], Young's account of image regimes [3], and contemporary debates on AI-based design ecologies [4]. In parallel, the paper acknowledges continuity with parametricism as a horizon of variation, differentiation, and relational control, while arguing that the latent regime displaces agency beyond code syntax toward negotiated semantics [1,2,8]. This comparison also foregrounds a critical tension that becomes acute in the latent regime, the plausibility gap between high visual-credibility and low tectonic-legibility, which the paper develops as an evaluative criterion in Sections 4 and 5.

Methodologically, the work combines conceptual and historical analysis with design research. Section 2 establishes the theoretical framework, Section 3 develops the regimes and their canonical cases, and Section 4 operationalises the latent regime through an authored case study describing the image-to-model workflow. Section 5 discusses implications for authorship, pedagogy, and disciplinary judgement, and Section 6 synthesizes conclusions and directions for future research.

Figure 2 summarizes this displacement of the geometric generator across regimes.

Figure 2 — Regimes of the Geometric Generator (Conceptual Framework)

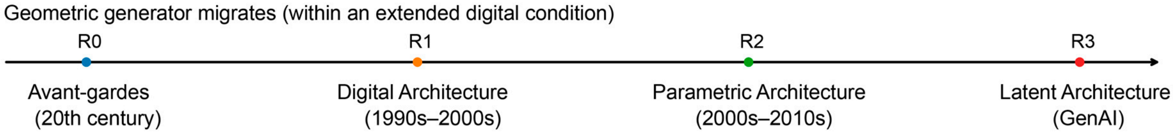

	Regime 0 Avant-gardes	Regime 1 Digital Architecture	Regime 2 Parametric Architecture	Regime 3 Latent Architecture
Generator (what drives form)	Physical artifact / model (form precedes programme)	Explicit geometry (mesh / NURBS model)	Rule-set / script (parameters + constraints)	Statistical latent model (prompted synthesis)
Control (how agency operates)	Direct shaping (material handling)	Explicit, geometric (editable operations)	Syntactic control (relations in code)	Semantic, indirect (iterative negotiation)
Workflow direction	Object → drawing/image (documentation)	Model → image (rendering pipeline)	Code → model → image (generative assembly)	Image → model (inference / reconstruction)
Legibility / auditability	High (tactile) interpretable artifact	High (auditable) readable geometry	High (auditable) readable rules	Low (opaque) black-box statistics
Architect's role	Maker / form inventor	Modeler / controller	Coder / systems designer	Editor–curator–translator (stabilises variants)

Note: Regime boundaries are heuristic; the diagram isolates shifts in the geometric generator and associated agency.

Figure 2. Framework of the geometric generator across regimes: design authority shifts from object-based composition (Regime 0) to digital modelling (Regime 1), parametric rule systems (Regime 2), and data/latent-space curation (Regime 3).

2. Theoretical Framework and State of the Art

2.1. Architecture and Image Regimes: Three Visual Revolutions

If, as Michael Young argues, contemporary architecture is increasingly "modelled after images", then any genealogy of digital design must begin from the history of image regimes rather than from the history of construction [3]. In this paper, image regimes are understood not only as representational shifts but as evidentiary regimes: they redefine what counts as valid early evidence in design, and therefore where the geometric generator is located. From this perspective, today's latent regime of AI-generated images is not an isolated novelty but the third episode in a longer sequence of visual revolutions, each triggered by a technological innovation that redefines how reality can be pictured and, in turn, how architecture can be conceived.

The first visual revolution is the invention of perspectiva artificialis in early quattrocento Italy. Brunelleschi's perspectival experiments and Alberti's subsequent codification in De pictura and De re-aedificatoria do not merely introduce a new drawing technique; they institute a conceptual model in which space becomes a homogeneous, mathematically ordered field that can be projected onto a plane [12,13]. Architecture, painting, and urban perception are reorganized around geometrical optics: the building is both designed and judged according to how it will appear within the cone of vision of a single, idealized observer. Representation remains mimetic, but the standard of realism is now set by a constructed visual device.

The second visual revolution, in the mid-nineteenth century, is the invention of photography. Mechanical capture of light destabilizes the older pact between painting and reality: if the camera can now "tell the truth" more quickly and more faithfully, painting is relieved from representational duty. The crisis is productive: it opens modernist explorations of abstraction, fragmentation, and non-Euclidean space in Cézanne, Picasso, Malevich, or

Duchamp [14,15]. In parallel, architecture becomes increasingly conditioned by photographic framing and by the circulation of images as a cultural apparatus [3,14,16,17].

Frederick Kiesler belongs squarely to this post-photographic modernity. His writings on "art-architecture", the theatre-city, and dynamic continuous space are informed by the pictorial and cinematic experiments of the historical avant-gardes [9–11]. The dialogue with Marcel Duchamp is decisive: Kiesler reads The Large Glass as an apparatus where spatial perception, technical devices, and the active role of the spectator are fused into a single environment [15]. Duchamp's readymades and his insistence on spectator participation erode stable categories of object and authorship, and Kiesler translates these intuitions into architecture by theorizing environments in which movement and inhabitation complete the work [9,10]. The Endless House can thus be read as an architectural heir to Duchamp: a continuous, almost blob-like object that dissolves distinctions between figure and ground, between art and architecture, and anticipates an immersive environment in which the inhabitant becomes, in a sense, a performer [9,10,15]. Crucially for this article's argument, Kiesler also provides a hinge case: an early instance in which the first and most persuasive "evidence" of a project is an autonomous artefact whose architectural rationalization comes later. This hinge will be used as an echo to measure the later migrations of the geometric generator from object to model, from model to code, and from code to latent prediction [3].

The third visual revolution is the contemporary regime of synthetic images, intensified by internet culture and, more recently, by generative AI. Since the late twentieth century, the proliferation of screens and networked platforms has produced an unprecedented volume and velocity of images. At first, digital photography and rendering reinforced mimetic logics, producing ever more realistic pictures of existing or projected spaces. But with diffusion models and large-scale datasets, images no longer need to refer to pre-existing realities: they are sampled, interpolated, and hallucinated within high-dimensional latent spaces learned from training data [1,4]. Unlike perspective or photography, which changed how reality is seen, this third revolution foregrounds images of things that do not yet, and may never, exist. For architecture, this means that the "model" of reference is increasingly an image that precedes built form and even geometrically explicit models.

Seen from this angle, Young's claim that architecture adjusts itself to regimes of images more than to regimes of objects or construction becomes almost literal [3]. Endless House thus functions as a historical hinge: as a product of the second visual revolution, it already treats the architectural object as a quasi-sculptural continuous artefact, relatively detached from a functional programme, and conceived to be exhibited, photographed, and inhabited also as an image [9–11]. This hinge role matters because it makes visible a recurring disciplinary tension: the autonomy of an early artefact versus the later work of architectural rationalization. In the latent regime, that tension reappears in a new form, because the first artefact can be a highly persuasive synthetic image that is not the representation of a prior model but a visual prediction, opening a structural gap between visual plausibility and tectonic legibility [1,4,18].

Taken together, these three visual revolutions show a recurrent pattern: each technological innovation that reorganizes the status of the image triggers a disciplinary crisis and, with it, a field of experimentation. Perspective was not only a technique but a new epistemic framework of space; photography was not only a medium but a rupture that forced modernity to reinvent regimes of visibility. Analogously, the current proliferation of synthetic images should not be reduced to an instrumental improvement of rendering or a productivity shortcut. Historically, visual crises do not merely add tools: they transform criteria of plausibility, authorship, and judgement, opening the ground for new avant-gardes. This is why generative AI must be read at the scale of a regime change, with consequences for the ontology of the architectural project—specifically for what counts as first design evidence and where the geometric generator is situated [1–3].

2.2. From Geometric Complexity to Code: The Digital Condition and the Displacement of the Geometric Generator

When digital tools entered architectural practice in the late twentieth century, they did so within an image-saturated milieu yet initially served a different purpose: to manage geometric, tectonic, and organizational complexity. Early CAD systems of the 1960s and 1970s primarily assisted drafting and documentation but already gesture toward a parametric understanding of form as sets of controllable variables [6]. From the standpoint of this paper, what matters is the evidentiary status of the generator: unlike Kiesler's pre-digital hinge, where the first persuasive "evidence" can be a singular physical artefact, the early digital condition stabilizes evidence through legible geometric constructs. Historical precedents of parametric thinking can be traced further back in Gaudí's hanging-chain models and catenary vaults or Luigi Moretti's mathematical explorations, where form emerges from manipulating continuous variables under structural and programmatic constraints [6]. In the 1990s, the availability of NURBS modellers and animation software catalyzes a first generation of digital projects that treat continuous surfaces and topological transformations as primary design problems. Greg Lynn's work, articulated in Animate Form, is paradigmatic: a shift from Cartesian X–Y–Z coordinates to U–V surface parameters, the use of animation as a formal processor, and an emphasis on growth and transformation over static form [5]. In parallel, Jesse Reiser and Nanako Umemoto develop what they call "novel tectonics": a morpho-dynamic model in which flows, fields, and external stimuli generate complex continuous structures whose articulation is inseparable from material and structural performance [7]. Together, these strands define a first digital condition in which the 3D model—deformable, animatable, multi-variant—becomes both the central design object and the primary generator of architectural geometry [2,5,7]. In contrast to Kiesler's sculptural artefact, the generator here is a reproducible digital genotype: a legible model that can be measured, edited, and iterated while preserving an explicit chain from geometric decision to formal outcome.

A second, closely related tendency pushes these logics toward explicit parametrisation. By the late 1990s and early 2000s, coding environments such as MEL (Maya Embedded Language, the scripting language of Autodesk Maya) and Python and, later, visual programming tools such as Grasshopper enable architects to encode relations, constraints, and data streams directly into scripts. Geometry becomes a dependent variable. Patrik Schumacher formalizes this constellation under the label parametricism, presenting it as a comprehensive agenda premised on continuous variation, correlation of multiple parameters, and mass customisation [8]. Zaha Hadid Architects popularize its aesthetic through large-scale projects whose fluid envelopes are produced and rationalized with advanced digital tools, while Frank Gehry's Guggenheim Bilbao—realized with CATIA—stands as an early technical and symbolic milestone in the construction of spline-based "free-form" geometry [19]. Within this landscape, practices such as Marc Fornes/THEVERYMANY explore finely discretized shells and lattice structures in which code, material behaviour, and fabrication co-produce highly differentiated envelopes, and where authorship is claimed primarily in the design of the protocol rather than in any single instance [20]. If Lynn marks a migration from object-like form to explicit digital geometry, Fornes exemplifies a further migration from explicit geometry to explicit relations: the generator is no longer the model-as-object but the model-as-rule-system.

For the purposes of this paper, the key point is not stylistic but ontological: in both the animate-form phase and the parametric phase, the geometric generator remains either the digital model itself or the underlying code and parameter space. Programmatic, structural, or environmental data may modulate the output, but the design process remains fundamentally model-first: geometry is authored as a manipulable entity or as a computed result of scripted relations, and only then translated into images and, eventually, construction [2,5,7,8,20]. Crucially, even when authorship is partially delegated to computation, the generator remains legible: it can be inspected, audited, and debated as a human-readable artefact (a model, a script, a parameter space). This is why Regimes 1 and 2 are treated here as consecutive and internally coherent displacements within the same evidentiary logic: despite their differences in agency, both preserve a legible, auditable generator (model or code) and an explicit chain of justification. This is precisely the condition that makes Regime 2 function, in the paper's wider argument, as a threshold rather than a rupture: it expands and accelerates variation, but it does not yet displace architectural agency into an opaque domain.

This emphasis on legibility sets up the contrast with the latent regime. In recent scholarship, the shift introduced by generative AI is often described as a move from syntactic control (geometry/code) toward semantic or connectionist design rationality, where form is accessed through statistical inference and language-based negotiation rather than explicit relations [21,22]. This condition prepares the contrast with the latent regime, where the generator ceases to be a legible design artefact (model or code) and migrates into an opaque statistical domain primarily accessed through semantic negotiation.

2.3. Generative AI, Automation, and Predictive Architectures

Contemporary debate on generative AI in architecture moves between two readings in tension. On one side, a continuity reading argues that generative models do not inaugurate an independent "new turn", but accelerate, automate, and scale logics already embedded in digital culture, particularly variation, mass customisation, and the production of formal families. In this line, Carpo insists that much of the "miracle" of generative AI lies in the automation of tacit knowledge and imitative competence, statistically learned from corpora, rather than in disciplinary invention ex nihilo [1]. The result is a machine capable of producing plausible proposals that, in many respects, extend the digital condition rather than replace it.

On the other side, a rupture reading holds that, even if AI continues previous processes, it alters a structural point of the project: the location of design agency. In earlier regimes, even when projects were highly complex, the geometric generator remained legible as a human artefact—either a manipulable model or an auditable codebase. By contrast, the latent regime relocates the generator into an opaque statistical domain whose internal operations cannot be reconstructed as explicit architectural relations, rules, or causal chains. This opacity is not a minor technical detail but a change in the epistemic status of the instrument. Accordingly, the ontological claim advanced in this paper is not that "architecture becomes immaterial", but that the project's first evidence, and thus the conditions of judgement and legitimacy, are reconfigured. What counts as valid early proof of a design, and what can be audited, explained, and justified, shifts with the generator's substrate. For that reason, its impact should be evaluated less as a "new tool" and more as a reconfiguration of decision-making and formal justification [2,18]. Two elements that could easily be lost, but are central to this section, are (i) the intentional asymmetry between regimes and (ii) cognitive friction in the latent regime (semantic agency, the curator's paradox, and the "interrogation" of a black box). Both are made explicit here, because this is where the argument must show that the shift is not merely technical but ontological and operational.

This paper frames that reconfiguration as a relocation of the geometric generator and, therefore, as an ontological shift internal to architectural practice. In brief, the generator migrates from explicit geometry to explicit relations, and finally, to latent prediction. In the model regime, the generator is the digital model, an explicit geometric artefact that can be inspected, measured, and revised. In the code regime, the generator becomes a parameter space and a script, a readable set of relations in which geometry is a dependent variable. In the latent regime, the generator is no longer a legible artefact, but a learned topology that is accessed primarily through semantic negotiation, prompts, and iterative selection. Figure 3 summarizes this shift and the resulting change in the project's evidentiary chain.

From legible, auditable artefacts (model/code) to opaque latent prediction, where early "evidence" of design becomes a synthetic image.

