

From Rendering to Semantization: Intensified Architectural Semantics with Generative AI

Generative AI is often framed in architecture as a faster rendering engine, able to turn sketches, models, and prompts into persuasive images. Such a view collapses two distinct operations: rendering and semantization. Rendering concerns appearance; semantization concerns the process through which geometry becomes meaningful, legible, and recognizable as architecture. Building on this distinction, the article proposes three regimes of semantization: natural, anticipatory, and intensified. While natural semantization emerges through use, material identity, and time, and anticipatory semantization is associated with architectural rendering, intensified semantization names the condition introduced by GenAI, in which meaning is accelerated, thickened, and partially destabilized through probabilistic image formation.

Methodologically, the paper develops sequential protocols through two case studies. The first

starts from a deformed carton package and follows its transformation from everyday object to neutralized meta-object, then to anticipatory architectural rendering, and finally to three branches of intensified semantization: typological, atmospheric, and affective. The second begins with the Barcelona Pavilion and subjects it to controlled semantic displacement in order to test thresholds of recognizability and defamiliarization. Across both cases, the article introduces semantic inflation as the failure mode in which intensification becomes overcoded, literal, or cliché-driven.

The paper concludes that GenAI should not be understood only as an instrument of improved visualization, but as a semantic laboratory for architecture. It enables branching, iterative, and comparable transformations through which architectural meaning can be produced, displaced, and evaluated.



Cintya Eva Sánchez Morales is an architect, researcher, and PhD candidate (University of Innsbruck/University of Granada). She teaches at the University of Granada and has taught at Lund University and Innsbruck. Co-founder of Studio Soqotra, her research explores abstraction, graphic communication, and generative AI in architectural design



José Carlos López Cervantes is an architect, researcher, and PhD candidate (University of Innsbruck/University of Granada). He teaches at the University of Granada and has taught at the University of Innsbruck, the University of Applied Arts Vienna, IAAC Barcelona and Lund University. His research focuses on architectural representation, digital design, and generative AI

Keywords:
Generative AI; Architectural Semantics; Rendering; Defamiliarization; Architectural Grammar

1. INTRODUCTION

Generative artificial intelligence has entered architectural culture with unusual speed. In professional practice, academic environments, and design competitions, it is often understood as an advanced image-making tool capable of transforming sketches, rough models, and textual prompts into persuasive architectural visuals. Within this widespread interpretation, GenAI is primarily framed as a new rendering engine: faster, more accessible, and less technically demanding than conventional visualization workflows. From this perspective, its main contribution lies in accelerating representation.

This reading, however, is only partially adequate. While it correctly identifies GenAI as a powerful visual instrument, it tends to collapse two distinct operations into one: rendering and semantization. Rendering concerns how architecture appears; semantization concerns how architecture becomes meaningful, legible, and culturally recognizable. A geometry does not become architectural only because it looks realistic. It becomes architectural when it begins to signify scale, use, tectonic plausibility, contextual belonging, material character, and inhabitation. For this reason, the present article argues that GenAI should not be understood only as a better rendering device, but as an agent of intensified architectural semantization.

This distinction is not entirely new. Architectural representation has long operated as a mechanism of semantic anticipation. Drawings, models, collages, and renders do not merely depict future objects; they progressively load them with cues that make them readable as buildings and as particular kinds of buildings. In this sense, rendering already exceeds pure visualization. It anticipates atmosphere, use, weathering, and cultural legibility before the building exists. A render does not simply show a project; it projects a world. This article retains that insight from the extended abstract, but develops it further by arguing that generative

AI does not simply continue this anticipatory role. It intensifies it. Rather than stabilizing one coherent future image, it can accelerate, thicken, and partially destabilize the semantic regime through which geometry is read as architecture.

The paper therefore proposes a three-part distinction. First, natural semantization refers to the meaning an object acquires through real-world use, material identity, cultural recognition, and time. Second, anticipatory semantization describes the operation through which architectural representation, and especially rendering, projects a plausible semantic future for a geometry by making visible its scale, atmosphere, materiality, and probable use. Third, intensified semantization names the specific condition introduced by GenAI: a process that is faster, denser, more iterative, and less fully controlled, in which meaning is not only projected by the author but also amplified and displaced by statistically learned visual associations.