	Regime 1 **Digital model as generator**	**Regime 2** **Parametric code as generator**	**Regime 3** **Latent space as generator**
Generator substrate What the "seed" is	Legible geometry (mesh / NURBS) Explicit, inspectable model Human-readable artefact "Source of truth" for renders	Legible relations (script / param space) Rules, constraints, variables migration of the generator Auditable generative grammar Instances as computed outputs	Opaque statistics (latent topology) Learned manifold from datasets Black-box internal causality Generator not directly "readable"
Access mode How agency operates	Syntactic control Edit geometry (vertices, surfaces) Deterministic operations	Syntactic control Write / debug code Parameter steering	Semantic negotiation Prompts, curation, iteration Selection + stabilisation protocols Indirect control over outcomes
Evidentiary status What counts as early evidence	Model-first evidence Geometry → image → documentation Auditability: high Justification via explicit form-making	Rule-first evidence Code/relations → instances → image Auditability: high Justification via readable generative logic	Image-first evidence Synthetic "predictive appearance" Plausibility gap: visual credibility vs tectonic legibility Post-hoc reconstruction Infer geometry + tectonics afterwards

Ontological shift: early evidence can be an image produced without prior geometry

Reading key: "legible" = inspectable by disciplinary reasoning; "opaque" = not directly reducible to explicit rules.

Figure 3. Ontological shift in the geometric generator across regimes. The generator migrates from legible and auditable artefacts, first explicit geometry (digital model) and then readable relations (parametric code), to an opaque latent topology accessed through semantic negotiation. In the latent regime, the project's earliest evidence can be a synthetic image produced without prior geometry, shifting verification to post hoc reconstruction and opening a plausibility gap between visual credibility and tectonic legibility.

However, it is important to clarify that the loss of direct control over the final object is not entirely new. The displacement of authorship was already underway in parametricism, where the architect no longer designs a single final object but designs a space of possibility, which is the set of geometries a system can generate under rules and constraints. Schumacher formalizes this as a regime of correlations and continuous variation [8]. Practices such as Fornes' make this shift explicit when they locate authorship primarily in the code and the process rather than in each particular instance, accepting that part of agency is delegated to algorithmic logic [20]. From this perspective, AI may be read as a radical intensification of that delegation: the possibility space is no longer a readable parameter space or an authored grammar, but a learned latent topology that is, to a large extent, inaccessible.

This displacement becomes especially visible when AI enters early ideation. Design ceases to begin with explicit geometry or a parametric system and begins instead with a synthetic image. In the experiments documented within "neural architecture" design ecologies, AI operates as a motor of visual variation that produces series of alternatives with high atmospheric- and material density, forcing the architect to operate through selection, editing, and stabilization rather than direct modelling [4]. Here the paper's central hypothesis becomes operative: the inversion of the workflow from model → image to image → model is not a trivial reordering of steps, but a relocation of the project's source into an apparatus of visual synthesis. For the first time in digital architecture, the primary artefact can be a photorealistic image of something with no prior geometric substrate: not a representation of a model, but a predictive appearance. In this paper, "predictive" does not refer to textual next-word prediction, but to the production of a visually coherent projection that precedes, and pressures, geometric and tectonic rationalization. This produces an unprecedented gap between visual plausibility and constructive viability, the *plausibility gap*, understood as the structural distance between what appears convincingly buildable and what can be made tectonically legible, auditable, and contractible. This distance is not an anecdotal mismatch. It becomes a disciplinary problem, because it targets exactly the site where architecture claims responsibility, namely the passage from appearance to accountable form. The GenAI blob in Section 4 takes as a research problem by forcing appearance to be translated back into geometry and, ultimately, into system.

In this context, tectonics and material culture re-emerge not as a moralistic resistance to the digital, but as a field of critical verification. AI tends to produce highly persuasive images that can conceal structural, material, or logistical incoherence, shifting the conflict from geometry to mediation: who controls the means, which knowledge becomes automated, which labour is invisibilized, and which criteria validate the passage from image to building. Carpo frames this as a technological "culture war" around building technology, where narratives of innovation polarize between celebrating formal liberation and calling for disciplinary returns to materiality, without accounting for how project devices actually change [18]. Within this paper's framework, the disciplinary question sharpens: it is not only what AI can produce, but what protocols of translation, auditability, and responsibility are needed to stabilize a predictive image as architecture. In practical terms, the plausibility gap makes judgement inseparable from verification procedures, and turns the definition of criteria, not only the production of images, into a central design act.

If the focus is widened, AI also reopens a historical problem: the relationship between canon, style, and repetition. Since training data defines a possibility space, the dataset functions as an operative canon that conditions what becomes recognizable, plausible, and "imaginable" by the system. Innovation therefore tends to appear as a plausible recombination inside the learned manifold rather than as a programmatic rupture [1,2]. This does not require condemnation, but it shifts the debate from romantic originality toward the critical control of archives, the selection of precedents, and the construction of criteria to evaluate variation. In other words, the latent regime turns curation—of both data and images—into an ontological operation of the project, because it conditions which architecture can be produced and at what disciplinary cost.

For these reasons, this article is less interested in AI's capacity to generate images than in its capacity to alter the status of the geometric generator, and thereby the conditions of authorship, judgement, and responsibility. If, in the code regime, control was syntactic and explicit—writing rules, parameterising relations, debugging scripts—in the latent regime control becomes semantic and indirect: negotiating descriptors, modulating prompts, curating iterations, and then reconstructing geometry from the returned appearance. This is the specific level at which the ontological shift is located in this paper: the generator changes, therefore the evidentiary order of the project changes, and with it the means through which architectural knowledge is produced, debated, and validated [23,24]. This allows the guiding question of the paper to be formulated precisely: is AI an inflexion point that inverts the design process, or the culmination and extreme intensification of a prior digital trajectory? The thesis advanced here is that both can be true at once: historical continuity in the logic of variation, and ontological rupture in the location and legibility of the generator. This ambivalence—continuity and break—is what the GenAI blob case study will test operatively in Section 4 by turning the image-to-model inversion into an explicit and assessable procedure.

3. Four Generative Regimes: From Object to Latent Space

This section translates the theoretical framework (Section 2) into a comparative model of four generative regimes, understood as historically recognizable configurations in which the geometric generator of a project changes, and with it the locus of architectural agency. In continuity with the thesis stated in the Introduction, the last three regimes belong to a single prolonged digital condition, from the 1990s to the present, while the first operates as an analogue precursor and disciplinary hinge. The "zero" case, Kiesler, is not digital, but it anticipates a form-first logic that reappears in 1990s digital architecture and is re-intensified today through synthetic images.

Importantly, the regimes are not treated here as four "equal" theoretical plateaus. Regime 0 and Regime 3 carry greater ontological weight by design, because they expose the two endpoints of the paper's argument about what counts as first evidence in architectural projects. Regimes 1 and 2 are therefore read as consecutive, internally coherent displacements within the same evidentiary logic: despite their differences in agency, both preserve a legible and auditable generator (model or code) and an explicit chain of justification. This intentional asymmetry is not a weakness of the framework but part of its claim about continuity versus rupture.

Each regime is illustrated through a blob-like case, not for stylistic reasons, but because keeping a morphological intuition constant makes the ontological displacement of the generator legible across regimes. In Section 3, the emphasis is conceptual and genealogical: what each regime is, why it exists, and why the selected examples represent it. The operational comparison, including step-by-step pipelines and the authored GenAI experiment, is developed in Section 4.

To enable a like-for-like comparison, Figure 4 assembles four diagrammatic reconstructions (Figure 4a–d), each corresponding to one regime and used as a probe to identify what, in each case, actually generates form.

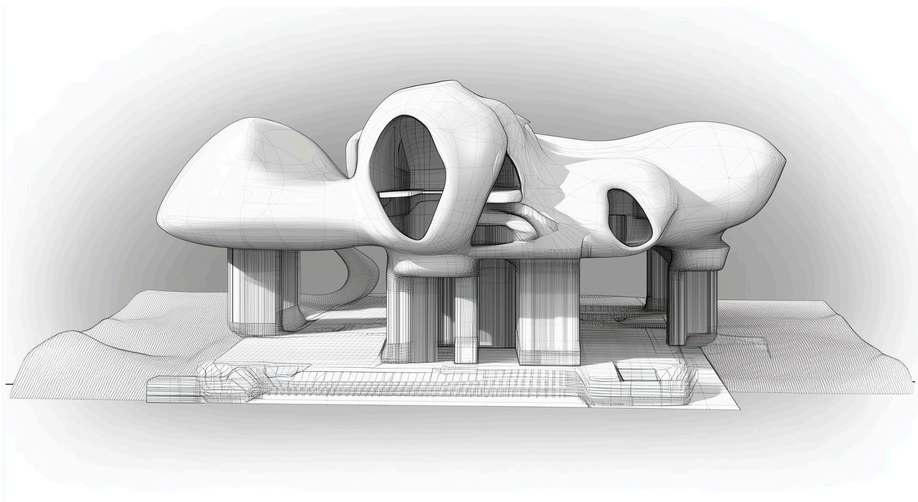

(**a**)

Figure 4. *Cont.*

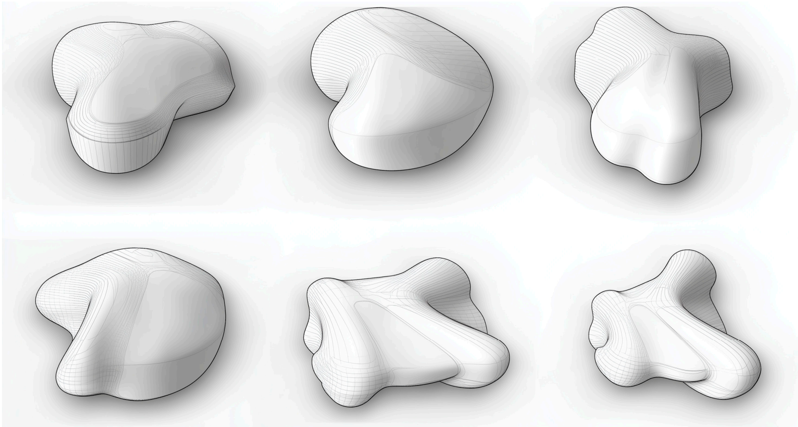

(**b**)

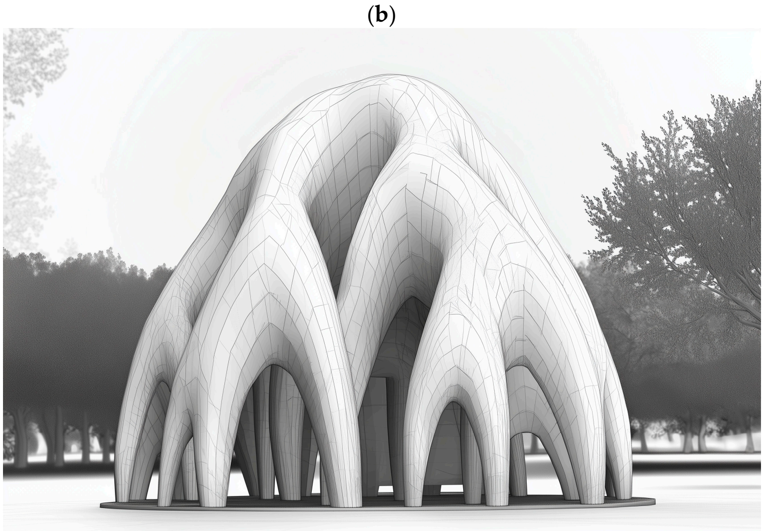

(**c**)

(**d**)

Figure 4. Comparative diagrammatic reconstructions used as probes across four generative regimes. Panels (**a**–**d**) correspond to Regimes 0–3 and are rendered as white models with thin linework to foreground the generator's status as evidence, rather than stylistic difference. (**a**) **Regime 0—Object as generator:** the architectural object operates as the primary locus of formal decision-making (Kiesler's Endless House). (**b**) **Regime 1—Digital model as generator:** a continuous digital geometry becomes the main site where form is produced and iterated (Lynn's Embryological House). (**c**) **Regime 2—Parametric rule system as generator:** scripted relations and adjustable parameters govern the assembly of form (Vaulted Willow). (**d**) **Regime 3—Latent regime as generator:** image-first synthesis drives form, while geometry must be reconstructed post hoc and validated tectonically.

3.1. Regime 0: Object as Generator, Frederick Kiesler's Endless House (1947–1960)

In the object regime, the geometric generator is neither a functional diagram, nor a set of rules, nor an editable digital model, but an autonomous physical artefact whose form precedes its architectural rationalization. Endless House condenses Kiesler's long research on continuous environments and "correlated" space, articulated around notions such as correalism, elasticity, and "art-architecture", where space is conceived less as a composition of rooms and more as a continuous medium correlating bodies, perception, and artefacts [9,10]. Regime 0 is summarized in Figure 5, where the object is framed as an artefact that concentrates formal decisions and disciplinary meaning (see also Figure 4a).

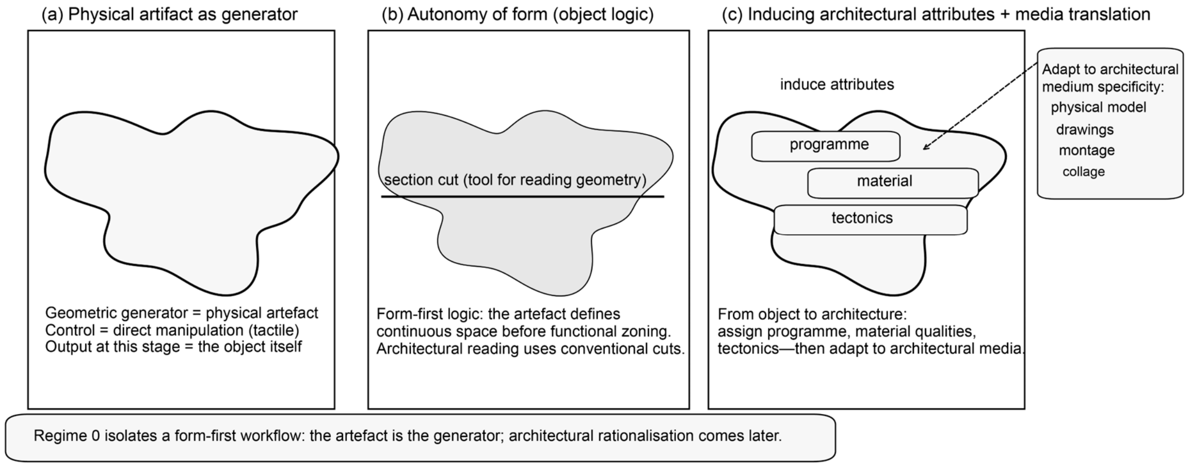

Figure 5. Regime 0—**Object/Artefact**. The architectural object acts as the geometric generator: a discrete artefact that concentrates formal decision-making and disciplinary meaning, and that can be read independently of downstream representational pipelines. (**a**) Physical artefact as generator (direct, tactile manipulation; the object is the primary output at this stage). (**b**) Autonomy of form (object logic): the artefact defines continuous space prior to functional zoning; architectural reading begins through conventional analytical operations (e.g., section cuts). (**c**) Translation from object to architecture: architectural attributes (programme, material, tectonics) are induced and the artefact is subsequently adapted to medium-specific representations (physical model, drawings, montage/collage), postponing rationalisation to later stages.

What matters for the present genealogy is that the model ceases to be a provisional representation and becomes the primary site of invention. The project was iterated for years through models and exhibition versions, culminating in the well-known museum-mediated condition that stabilized it simultaneously as a total object and as an image-making device [11].

Three features justify reading Endless House as an object regime. First, form precedes programme. The project does not originate in a brief or typology; a continuous spatial shell is given first, and domestic programme is retrofitted later. This inversion is not an eccentricity but aligns with a post-photographic modernity that experiments with the status of objects and spectators [9–11]. Second, the model functions as the work, rather than as a draft. Agency shifts from deriving form from requirements to reading an object and translating it into architecture afterwards, through structural, programmatic, and representational decisions. Third, the project crystallizes a performative conception of space, where continuity is inseparable from movement and perception; Kiesler's affinity with Duchamp's destabilization of authorship and spectatorship provides a relevant intellectual background for this spatial logic [15].

This is why Endless House operates here as a disciplinary anchor: not as a pre-digital anomaly, but as a hinge that connects contemporary questions of geometric generation to a tradition of object-oriented design capable of preceding programme and demanding architectural translation afterwards [9–11].

3.2. Regime 1: Digital Model as Generator, Greg Lynn's Embryological House (1997–2001)

The first properly digital regime displaces the generator from the physical object to the digital model as a manipulable, productive entity. The key shift is not "drawing with a computer" but treating computation as a medium for generating and varying continuous geometries. In Lynn's work, this condition is theorized as an animate form, where form is understood as process, variation, and transformation, and concepts drawn from animation and surface mathematics become operators of design [5]. Regime 1 is summarized in Figure 6, which frames the digital model as the generator and distinguishes genotype (rule/structure) from phenotype (resulting form) (see also Figure 4b).

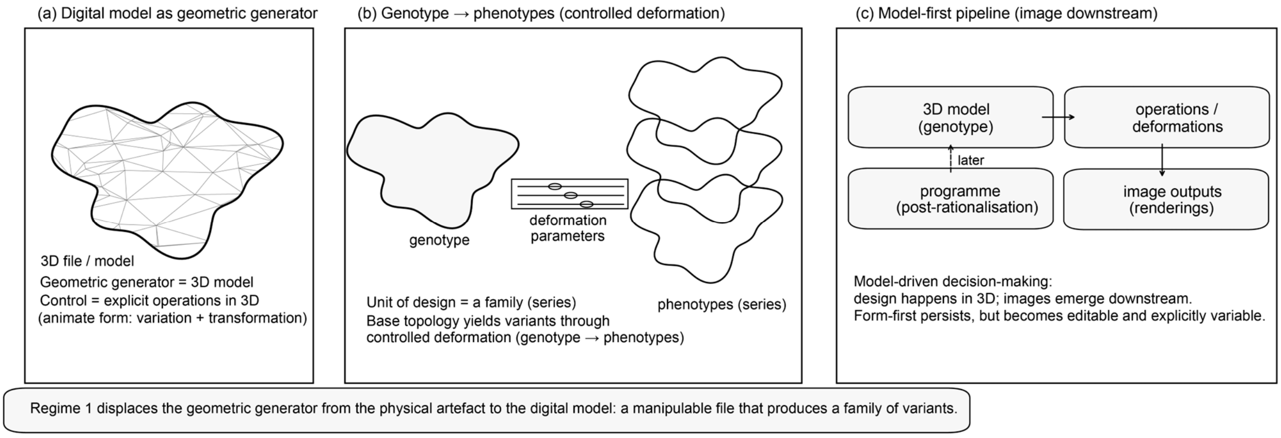

Figure 6. Regime 1—Digital Model (Genotype/Phenotype). The digital model becomes the geometric generator: a manipulable mesh and its underlying organizational logic (genotype) give rise to a family of formal outcomes (phenotype) through iterative modelling operations. (**a**) Digital model as geometric generator; grey triangulation lines indicate the underlying mesh/wireframe structure of the 3D model. (**b**) Genotype → phenotypes: deformation parameters are represented as sliders; circles mark adjustable control values driving controlled variation across the series. (**c**) Model-first pipeline: arrows indicate the direction of the workflow (3D model → operations/deformations → image outputs), while the dashed arrow ("later") indicates programme post-rationalisation occurring after model generation. Graphic key: grey lines indicate the underlying mesh/wireframe structure of the digital model; circles denote parametric control values (adjustable deformation parameters) that propagate changes across the model; arrows indicate the direction of the workflow (model → deformation/variation → image output).

Embryological House represents this regime because the object of design is no longer a single building, but a family. The "embryological" metaphor is not merely rhetorical; it describes a generative structure in which a geometric genotype yields multiple phenotypes through controlled deformation of a continuous model. Here, the geometric generator is the 3D file and its internal parameters of deformation, rather than plan, physical prototype, or textual specification [2–5].

Three features make this paradigm clear. First, genotype and phenotypes: the unit of design becomes a base topology capable of producing series, aligning with the digital logic of variation [2]. Second, model-driven decision-making: design decisions occur directly in three dimensions, while images emerge downstream as outputs of the model. Third, programme as post-rationalization: as in Kiesler, a form-first tendency persists, but the "hinge" changes its substrate, the autonomous artefact is no longer a singular physical object but a reproducible digital genotype that can be iterated as a family [5]. The regime remains model-first: explicit geometry precedes image, even if the image becomes culturally central.

3.3. Regime 2: Parametric Code as Generator, Marc Fornes/THEVERYMANY's Vaulted Willow (2014)

The parametric regime introduces a further displacement: the geometric generator is no longer the visible model but code and parameter space, a system that outputs geometry as a computed result. Schumacher formalizes this ambition under the term parametricism, emphasizing multi-variable correlation and systemic variation [8]. In parallel, the broader discourse on digitally mediated tectonics insists that these geometries are not only stylistic but reorganize the relation between form, performance, and fabrication [7]. Regime 2 is summarized in Figure 7, where parametric code and explicit rule-sets become the generator and mediate between form and fabrication (see also Figure 4c).

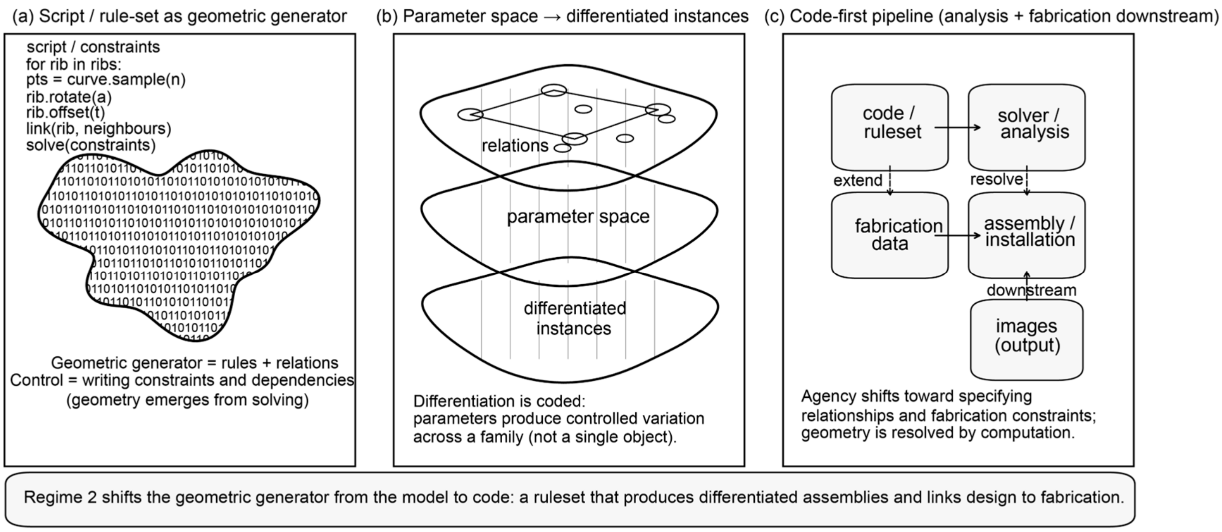

Figure 7. Regime 2—Parametric Code/Rules/Fabrication. The geometric generator shifts from a single model to a rule-based system: scripted relations and parameters define an expandable design space, while constraints and fabrication logics stabilize outcomes into buildable configurations. (**a**) Code/ruleset as generator: the parametric definition (dependency graph) encodes relations and constraints that drive geometry. (**b**) Solver/analysis stage: numeric evaluation and constraint-solving stabilise outputs by selecting viable instances from the design space. (**c**) Fabrication and assembly output: discretisation/panelisation logic is translated into part identifiers and constructible configurations for production and coordination. Graphic key: arrows indicate direction of dependency/workflow; circles denote parametric nodes/variables (control values) whose changes propagate through the system.

Within this landscape, Vaulted Willow is particularly useful because it makes explicit a key thesis of the regime: authorship resides in the system, not in a single final instance. The project foregrounds computational protocols that organize geometry through rule-based discretization, fabrication logic, and assembly constraints [20].

Three features justify its reading as a code regime. First, the script as primary object: the project is organized through computational rules that integrate skin, structure, and pattern as a single system. Second, geometry as derived output: form is not sculpted directly but computed from constraints and relations; changes propagate across the system [8]. Third, tectonic pre-rationalization: discretization, tolerances, and assembly are embedded in the generator rather than added later [7,20].

Crucially, this regime already entails a structural reduction in direct control over the final object, not as a loss of agency, but as a shift in agency from designing an object to designing a possibility space and then selecting viable instances. If Kiesler's object was a singular sculptural "given" that demanded retroactive architectural translation, here the generator becomes an explicit and auditable grammar, the form is rationalized before it exists, and translation is pre-embedded as constraints.

3.4. Regime 3: Latent Space as Generator, Toward Predictive Image-Based Generators

The latent regime pushes that delegation further, not merely by increasing complexity, but by changing the status of the generator itself. In contemporary generative systems, the geometric generator is no longer a legible artefact, model or human-authored code, but a high-dimensional statistical manifold learned from massive corpora of images and text [1,18]. Formal production is accessed through semantic descriptors and iterative sampling rather than through explicit syntactic relations. Regime 3 is summarized in Figure 8, which breaks the latent workflow into semantic prompting, diffusion-based sampling with iterative curation, and the translation of selected images back into explicit geometry (see also Figure 4d).

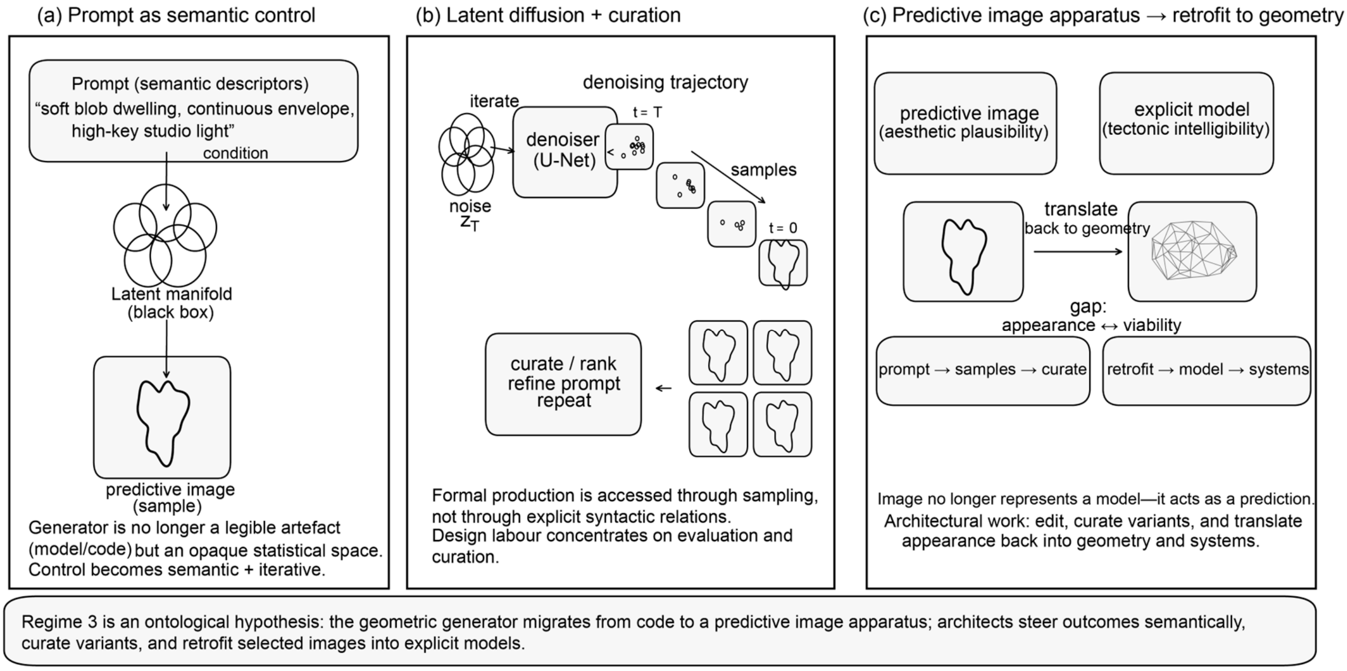

Figure 8. Regime 3—Latent space: latent diffusion, sampling, and semantic curation. (**a**) Prompts act as semantic control over an opaque latent manifold. (**b**) Form is accessed through diffusion-based sampling and iterative curation rather than explicit syntactic relations. (**c**) The predictive image becomes the primary design artefact and is retrofitted into geometry, exposing a gap between aesthetic plausibility and tectonic intelligibility. Graphic key: arrows indicate the direction of the pipeline (prompt → sampling/curation → image → geometry/validation); numbers mark the main operational stages; circles denote decision points/curatorial interventions (selection, rejection, refinement).

Section 3 defines this regime conceptually, without yet unfolding the authored case study, which is developed in Section 4. Three conditions delineate the regime. First, the initial artefact tends to be a synthetic image of high visual plausibility, not explicit geometry. This shifts the point of departure and opens a new gap between aesthetic plausibility and tectonic intelligibility, because the image no longer represents a model, it operates as a visual prediction [4,18]. Second, the generator is a black box whose parameters do not translate easily into architectural rules. Parametricism already shifted authorship toward systems, but latent space intensifies delegation because control becomes semantic and indirect [1,2].

Third, the regime reorganizes the architect's role around editing, curating variations, and translating appearance back into geometry and systems, a convergence of Carpo's diagnosis of imitation automation, the design ecology of neural-architecture workflows, and tectonic concerns about persuasive images and material responsibility [1,4,18].

At this point, the "Kiesler hinge" closes: the latent regime appears radically new because it inverts the workflow (image to model), yet it also reactivates an older form-first logic in which an autonomous artefact precedes rationalization. The difference is that the artefact is now a statistically produced prediction rather than a stubborn object, and the ensuing translation is performed under conditions of opacity and overabundance. This is where cognitive friction becomes structural: semantic agency must be exercised as a slow judgement against a fast generator, the curator's paradox, and the repeated "interrogation" of a black box whose internal reasons cannot be fully audited.

In sum, Section 3 establishes the latent regime as an ontological hypothesis: the generator migrates from code to a predictive image apparatus. Section 4 tests and operationalizes this definition through the GenAI blob case study, making explicit how one moves from image to model and evaluates the resulting gap between appearance and viability. Having defined the four regimes, Figure 9 treats the "blob" as an invariant morphology in order to isolate how each regime defines and legitimizes form-making.

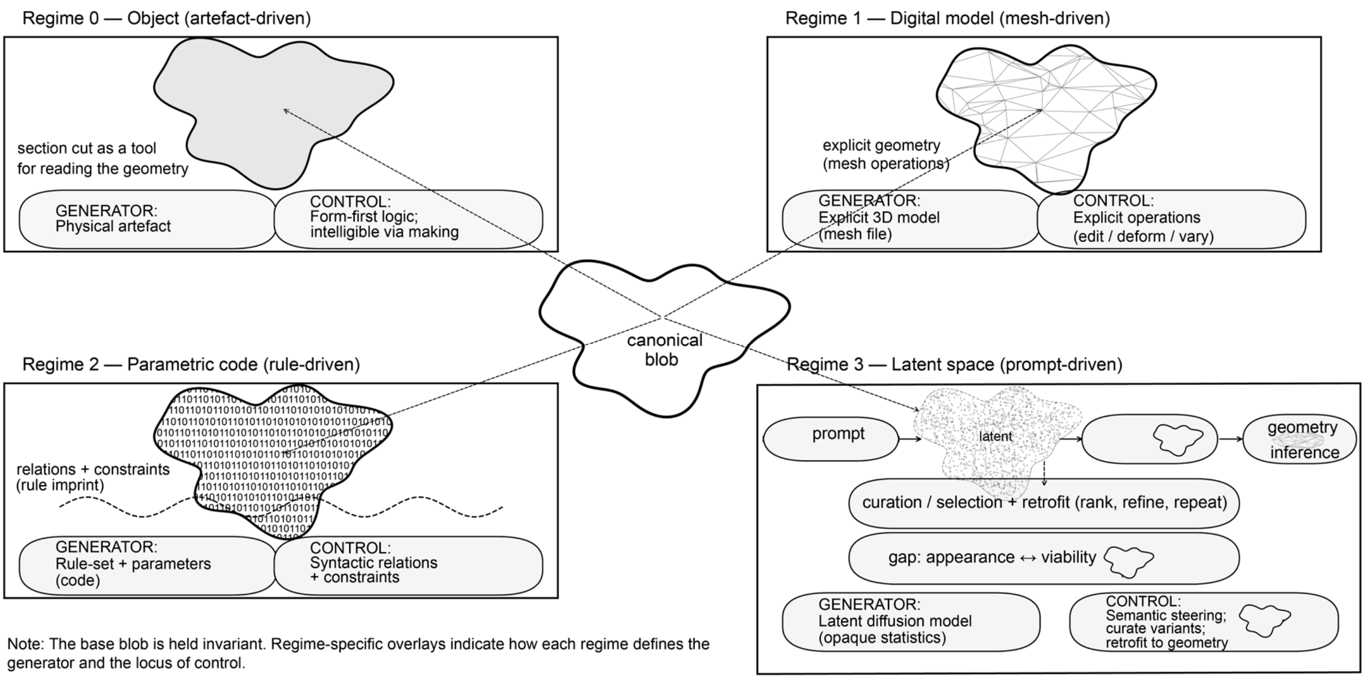

Figure 9. One morphology, four generative regimes ("family of blobs"). The same canonical blob is held invariant to compare how each regime defines the generator and the locus of control: object/artefact (Regime 0), explicit mesh operations (Regime 1), rule-driven parametric relations (Regime 2), and prompt-driven latent-space steering with selection and geometry inference (Regime 3). Regime 0—Object/artefact: form is anchored in a physical artefact that precedes programme rationalisation. Regime 1—Digital model: variation is produced through explicit mesh/topological operations on a legible 3D model. Regime 2—Parametric code: form emerges from rule-based relations, constraints, and discretization/panelization logic. Regime 3—Latent regime: form is initiated as an image-first artefact guided by semantic prompting, then reconstructed into geometry post-hoc. Graphic key: numbers identify regimes and/or ordered stages in the comparative schema; arrows indicate the direction of the generative/translation logic within each regime; dashed arrows/lines denote downstream or deferred steps (post-hoc rationalisation/reconstruction/validation); dots mark key decision points (selection, parameter adjustment, or constraint checkpoints) where agency is exercised in the workflow.

4. Comparative Cases: A Family of Blobs Across Regimes

This section turns the genealogy of Section 3 into an explicit comparative device. If Section 3 defined what each regime is and why the examples represent it, Section 4 asks how the same examples operate as cases: what kind of "initial artefact" they mobilize, which decision sequence they activate, and where the critical friction emerges between visual plausibility and tectonic legibility. The aim is not faithful reconstruction of the original projects, but a homogeneous set of diagrammatic reconstructions, consistent viewpoint, consistent graphic economy, enabling comparison of processes rather than styles. In that sense, the comparison is explicitly ontological: it tracks how the project's first "evidence" is produced, how it becomes legitimate, and how (or whether) it remains auditable as an evidentiary chain.

Visually, the section relies on two complementary families of material:

(1) rights-free white diagrammatic reconstructions of the historical cases (Figure 10a–c), replacing original imagery for copyright reasons;

(2) an authored Regime 3 "latent blob" developed as a GenAI design-research experiment (Figure 10d), where the interest is not "a new blob" but a regime shift: the geometric generator displaced toward an image-first artefact of high plausibility [1–4,18]. This is also where the framework's intentional asymmetry becomes operative: Regimes 1–2 remain internally continuous because they preserve a legible generator (model/code), whereas Regime 3 introduces a qualitative discontinuity by relocating first evidence to an image produced without prior geometry.

Figure 10 assembles the four comparative cases as a homogeneous set.

Having established the four cases (Figure 10), Figure 11 compresses each regime into a comparable operational sequence, so the subsections that follow can read each case against the same decision structure.

The four regime pipelines are summarized in Figure 11.

Regime 3 case-generation protocol (traceable sampling and rubric-based selection). To construct the Regime 3 comparative case, we executed a small design-research protocol using a text-to-image generative AI system (DALL·E, accessed via ChatGPT 5.2) under controlled representational conditions (fixed viewpoint, lighting, and context). The image set was produced through 22 batches of four outputs each (n = 88), keeping the representational conditions fixed while allowing geometric and tectonic variability across candidates. We retained the full generation trail and scoring sheet to ensure auditability of the selection process. Rather than treating a single output as an illustrative "hero image", we treated the image-first artefact as the outcome of traceable sampling and curation. A fixed semantic prompt was used to generate candidate images under constant viewpoint, horizon line, atmospheric conditions, and coastal context, so that differences across outputs could be attributed primarily to latent sampling and semantic steering rather than to changes in framing. In total, 22 batches were generated with four images each (n = 88). A curated subset was then assembled as a contact sheet (Figure 12), and one sample was selected as the representative image-first artefact (Figure 13) used as the Regime 3 case in the homogeneous comparative set (Figure 10d). The complete generation trail and scoring sheet were retained to ensure auditability of the selection.

https://doi.org/10.3390/architecture6010025

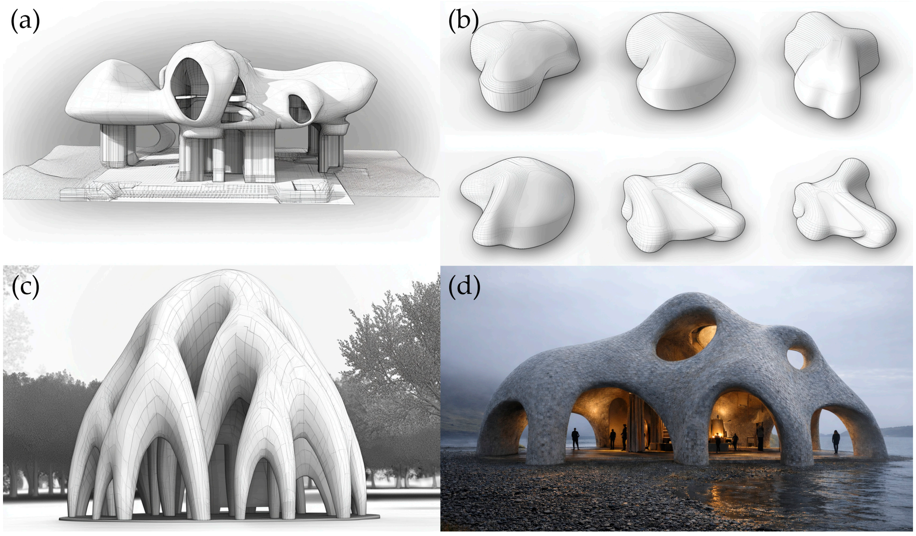

Figure 10. Comparative cases used in Section 4 (homogeneous viewpoint and graphic economy). (**a**) Regime 0—Endless House (diagrammatic reconstruction). (**b**) Regime 1—Embryological House (diagrammatic reconstruction). (**c**) Regime 2—Vaulted Willow (diagrammatic reconstruction). (**d**) Regime 3—GenAI-generated image-first artefact (selected latent-space sample), curated from a documented archive (n = 88; see Figure 12) using the rubric in Table 1 and subsequently analyzed through the friction map in Figure 14. Unlike Regimes 0–2, this sample does not derive from an explicit object/model/code; it precedes auditable geometry and therefore functions as the founding evidence of an image-first workflow whose reconstruction and tectonic validation are required post hoc.

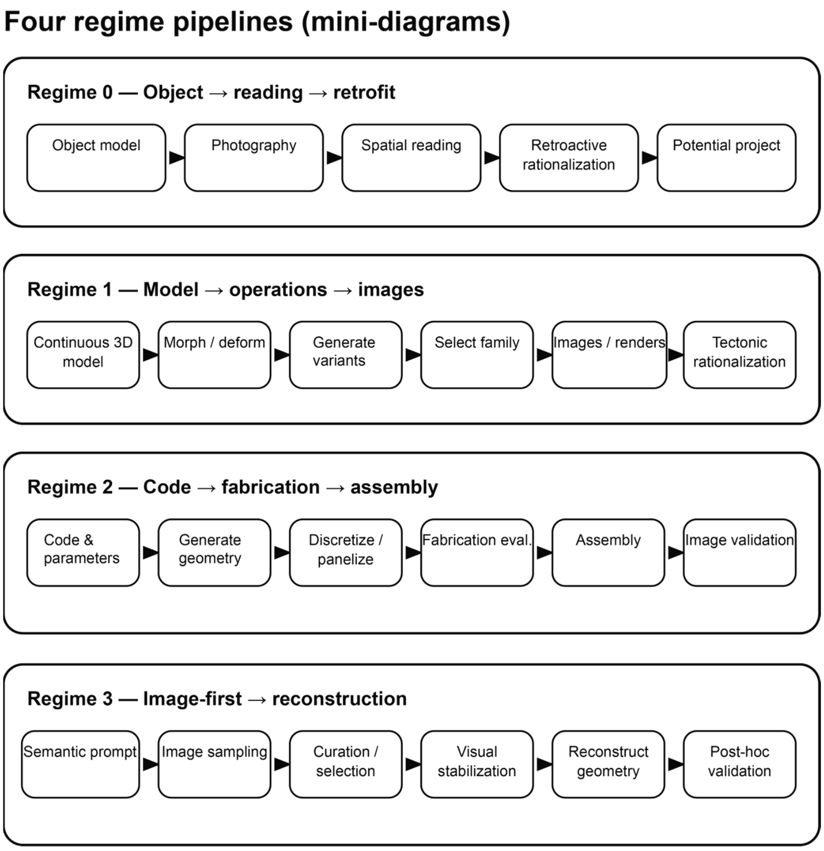

Figure 11. Four regime pipelines (mini-diagrams). Summary of the operational sequence in each regime, highlighting where the generator is located (object/model/code/prompted image) and where translation, evaluation, and retrofit are performed. The diagram also indicates where the "audit trail" is strongest (legible model/code) versus where it becomes inferential (image-first curation and post hoc reconstruction requirement).

Selection rubric (Table 1). Selection was performed using an explicit rubric designed to operationalize "evidentiary legibility" in the latent regime. The rubric scores five dimensions: Formal Legibility (FL), Syntactical Coherence (SC), Material–Tectonic Cues (MT), Semantic Alignment (SA), and Comparability (C) with the blob genealogy used across regimes. Importantly, this rubric formalizes a first moment of judgement—whether an image can legitimately function as early project evidence—prior to any explicit geometric reconstruction.

Table 1. Selection rubric for Regime 3 (evidentiary legibility prior to reconstruction).

Criterion	Abbrev.	What Is Evaluated (Operational Definition)	Typical Signals in the Image
Formal Legibility	FL	Degree to which the image contains observable cues of spatial and organizational order beyond atmosphere alone.	Volumetric hierarchy, readable voids/oculi, implied thickness, spatial layering, inhabitable scale cues
Syntactical Coherence	SC	Internal consistency of the overall form and its compositional logic (continuity, repetition, absence of contradictory signals).	Continuous shell logic, stable openings, coherent transitions, no conflicting geometrical cues
Material–Tectonic Cues	MT	Presence of material articulation and constructive hints that suggest (without resolving) a tectonic logic.	Seams, ribbing, panel-like segmentation, tolerances/fasteners cues, differentiated surface regimes
Semantic Alignment	SA	Alignment with the regime-specific intent (latent image-first artefact combining high plausibility with tectonic underdetermination) and avoidance of drift to unrelated imaginaries.	Reads as an image-first shell/pavilion, not as modular housing, object art, or unrelated typologies
Comparability	C	Ability to remain within the "blob genealogy" used for cross-regime comparison (0–3) so that differences read as regime shifts rather than stylistic outliers	Recognizable blob continuity, apertures/interiority, graded discretization echoes, discipline-consistent lineage

Scores were assigned on a 1–5 scale (5 = highest evidentiary legibility).

The blob is not a depiction of the latent space itself, but a single image sample produced from it under semantic constraints; in terms of the workflow in Figure 8, Figure 10d corresponds to the sampling stage prior to any explicit geometric reconstruction. The "image-first" status is therefore methodological, not representational: it defines where evidence starts and where verification is forced to begin.

4.1. Regime 0: Object as Generator, Endless House as Proto–Ready-Made

In Regime 0, the geometric generator is a material artefact that precedes disciplinary explicitness. Endless House functions as a limit case: the model is not merely representational but constitutive, a formal reality that forces programme, structure, and representation to adapt retroactively. In this sense, the project produces an early inversion of rationalist workflows: form does not result from a functional diagram but demands architectural translation afterwards [9–11].

Operational sequence (Regime 0; see Figure 11): object model, photography, spatial reading, retroactive rationalization, potential project.

Comparative criteria: agency concentrates on manipulating the object's continuity and interiority; evidence is primarily photographic and exhibitionary; tectonic viability is deferred or suspended. This makes Endless House a powerful anchor: the challenge of translating autonomous form into architecture does not begin with AI. What changes in the latent regime is not autonomy, but the generator's status, from a stubborn object to an opaque prediction [1,2,18]. This "Kiesler echo" will be used repeatedly as a measuring device: it clarifies that what is new in Regime 3 is not that form can precede function, but that the first artefact can be a statistically produced image whose material commitments are undecidable until reconstruction.

Figure 12. Regime 3—Contact sheet and rubric-based evaluation (excerpt from the 88-image archive). From 22 controlled batches (four outputs per batch; n = 88), a curated subset is assembled to make the sampling space visible under fixed viewpoint and context. Each candidate is scored using the five-criterion rubric (Table 1): Formal Legibility (FL), Syntactical Coherence (SC), Material–Tectonic Cues (MT), Semantic Alignment (SA), and Comparability (C). Each criterion is evaluated on a scale from 0 to 5, where 5 represents the highest score and 0 the lowest. The sheet illustrates typical outcomes of latent sampling: some images achieve high atmospheric plausibility but drift semantically away from the blob genealogy (low SA/C), while others intensify echoes of earlier regimes (object-like continuity or code-like discretization) at the expense of organizational legibility. This visualizes curation as an evidentiary operation rather than a stylistic preference.

Figure 13. Regime 3—Selected image-first artefact (latent-space sample) used as the comparative case in Figure 10d. The sample is selected not as a "best-looking" rendering but as the most defensible starting evidence under the rubric: it remains within the blob lineage (high comparability), sustains the regime-specific semantic target (high semantic alignment), and provides readable material–tectonic cues without resolving them metrically. In the workflow of Figure 8, this image corresponds to sampling-and-curation prior to any explicit geometric reconstruction; its "image-first" status is therefore methodological (where evidence begins), not representational (a depiction of latent space).

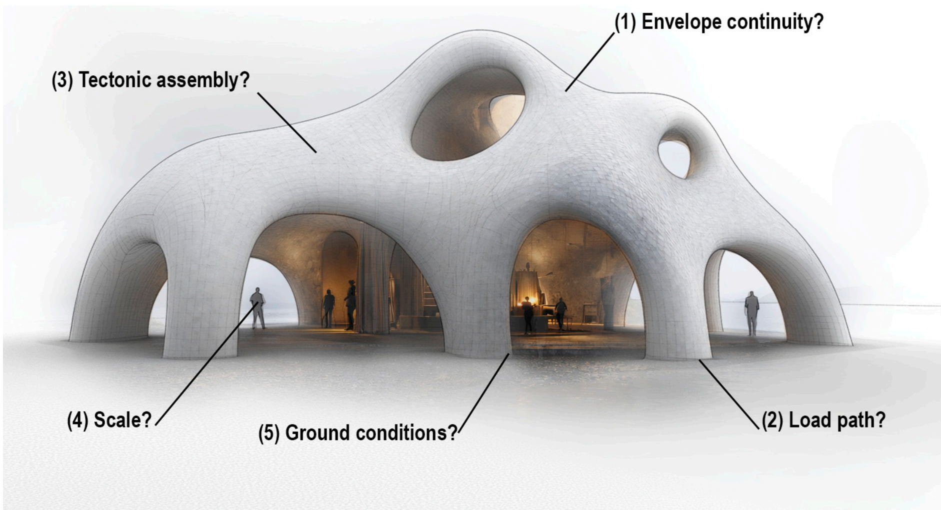

Figure 14. Regime 3—Friction map: operational diagnosis of the plausibility–tectonics gap. The selected artefact (Figure 13) is annotated to make translation pressure inspectable: callouts mark where visual plausibility forces underdetermined decisions during post hoc reconstruction and tectonic validation. Five categories are indicated: (1) geometric ambiguity (what geometry is implied), (2) structural ambiguity (how it works and where it bears), (3) material/tectonic ambiguity (thickness, assembly, joints), (4) metric/scale ambiguity (module, dimension, tolerance), and (5) detail ambiguity (ground contact, oculi edges, drainage, continuity). The map turns the gap from a general claim into a readable procedure: the image remains visually convincing, yet the architectural object remains incomplete as an evidentiary chain until these decisions are reconstructed and validated.

4.2. Regime 1: Digital Model, Embryological House and the Family as the Design Object

Regime 1 shifts the generator from physical artefact to a continuous digital model. Here, the project is not a singular form but a system of variations, a family derived from a shared geometric genotype. The key change is epistemic: the design object becomes the 3D file and its capacity for deformation, morphing, and controlled mutation [2,5].

Operational sequence (Regime 1; see Figure 11): continuous 3D model, morphing and deformation, generation of variants, family selection, images, and tectonic rationalization.

Comparative criteria: the generator is legible and auditable as explicit geometry; variation becomes a design criterion; the image remains derived from a model rather than foundational. This aligns with the historical understanding of digital architecture as a framework for managing geometric complexity while maintaining explicit control of the generator [2,6]. In contrast to Kiesler's singular object, the "form-first" tendency here is stabilized as a reproducible digital genotype, which preserves a clear audit trail from geometric operation to representational output.

4.3. Regime 2: Parametric Code, Vaulted Willow and the Authorship of the System

Regime 2 shifts the generator from explicit geometry to a system of rules, scripts, parameters, and protocols. Geometry becomes computed output, dependent on a possibility space. This transforms authorship: the architect becomes less a designer of a final object and more a designer of a system that produces objects. This partial loss of direct control is therefore not introduced by AI; it is structurally present in parametricism and in code-driven practices where the script is the primary design object [8,20].

Operational sequence (Regime 2; see Figure 11): code and parameters, geometric generation, discretization and panelization, material and fabrication evaluation, assembly, and image.

Comparative criteria: control is syntactic and explicit; tectonics is embedded through discretization and tolerances; the image remains subordinate, validating outcomes but not founding them. The regime already anticipates distributed authorship between designer, algorithm, machine, and assembly. The difference with the latent regime is not "more complexity", but a jump to a generator that is no longer auditable as a human-authored artefact, but statistical [1,2,18]. This is why Regimes 1 and 2 are treated as a threshold rather than a rupture: despite the shift from model to code, both preserve a legible generator and an explicit chain of justification.

4.4. Regime 3: Latent Space, from Prompt to Image, from Image to Geometry

Regime 3 is defined by a historical inversion: the first artefact is neither object, model, nor code, but a photoreal synthetic image of something with no prior explicit geometric substrate. The image does not represent a model; it produces a predictive appearance and therefore opens an unprecedented gap between visual plausibility and constructive legibility [1,4,18].

The latent blob is not presented as "another blob", but as an artefact designed to make the regime-shift visible. It is constructed to contain traces of the previous three regimes without collapsing into any of them: continuous interiority and large apertures, as memory of the object; smooth morphing transitions, as memory of the model; graded discretization and panel logic, as memory of the code; plus a distinctly contemporary condition of high atmospheric plausibility and subtle "too-coherent" cues of model mediation without resorting to obvious glitches [3,4,8,20]. In other words, it stages the genealogy as tension: the latent artefact appears to "inherit" earlier regimes while breaking their evidentiary logic.

Operational sequence (Regime 3; see Figure 11): semantic prompt, image generation, curation and selection, visual stabilization, post hoc geometric reconstruction requirement, and post hoc tectonic validation requirement.

The methodological pivot becomes explicit: design agency migrates toward semantic negotiation; control ceases to be direct and syntactic and becomes indirect and editorial. This is also where cognitive friction becomes structural: the architect performs an iterative "interrogation" of a black box through language, while confronting the curator's paradox, near-infinite variant generation versus the slow tempo of critical judgement required to select, stabilize, and justify a proposal.

Regime 3 protocol (documented sampling, rubric-based curation). To avoid treating Figure 10d as an ad hoc illustration, the Regime 3 case was produced through a traceable sampling-and-curation procedure. A fixed semantic prompt (reported below) was used under controlled representational constraints (constant viewpoint, atmosphere, and coastal setting), so that the output family remained comparable as a "blob lineage" rather than drifting across unrelated typologies. We generated 22 batches of four images each ($n = 88$). From this archive, a curated subset was assembled as a contact sheet (Figure 12) and evaluated using an explicit rubric (Table 1). One image was then selected as the representative image-first artefact (Figure 13) and used as the Regime 3 comparative case in Figure 10d. This makes the selection itself an evidentiary act: the architect's agency is exercised through criteria, not through direct geometric control.

Prompt (Regime 3): A speculative pavilion generated via latent interpolation: a continuous inhabitable shell with cave-like interior voids and large oculi. Smooth morphing geometry as if frozen mid-animation. A non-uniform gradient skin transitioning from seamless plaster-like mass to discretized tessellated shingles and ribbed striations; partial panelization with visible tolerances and fasteners in some zones, other zones smooth and monolithic. High visual plausibility yet tectonically ambiguous. Subtle latent artefacts (soft halos, coherent micro-noise, "almost-too-perfect" atmosphere). Cinematic foggy dawn, minimal landscape, neutral palette, photographic realism, crisp depth of field, architectural visualization. No text, no logos, no recognizable references. Fixed viewpoint.

Selection rubric (why these criteria). The rubric operationalizes "evidentiary legibility" in the latent regime: it scores whether an image can function as admissible early-stage evidence before any explicit geometry exists. The five criteria are Formal Legibility (FL), Syntactical Coherence (SC), Material–Tectonic Cues (MT), Semantic Alignment (SA), and Comparability (C). This separation is deliberate: at this stage, the question is not yet whether the image is buildable, but whether it is readable enough—as a disciplined artefact—to justify subsequent reconstruction work and tectonic verification.

Clarification. The blob is not a depiction of the latent space itself, but a single image sample produced from it under semantic constraints; in terms of the workflow in Figure 8, Figure 13 corresponds to the sampling-and-curation stage prior to any explicit geometric reconstruction. The "image-first" status is therefore methodological, not representational: it defines where evidence starts and where verification is forced to begin.

From selection to diagnosis: making the plausibility–tectonics gap inspectable. The selection rubric (Table 1) establishes a first moment of judgement in the latent regime: whether a predictive image is admissible as early-stage evidence before any explicit geometry exists. However, this evidentiary legibility should not be conflated with buildability. The plausibility–tectonics gap only becomes inspectable once the selected image is treated as an origin that demands translation into explicit geometry and tectonic commitments. To operationalize this second moment, we introduce a friction map (Figure 14): a structured annotation of the chosen Regime 3 sample that identifies where visual plausibility forces underdetermined decisions during post hoc reconstruction and validation. Table 2 complements the figure by formalizing the friction categories as a compact checklist, making the gap readable as a procedure rather than a generic claim [4,18].

Table 2. Friction-map checklist for Regime 3 (plausibility–tectonics diagnosis).

Friction Category	What Is Underdetermined in the Image-First Artefact	What Must Be Decided/Reconstructed (Examples of Required Commitments)
(1) Geometric ambiguity	The image suggests a continuous shell, but the exact geometry (curvature continuity, thickness strategy, edge conditions at oculi and arches) remains undecidable from appearance alone.	Define explicit surfaces/solids (e.g., NURBS mesh), continuity class, shell thickness logic, opening profiles, and geometric tolerances for fabrication.
(2) Structural ambiguity	The pavilion "reads" as stable, yet the load path (where it bears, how forces transfer through the shell, what acts as arch/rib/diaphragm) is not specified.	Establish structural system (shell action vs. ribbed frame vs. hybrid), supports/foundations, reinforcement strategy, and a verifiable load-transfer diagram.
(3) Material/tectonic ambiguity	Surface cues imply masonry-like or tiled assembly, but materiality is semantically plausible rather than tectonically defined (units, joints, fixings, waterproofing, thermal logic).	Specify material system (units/panels, joints, substructure), assembly sequence, detailing of seams, waterproofing/drainage, and performance requirements.
(4) Metric/scale ambiguity	Human figures provide approximate scale, but modules, spans, radii, and tolerances are not metrically anchored; the image does not define a reliable dimensional system.	Fix dimensional grid/module, spans and clearances, thickness-to-span ratios, tolerances, and construction constraints that stabilize the design beyond pictorial scale cues.
(5) Detail ambiguity	Critical junctions are visually smooth, but key details are unresolved (ground contact, oculi edges, transitions between interior lining and exterior skin, drainage paths).	Resolve edge details, base conditions, moisture/drainage strategy, continuity of finishes, interfaces between systems, and maintainability/ safety requirements.

Table 2 Friction categories for diagnosing the plausibility–tectonics gap in Regime 3. The table formalizes five recurrent types of translation pressure identified through the friction map (Figure 14). Each category corresponds to a dimension in which a visually plausible latent-space image remains underdetermined when treated as architectural evidence, requiring explicit geometric reconstruction and tectonic validation. The table functions as a procedural checklist rather than an evaluation score, making the plausibility–tectonics gap readable and inspectable as part of the image-first workflow.

To avoid leaving the plausibility–tectonics gap as a purely qualitative claim, we introduce a simple ordinal Gap Index (GI) that summarizes translation pressure as read from the selected image and comparable candidates. Unlike the selection rubric (Table 1), which evaluates whether a predictive image is admissible as early-stage evidence, the GI does not assess buildability in an engineering sense. Instead, it operationalizes the plausibility–tectonics gap as the degree of underdetermination that the image imposes when one attempts post hoc translation into explicit geometry, structural logic, tectonic assembly, scale, and detail. In this sense, GI functions as a compact audit of the friction categories formalized in Table 2.

GI is computed by scoring each of the five friction categories—(1) geometric ambiguity, (2) structural ambiguity, (3) material/tectonic ambiguity, (4) metric/scale ambiguity, and (5) detail ambiguity—on a 0–3 scale (0 = low translation pressure; 3 = very-high translation pressure). The sum yields a GI in the 0–15 range (Table 3). This index enables a controlled comparison across candidates from the same batch archive: some images may appear "more plausible" yet still demand high translation pressure, while others may be straightforward to reconstruct precisely because they collapse into readily available geometric primitives. Crucially, this makes explicit the methodological distinction on which our argument depends: evidentiary legibility (Table 1) and translation pressure (GI) are different assessments, and the plausibility–tectonics gap emerges in the interval between them.

Table 3. Ordinal plausibility–tectonics Gap Index (GI) derived from friction categories (0–3 scale). Each category is scored 0–3 (0 = low translation pressure; 3 = very high translation pressure). GI (0–15) summarizes translation pressure for comparable candidates. The symbol "#" identifies the selected reference case discussed from Figure 12.

Candidate (Contact Sheet ID)	Geometric	Structural	Material/Tectonic	Metric/Scale	Detail	GI Total	Gap Class
Selected sample (Figure 13/Regime 3)	2	3	2	2	2	11	High
Contact sheet #14 (primitives: cylinders/cubes)	0	0	1	1	1	3	Low
Contact sheet #7 (highly suggestive / difficult to execute)	3	3	3	2	3	14	High
Contact sheet #10 (highly suggestive / difficult to execute)	3	3	2	3	3	14	High
Contact sheet #6 (echoes Regime 0/2; moderate traceability)	2	2	2	2	2	10	High-Medium
Contact sheet #2 (echoes Regime 0/2; moderate traceability)	2	2	2	1	2	9	Medium

Suggested GI classes (optional): Low = 0–4; Medium = 5–9; High = 10–15. Adjust candidate IDs and scores to match your contact sheet.

Read together, Tables 1–3 and Figure 14 separate two moments that are often conflated in GenAI discourse: (i) evidentiary legibility (whether a predictive image is admissible as early-stage project evidence) and (ii) translation pressure (the set of underdetermined decisions that emerge when that image is treated as an origin requiring reconstruction and tectonic validation). This distinction clarifies that an image may score highly in the selection rubric while still exhibiting substantial tectonic indeterminacy (high GI), whereas low-legibility samples may amplify novelty potential at the cost of higher reconstruction effort and risk. Put differently, high traceability combined with a low GI tends to yield incremental novelty, because translation remains close to conventional solution spaces; very-high GI values can indicate greater novelty potential, but often coincide with weaker traceability, making the image a fragile evidentiary starting point. Our selected Regime 3 sample is therefore positioned deliberately in an intermediate zone: it preserves sufficient traceability to function as legitimate evidence, while maintaining a non-trivial GI that keeps open the space for novel tectonic, programmatic, and compositional development through reconstruction and validation. In the latent regime, novelty and responsibility are thus negotiated precisely within the gap between prediction and validation [1,2,4,18].

4.5. What the Comparison Produces: Inversion as a Critical Criterion

Figure 15 makes this inversion explicit as a comparative criterion, separating representational direction (generator → image versus image → geometry) from questions of style.

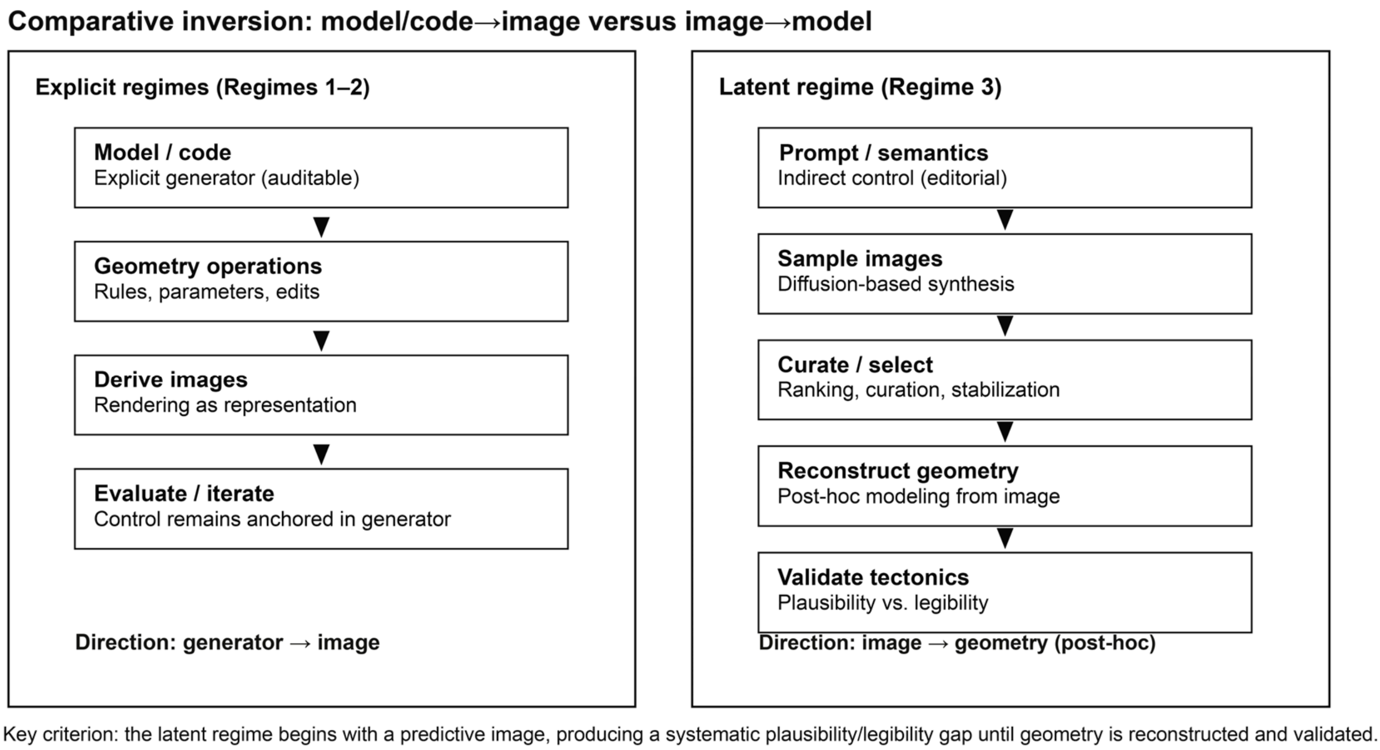

Figure 15. Comparative inversion: model/code → image versus image → model. In explicit regimes (1–2), images are derived from auditable generators; in the latent regime (3), a predictive image is produced first and only later reconstructed into explicit geometry, opening a systematic plausibility/legibility gap until post hoc validation.

Two operational outcomes follow. First, the comparison locates authorship shifts precisely: they already occur in the code regime, where the architect designs a possibility space, but the latent regime intensifies the shift because the generator is no longer legible as a human-authored system [2,8,20]. Second, it clarifies why the plausibility gap is not merely a "problem" but a criterion: until reconstruction and validation occur, the project's evidentiary chain is incomplete, so judgement is forced to operate under uncertainty (what looks credible may remain tectonically undecidable). This distinction is operationalized in Regime 3 through the separation between selection (Table 1; Figure 12) and diagnosis (Figure 14; Table 2), which makes the gap readable as a procedure rather than a general assertion.

Furthermore, it yields a disciplinary evaluation criterion: the gap between visual plausibility and tectonic legibility becomes the critical site for testing GenAI. The latent blob therefore functions not as a formal solution, but as a test of whether—and how—a predictive artefact can be reconstructed into consistent geometry and a viable system, and which protocols of judgement are required for that translation [4,18]. This is the point where the comparison begins to "hurt" in practice: a client can commit to an image before its tectonic commitments are knowable; a studio critique can reward atmosphere before auditability exists; professional responsibility becomes displaced across opaque tooling and retroactive engineering (Section 5 develops these implications concretely.) In this sense, the friction map (Figure 14) specifies where translation pressure concentrates (geometry, structure, tectonics, scale, and detail), clarifying why "high plausibility" does not complete the evidentiary chain [2,4,18].

To close the comparison, Figure 16 summarizes how each regime differs in auditability (legibility of the generator and its trace) and reproducibility, clarifying why the latent regime is methodologically powerful yet harder to verify.

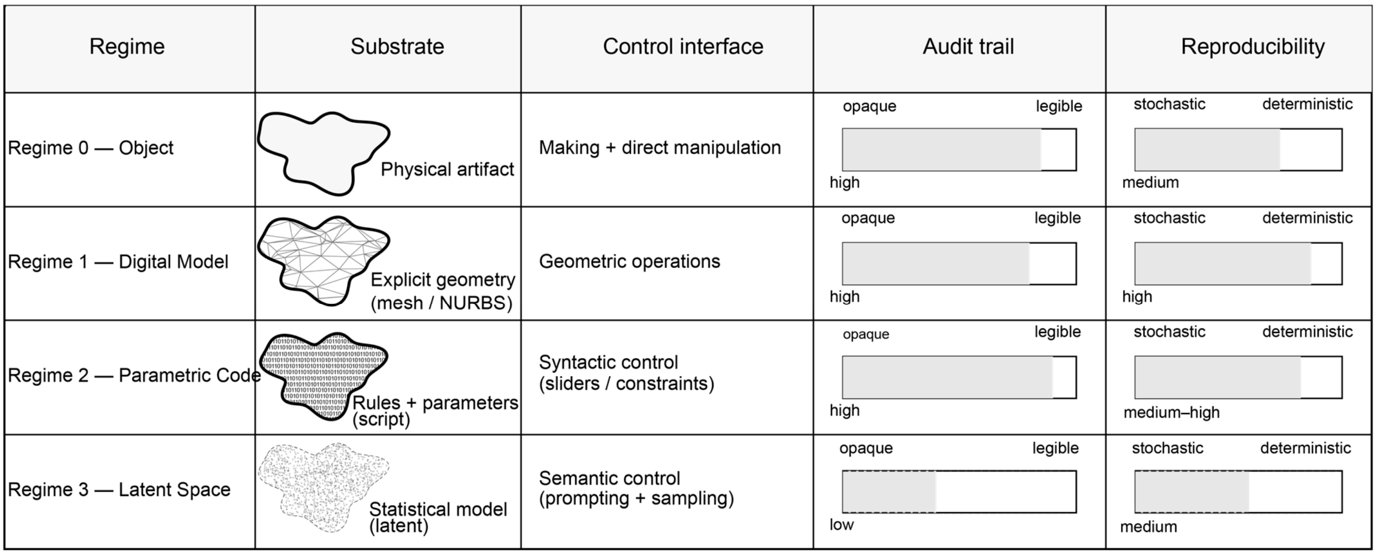

Note: The bars are conceptual (not measured). Regime 3 is shown as more opaque and more stochastic due to sampling and black-box weights.

Figure 16. Auditability and reproducibility of the geometric generator (schematic). Conceptual comparison across regimes showing how substrate and control interface affect the legibility of the audit trail and the degree of stochasticity in reproduction (not measured). The schematic also distinguishes reproducibility of an artefact (repeatable model/code execution) from reproducibility of an appearance (stochastic sampling under fixed prompts).

5. Discussion: Predictive Architecture, Authorship, and Pedagogy After the Latent Regime

This section extracts implications from the comparative framework. The key issue is not that AI produces blobs, but that it relocates the geometric generator into a statistical apparatus and makes the first project evidence a highly plausible synthetic image. Three consequences follow: a predictive condition of design, a mutation of authorship, and a reorganization of learning and canon formation [1–4]. Importantly, these consequences are not stylistic but evidentiary: they affect what counts as admissible early proof of a project, how validity is staged, and where responsibility accumulates in the chain from the image to buildable system. To avoid leaving these implications at the level of general scenarios, we anchor Section 5 in the Regime 3 micro-case developed in Section 4: the curated image-first artefact (Figure 13), its friction map diagnosis (Figure 14), and the two-stage evaluation protocol (Tables 1–3) that separates evidentiary legibility from translation pressure.

To clarify what changes operationally (not stylistically) in the latent regime, Figure 17 summarizes how design-labour shifts from making geometry to judging, curating, and translating outputs.

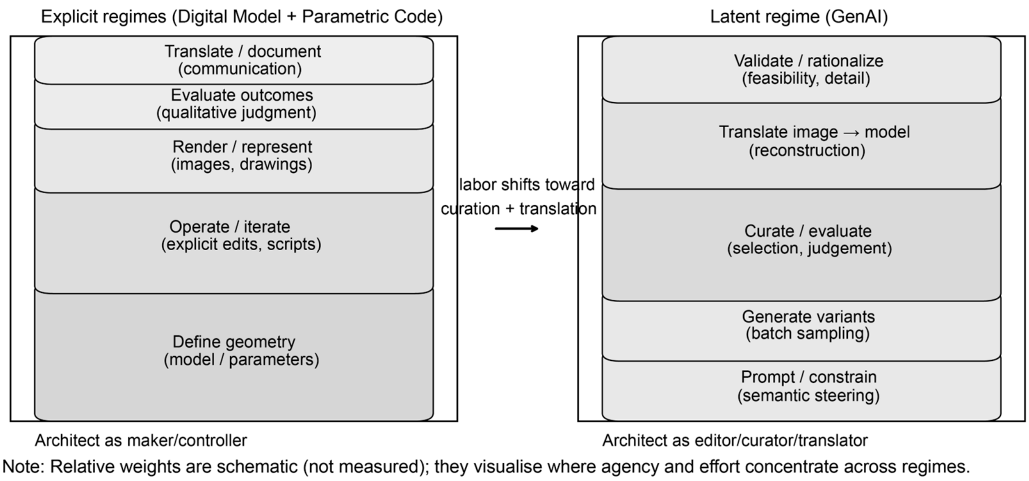

Figure 17. Reallocation of design labour: from making geometry to judging outputs. In explicit regimes (digital model and parametric code), effort concentrates on defining and operating auditable generators (model/parameters) and representing results. In the latent regime, effort shifts toward prompting, batch sampling, curation, and post hoc translation from image to model, with validation and detailing occurring downstream. Note: relative weights are schematic (not measured).

5.1. From Digital to Predictive Architecture, Why Inverting the Workflow Changes Validity

In Regimes 1 and 2, even when the form is extreme, plausibility is grounded in a legible artefact, model or code, and the image is derived. In the latent regime, plausibility becomes primarily visual: images appear resolved before explicit geometry exists. This produces a new friction: the image looks "done", while tectonic logic may not be [1,4,18]. Operationally, this relocates "project progress" from metric evidence (plans, sections, parametric dependencies) to pictorial evidence (a persuasive synthetic view), which can travel faster than verification. Section 4 operationalized this shift by distinguishing two moments that are often conflated: (i) selection/curation as an act of judgement that admits an image as early-stage evidence (Table 1), and (ii) diagnosis of translation pressure once that image is treated as an origin requiring reconstruction and validation (Figure 14; Table 2).

This tension can be summarized as a regime-dependent "plausibility gap": as the generator becomes less auditable, visual resolution tends to arrive earlier than constructive legibility, widening the distance between what looks finished and what can be convincingly reconstructed. Figure 18 visualizes the regime-dependent plausibility gap, clarifying that visual resolution can precede tectonic legibility, an effect that becomes most acute in the latent regime. In practice, this gap becomes immediately tangible in three recurring situations: (i) client alignment, when a client commits to an image whose reconstruction later requires major redesign, extra structure, or prohibitive cost, (ii) studio critique, when early evaluation rewards atmosphere and coherence of appearance before an audit trail exists, and (iii) professional accountability, when responsibility is distributed across opaque generation, human selection, and downstream engineering validation. Rather than introducing additional external anecdotes, we show how these three situations are already implicit in the Regime 3 micro-case: the friction map marks where commitments would have to be made (structure, assembly, scale, detail), and the gap index makes comparable "translation pressure" readable across candidates from the same archive.

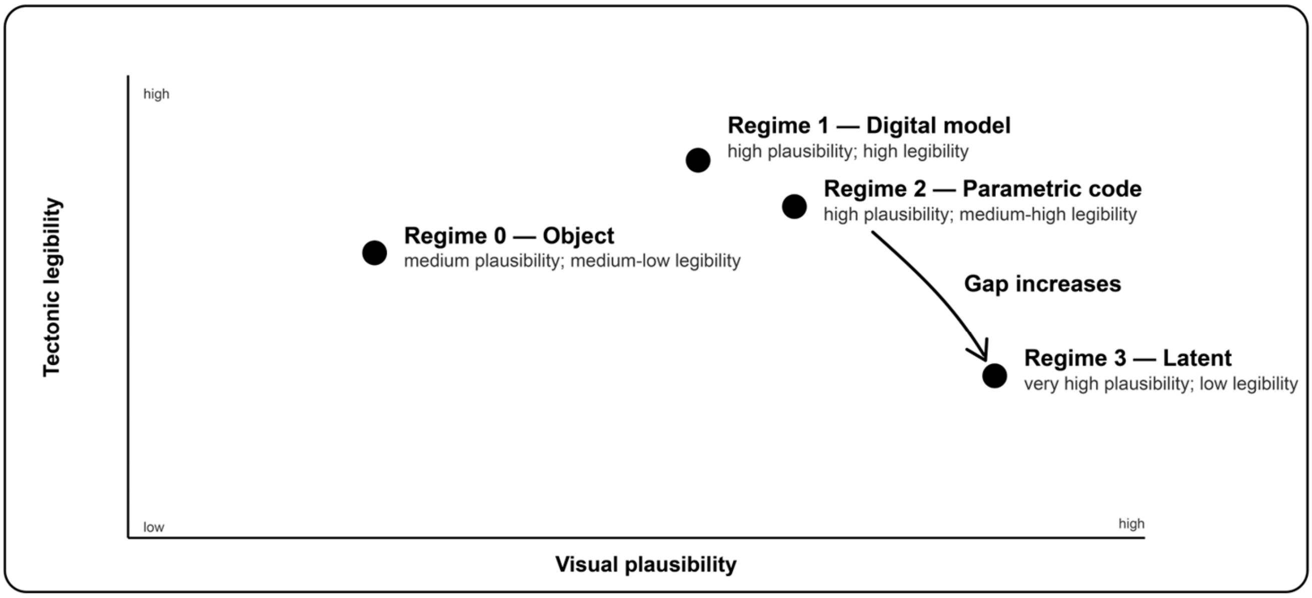

Note: positions are conceptual (not measured). Latent workflows can yield early visual resolution before an explicit model exists.

Figure 18. Plausibility gap (schematic): visual plausibility versus tectonic legibility across regimes. In explicit regimes (1–2), plausibility remains anchored in an auditable generator (model or code), so legibility can be evaluated in parallel. In the latent regime (3), the first evidence is a high-plausibility synthetic image produced before an explicit model exists, increasing the risk of a systematic plausibility/legibility mismatch until post hoc reconstruction and validation. Note: positions are conceptual (not measured). The diagram should be read as an evidentiary mismatch: appearance reaches "closure" earlier than tectonic justification. Graphic key: arrows indicate the widening displacement between visual plausibility and tectonic legibility as the generator becomes less auditable toward the latent regime.

Predictive architecture is not a style but an operative condition: projects begin as visual predictions and are stabilized afterwards as geometry, detail, and system. This requires rethinking what counts as "project progress" and which evidence is acceptable in early phases [3,4]. Concretely, it motivates explicit protocols of translation, for example, requiring that any image used as a decision artefact is accompanied by a stated reconstruction hypothesis, a set of geometric commitments, and a verification plan to close the plausibility gap. In our protocol (Section 4.4), this requirement is formalized as a two-stage chain: an admissibility judgement (Table 1) followed by an explicit articulation of underdeterminations (Figure 14; Table 2) and a compact audit of translation pressure (Table 3) [2,4,18].

5.2. Authorship and Control, from Syntactic to Semantic Agency

Loss of direct control does not begin with AI; it is anticipated by the parametric regime, where the architect designs a possibility space and claims authorship in rules and protocols [8,20]. The qualitative difference in the latent regime is twofold: control becomes semantic and indirect, and the generator is no longer auditable as a human artefact. This opacity alters the instrument's epistemic status and therefore the responsibility of judgement [1,2,18]. Figure 19 summarizes this shift by separating the *control interface* (how the architect intervenes) from the *audit trail* (how legible and reconstructable the generator remains across regimes).

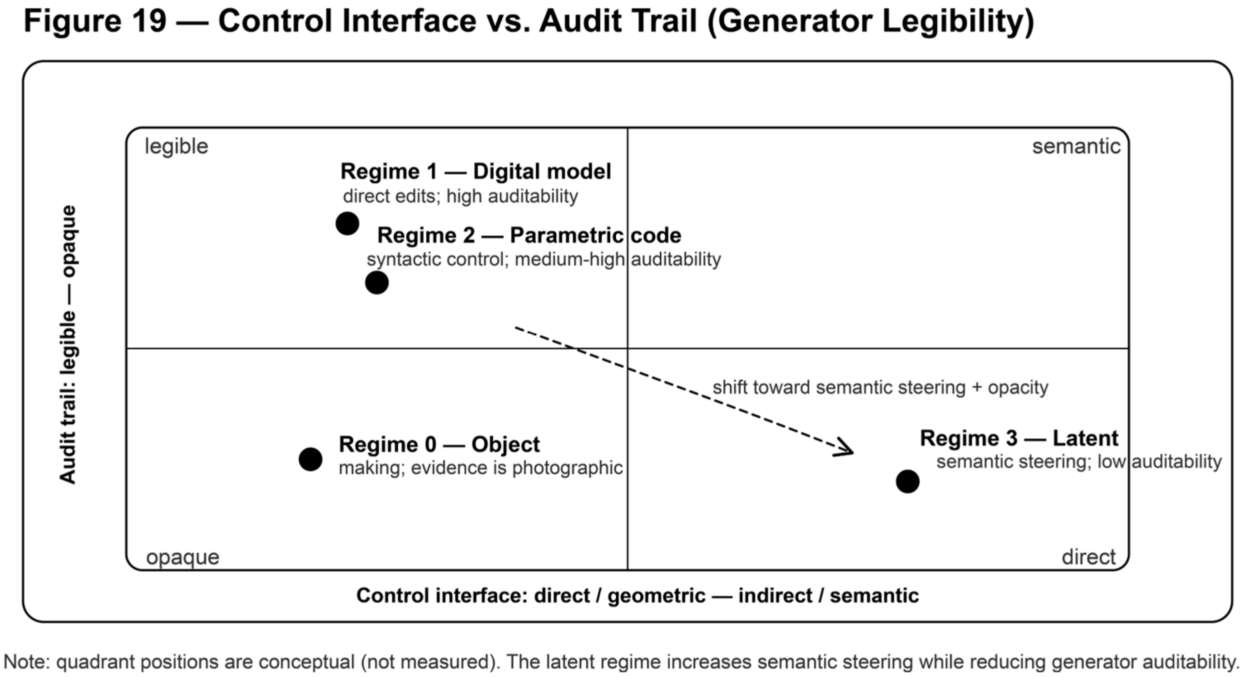

Figure 19. Control interface versus audit trail (schematic): direct/syntactic control and generator legibility across regimes. Regimes 1–2 retain an auditable generator (explicit geometry or code), even when agency is partially delegated to rules and protocols. Regime 0 relies on direct-making but leaves a weaker computational audit trail, with evidence often mediated through photography and exhibition. In the latent regime (3), control becomes predominantly semantic (prompting, sampling, curation) while the generator becomes comparatively opaque, increasing the burden of post hoc reconstruction and critical judgement. Note: quadrant positions are conceptual (not measured). Graphic key: dashed arrows indicate conceptual trajectories of regime transition (shifts in control interface and auditability), not measured flows—i.e., the direction in which agency migrates toward semantic control while generator legibility decreases.

The shift therefore changes authorship at the level of skill and cognition, not only at the level of tasks. Semantic agency introduces cognitive friction: the architect must interrogate a black box through language, iteratively probing what the system can and cannot stabilize, while absorbing the curator's paradox, the system produces variants faster than judgement can responsibly select them. This can lead either to analytic paralysis (too many plausible options) or to superficial selection (choosing by atmosphere because there is no time to verify).

In this regime, "virtuosity" relocates once more: no longer mastery of drawing, modelling, or even coding alone, but the capacity to construct a robust critical scaffold to validate, document, and justify the passage from image to geometry and from geometry to buildable system. In other words, authorship increasingly lies in producing criteria for judgement and in maintaining an explicit audit trail under conditions of opacity. The Regime 3 protocol makes this non-metaphorical: curation is documented as an evaluative act (Table 1), and the burden of judgement is displaced toward explicitly naming what the image does not determine (Figure 14; Table 2), i.e., the points where responsibility will concentrate during reconstruction.

5.3. From Designing Objects to Designing Canons, Archives, and Datasets

If latent space is learned from corpora, then the dataset functions as an operative canon: it defines recognizability, plausibility, and what the model can imagine. This shifts part of the debate from form to archive. Innovation tends to appear as plausible recombination within the learned manifold rather than programmatic rupture [1,2]. The implication is not necessarily negative, but it demands critical practices of curation, traceability, and explicit taxonomies in order to evaluate variation and bias [1,3]. Figure 20 condenses this argument by treating the dataset as an operative canon: curation and annotation decisions shape the learned manifold, which in turn conditions what kinds of image families appear most "natural" and reproducible.

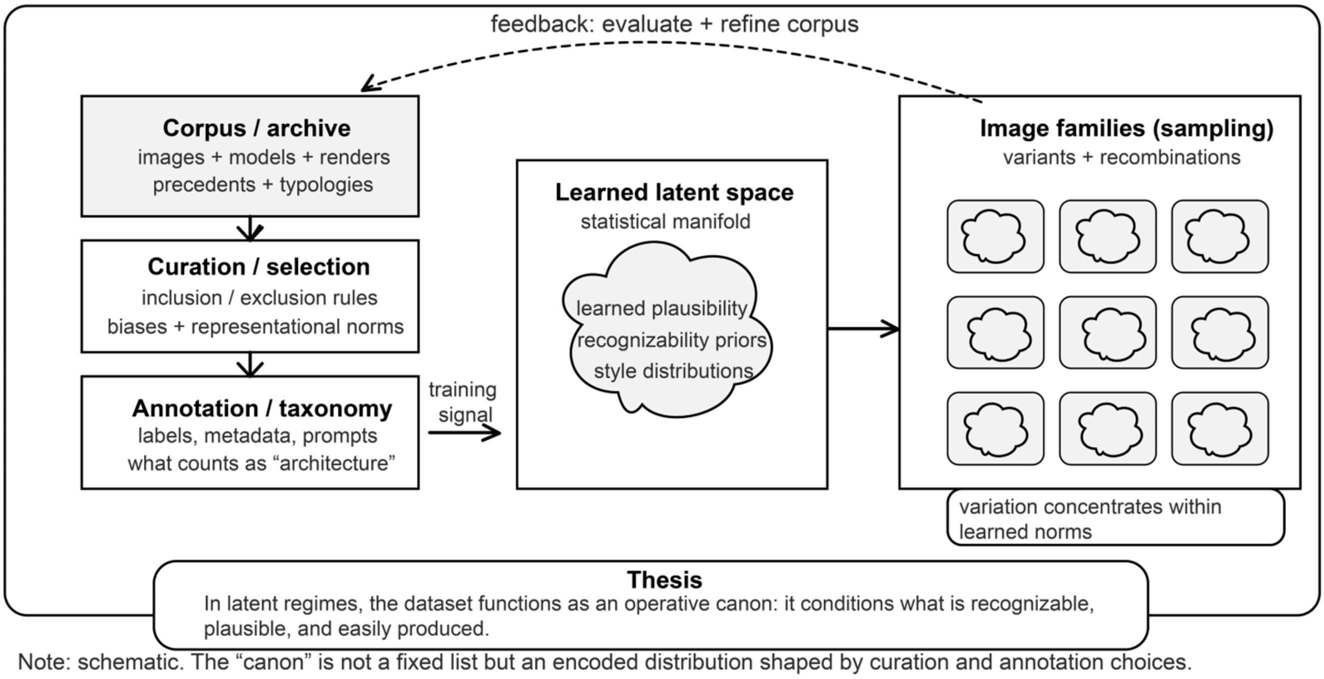

Figure 20. Dataset as canon (schematic): how corpus choices shape latent-space plausibility and image families. A curated and annotated archive functions as an implicit canon by encoding representational norms and distributions; sampling then tends to concentrate variation within those learned priors. This shifts part of architectural critique from form alone to the construction, documentation, and accountability of the underlying archive. Note: schematic; the "canon" is an encoded distribution shaped by curation and annotation choices. Graphic key: arrows indicate the direction of influence/flow (curation and annotation → learned latent priors → sampling → resulting image families), while grey fields/lines denote the implicit statistical distribution ("canon")—i.e., the regions of higher plausibility/probability within which variation tends to concentrate. The dashed arrow indicates the feedback loop, whereby outputs (selected image families, evaluation criteria, or downstream performance/validation) inform subsequent curation, annotation, and corpus refinement.

In this sense, curation is not an auxiliary activity but an ontological operation: it determines which architectures become statistically "thinkable" and which remain improbable, and it defines what the system returns as plausible without explanation. This also reframes critique toward documentation practices, provenance of training material, and the explicit declaration of representational priors when images are used as evidence in design decision-making. At the scale of a single design-research experiment, Table 1 already functions as a miniature version of this principle: it declares the representational priors by which certain images are admitted as "architecturally legible" evidence and others are rejected as drift, thereby making curation itself inspectable.

5.4. Pedagogy and Practice, from Virtuoso Production to Critical Editing in Multi-Agent Ecologies

In the latent regime, the core competence shifts from producing a single image to navigating families, filtering, editing, and stabilizing outputs. This is already visible in AI-driven design ecologies in which agency is distributed among image generators, modelling tools, evaluation routines, fabrication, and human judgement [4]. Pedagogically, this suggests reorganizing learning around three capacities: critical visual literacy, image-to-geometry translation, and tectonic verification behind persuasive appearance [3,4,7]. Figure 21 diagrams this as a multi-agent workflow: prompt-driven generation produces families, but architectural work concentrates in curation, reconstruction, and verification—an iterative loop that integrates tools, expertise, and judgement.

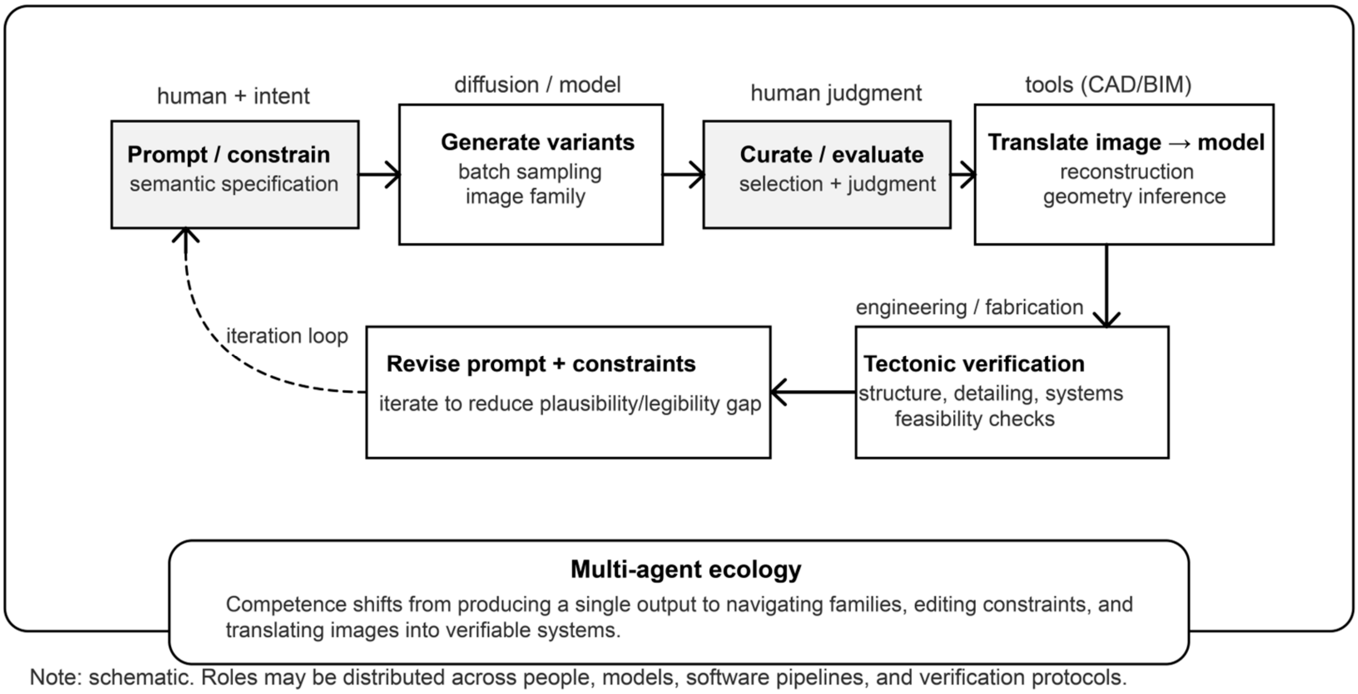

Figure 21. Agent ecology (schematic): distributed workflow from prompting to reconstruction and tectonic verification. In the latent regime, production shifts from crafting a single object to navigating variant families (batch sampling), curating outputs, translating selected images into explicit geometry (CAD/BIM reconstruction), and validating tectonic viability through systems and detailing checks. The loop closes by revising prompts and constraints to reduce the plausibility/legibility gap across iterations. Note: schematic; roles may be distributed across people, models, software pipelines, and verification protocols. Graphic key: arrows indicate the direction of the workflow and information handoffs between stages/agents (prompting → sampling → curation → reconstruction → verification), while dashed arrows indicate feedback loops (verification outcomes and reconstruction constraints informing subsequent prompts, sampling strategies, and selection criteria). Grey elements denote human intervention points (judgement, selection, editing, interpretation, and decision-making) within the otherwise tool-mediated pipeline.

The pedagogical consequence can be stated concretely: evaluation criteria must distinguish persuasion from verification. A studio rubric can require that any AI-generated image presented as a design decision is paired with (i) a reconstruction log (what was inferred, what remains ambiguous), (ii) at least one explicit geometric test (sections, structural hypothesis, envelope logic), and (iii) a brief risk statement describing where the plausibility gap is likely to appear. In our terms, this becomes teachable as a two-stage assessment: evidentiary legibility (Table 1) plus translation pressure (Tables 2 and 3), so that students are evaluated not on atmosphere alone but on whether they can make the plausibility gap legible and actionable [4,18].

Likewise in practice, the same logic supports clearer client communication: an image is presented explicitly as a predictive artefact, not as a commitment, until a reconstruction milestone is reached. This reduces avoidable friction where a visually "finished" proposal later collapses under cost, regulation, or constructability constraints.

Finally, the multiagent ecology implies a redistribution of responsibility: selecting a plausible image is not neutral. It is already a decision that constrains downstream engineering, detailing, and liability, which strengthens the case for keeping an auditable trail of prompts, selections, reconstructions, and verification steps. In Regime 3, the audit trail begins with selection and only becomes complete once underdeterminations have been declared and closed through reconstruction and validation; the micro-case (Figures 13 and 14; Tables 1–3) provides a compact template for such documentation [2,4,18].

6. Conclusions

This article does not evaluate generative AI as a new rendering tool, nor as a simple third turn inside a linear narrative of digital innovation. The claim is more specific: AI reconfigures architectural design because it shifts the geometric generator into latent space and makes the project begin, in an increasingly systematic way, with a synthetic image that precedes the model [1–4]. The resulting transformation is ontological at a disciplinary level: it changes the structure of project legitimacy, that is, what can count as early evidence, how judgement is staged, and how an audit trail can be maintained from appearance to buildable system. Accordingly, the paper's contribution is not a stylistic category but a methodological framework for treating predictive images as architectural evidence—separating admissibility, translation pressure, and responsibility along the image $\rightarrow$ geometry $\rightarrow$ tectonics chain.

6.1. Main Conclusion: Four Regimes, One Ontological Displacement

The genealogy proposed traces a coherent sequence: object, model, code, latent space. The blob family allows the paper to argue that the decisive change is not morphological, continuity, porosity, or discretization, but ontological: where the first project evidence resides and what kind of artefact founds the decision [2,3,8–11,20]. Regimes 1 and 2 are therefore treated as consecutive and internally coherent displacements within the same evidentiary logic: despite differences in agency, both preserve a legible, auditable generator (model or code) and an explicit chain of justification. The latent regime constitutes the qualitative break because it allows an image to function as founding evidence prior to explicit geometry, relocating verification to post hoc reconstruction. In this setting, architectural judgement is no longer primarily the verification of an explicit generator; it becomes the disciplined management of uncertainty—deciding which predictive images can count as evidence, and under which conditions they must be stabilized into geometry and tectonic commitments.

6.2. Operational Finding: The Plausibility–Tectonics Gap as a Critical Test

The latent regime intensifies a new gap: images of extreme plausibility without guaranteed constructive coherence. This does not invalidate the regime; it defines its critical testing ground. The GenAI blob case is proposed as an experiment precisely to evaluate how geometry and tectonics can be reconstructed from predictive appearance, and which disciplinary protocols are needed for that translation [4,18]. In this sense, the plausibility gap is not a rhetorical diagnosis but an evaluative criterion: it specifies where design research must operate (translation, auditability, verification) and where responsibility concentrates when a persuasive image outruns tectonic legibility. Rather than leaving the plausibility–tectonics gap as a qualitative claim, the revised framework makes it operational and inspectable: selection rubrics identify admissibility as evidence, friction mapping externalizes the specific sites of underdetermination, and an ordinal gap index can summarize translation pressure across comparable candidates.

6.3. Implications: Authorship, Canon, Responsibility

Authorship had already shifted in parametricism toward system design. AI radicalizes this displacement through the generator's opacity and the shift from syntactic to semantic control. In parallel, the dataset becomes a disciplinary locus of power, an operative canon that conditions what is imaginable. These transformations require rethinking responsibility, criteria of judgement, and legitimation in contemporary design culture [1,2,8]. The architect's expertise therefore relocates again: beyond producing geometry (and beyond writing code), toward constructing and maintaining criteria, curatorial decisions, and a defensible audit trail under conditions of semantic control and black-box mediation. Virtuosity relocates toward the construction of critical scaffolds—protocols, criteria, and traceable audit trails—capable of stabilizing predictive images into explicit geometry and tectonic rationality under conditions of opacity.

6.4. Future Work

Three lines of continuation follow directly: developing reproducible image-to-model reconstruction protocols combining manual modelling and neural auxiliary tools; investigating dataset curation and architectural taxonomies to understand how training canons bias outputs; and formulating pedagogical frameworks for critical operation within multi-agent ecologies integrating image, geometry, fabrication, and evaluation [1,2,4,6,7]. A fourth, closely related line concerns accountability: specifying how prompts, selections, reconstructions, and validation steps can be documented as a professional record when early "evidence" is a predictive image rather than a legible model. Figure 22 synthesizes the paper's contributions in one view: the four regimes and the displacement of the geometric generator, the workflow inversion that defines the latent regime, and the downstream implications for evaluation and authorship.

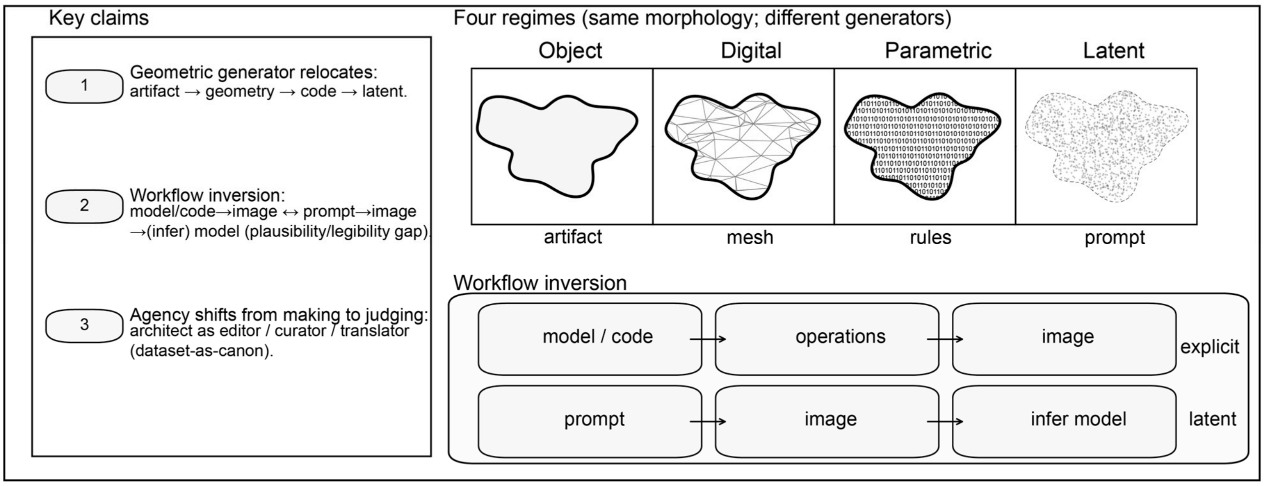

Figure 22. Summary of contributions (schematic): predictive architecture as a shift in the geometric generator. The figure condenses three claims developed across the paper: (1) the geometric generator relocates from artefact → explicit geometry → code → latent statistical space; (2) a workflow inversion differentiates the latent regime, where a predictive image is produced first and only later reconstructed into explicit geometry; and (3) design agency shifts from producing geometry to curating, translating, and validating outputs under conditions of visual plausibility and post hoc tectonic verification. Note: schematic (not measured). The summary also implies an evidentiary consequence: as the generator becomes less auditable, the burden of verification and documentation shifts downstream, making the plausibility gap the primary site of disciplinary judgement.

Author Contributions: Conceptualization, J.C.L.C. and C.E.S.M.; methodology, J.C.L.C. and C.E.S.M.; software, J.C.L.C. and C.E.S.M.; validation, J.C.L.C. and C.E.S.M.; formal analysis, J.C.L.C. and C.E.S.M.; investigation, J.C.L.C. and C.E.S.M.; resources, J.C.L.C. and C.E.S.M.; data curation, J.C.L.C. and C.E.S.M.; writing—original draft preparation, J.C.L.C. and C.E.S.M.; writing—review and editing, J.C.L.C. and C.E.S.M.; visualization, J.C.L.C. and C.E.S.M.; supervision, J.C.L.C. and C.E.S.M.; project administration, J.C.L.C. and C.E.S.M.; funding acquisition, J.C.L.C. and C.E.S.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data supporting the reported results are largely included within the article (including the Regime 3 prompt, the evaluation rubric, and the reported scoring/metrics). Additional materials underpinning the design-research protocol (e.g., batch-generation log/audit trail, extended scoring sheets, and full contact-sheet archive) are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Carpo, M. Generative AI, Imitation Learning, and the Automation of Tacit Knowledge. *a+u Archit. Urban.* **2024**, *646*, 104–109.
2. Carpo, M. *The Digital Turn in Architecture 1992–2012*; Wiley: Chichester, UK, 2012.
3. Young, M. *Reality Modeled After Images: Architecture and Aesthetics After the Digital Image*; Routledge: New York, NY, USA, 2022.
4. del Campo, M.; Manninger, S. Strange but Familiar Enough: The Design Ecology of Neural Architecture. *Archit. Des.* **2022**, *92*, 38–45. [CrossRef]
5. Lynn, G. *Animate Form*; Princeton Architectural Press: New York, NY, USA, 1999.
6. Kolarevic, B. *Architecture in the Digital Age: Design and Manufacturing*; Spon Press: New York, NY, USA, 2003.
7. Reiser, J.; Umemoto, N. *Atlas of Novel Tectonics*; Princeton Architectural Press: New York, NY, USA, 2006.
8. Schumacher, P. Parametricism: A New Global Style for Architecture and Urban Design. *Archit. Des.* **2009**, *79*, 14–23. [CrossRef]
9. Phillips, L. (Ed.) *Frederick Kiesler*; Whitney Museum of American Art: New York, NY, USA, 1989.
10. Kiesler, F. *Inside the Endless House*; Simon and Schuster: New York, NY, USA, 1966.
11. The Museum of Modern Art (MoMA). Visionary Architecture (Exhibition, 29 September–4 December 1960). Available online: https://www.moma.org/calendar/exhibitions/2554 (accessed on 10 December 2025).
12. Alberti, L.B. *On Painting*; Spencer, J.R., Trans.; Yale University Press: New Haven, CT, USA, 1956.
13. Alberti, L.B. *On the Art of Building in Ten Books*; Rykwert, J.; Leach, N.; Tavernor, R., Trans.; MIT Press: Cambridge, MA, USA, 1988.
14. Sontag, S. *On Photography*; Farrar, Straus and Giroux: New York, NY, USA, 1977.
15. Tomkins, C. *Duchamp: A Biography*; Henry Holt: New York, NY, USA, 1996.
16. Flusser, V. *Towards a Philosophy of Photography*; Originally published in German as Für eine Philosophie der Fotografie [M1.1]1983; Reaktion Books: London, UK, 2000.
17. Colomina, B. *Privacy and Publicity: Modern Architecture as Mass Media*; MIT Press: Cambridge, MA, USA, 1994.
18. Carpo, M. AI Tectonics, or the Culture Wars of Building Technology. *Archit. Intell.* **2025**, *4*, 2. [CrossRef]
19. Jodidio, P. *Guggenheim Bilbao*; Taschen: Cologne, Germany, 2008.
20. MARC FORNES/THEVERYMANY. Vaulted Willow. Project Page. 2014. Available online: https://theverymany.com/public-art/11-edmonton (accessed on 10 December 2025).
21. Chaillou, S. *AI + Architecture: Towards a New Approach*; Birkhäuser: Basel, Switzerland, 2022.
22. Huang, S.-Y. The Connectionist Turn: How Contemporary Generative AI Reshapes Architectural Rationality. *Architecture* **2025**, *5*, 132. [CrossRef]
23. Benjamin, W. The Work of Art in the Age of Mechanical Reproduction. In *Illuminations*; Schocken Books: New York, NY, USA, 1968.
24. Crary, J. *Techniques of the Observer: On Vision and Modernity in the Nineteenth Century*; MIT Press: Cambridge, MA, USA, 1990.