At this point, however, a clarification becomes necessary. The language of architectural semantics cannot be transferred uncritically from classical semiotics to generative image models. If semiotic accounts remain useful to explain how architecture becomes legible through culturally coded attributes, diffusion models do not “understand” meaning as human interpreters do. They do not decode signs, nor do they reason tectonically. They generate statistically plausible image formations based on learned visual correlations. In this context, semantics is not an intentional communicative act performed by the machine, but an emergent effect of probabilistic patterning. The article therefore moves from a classical account of architectural legibility toward a revised concept of statistical semantics more adequate to latent-space image production.

Methodologically, the paper builds on two experimental trajectories. The first starts from a deformed everyday carton package and follows its transformation through a staged protocol: from

naturally meaningful object, to neutralized meta-object, to anticipatory architectural rendering, and finally to several branches of intensified semantization. The second begins with a canonical and already legible architectural object, the Barcelona Pavilion, and subjects it to controlled semantic displacement in order to test thresholds of recognition and defamiliarization. Across both cases, the aim is not simply to produce compelling images, but to examine under what conditions geometry begins to be read as architecture, how that reading can be intensified, and when it risks collapsing into excess.

A key revision introduced in this article concerns precisely this methodological dimension. Rather than presenting GenAI outputs as isolated final images, the paper treats them as sequential protocols of semantic transformation. This is particularly important in the case of intensified semantization, which is here analysed through three branches: typological, atmospheric, and affective. The typological branch investigates how the same source geometry can acquire increasingly hybrid or overdetermined architectural identities. The atmospheric branch examines how lighting, climate, and perceptual density alter architectural reading. The affective branch explores how AI-generated architecture may trigger uncanny, embodied, or emotionally charged responses beyond stable coding. This staged approach makes the process itself visible and, at the same time, allows the identification of failed branches—cases in which intensification turns into what this paper calls semantic inflation, that is, the overcoding of the image through cliché, literal symbolism, or narrative excess.

This methodological shift also helps respond to a broader critical concern. If GenAI no longer merely produces better images, but reorganizes the way architectural meaning is generated and negotiated, then the problem is not only representational but epistemic. Representation is no longer simply the final confirmation of a project, but a field in which multiple semantic possibilities are proposed, compared, selected, and sometimes

mistrusted. For this reason, the future of architectural representation cannot be reduced to improved realism or faster rendering. It must also be understood as a transformation in the relation between image, meaning, judgment, and design agency.

The article is structured accordingly. It first distinguishes rendering from semantization and outlines the three regimes of semantization. It then introduces a theoretical bridge between classical semiotics and probabilistic image formation. After that, it presents the methodological framework of sequential protocols, including the notion of semantic inflation as a risk internal to intensified semantization. The paper then develops two case studies: the staged semantization of the deformed carton object, and the defamiliarization of the Barcelona Pavilion. It concludes by reflecting on semantic drift, inflation, and the future of representation, arguing that GenAI should be approached not only as a visual accelerator, but as a semantic laboratory for architecture.

2. RENDERING, SEMANTIZATION, AND STATISTICAL SEMANTICS

2.1 RENDERING AS ANTICIPATORY SEMANTIZATION

Architectural rendering is commonly understood as a technique of visualization. It translates drawings, models, and design intentions into persuasive images through light, texture, atmosphere, and composition. Yet such a definition remains incomplete. Rendering does not merely make a project visible; it also makes it legible. It projects not only the appearance of a future object, but also the conditions through which that object may be recognized as architecture and as a particular kind of architecture. In this sense, rendering operates as a mechanism of anticipatory semantization.

This distinction matters because architecture is

never received as pure geometry. Even the most abstract form becomes legible only through a set of semantic cues: scale, materiality, use, tectonic plausibility, context, and inhabitation. A wall, for example, is never just a plane. It may signify enclosure, weight, age, craft, climate, domesticity, monumentality, or fragility depending on how it is articulated and perceived. The same geometric configuration may therefore support very different architectural readings according to the semantic layers through which it is presented (Eco, 1968; Jencks, 1969).

Rendering has long participated in this layering process. Through carefully staged imagery, architecture is endowed in advance with visual attributes that, in the built world, would normally emerge over time: weather, occupation, surface depth, atmospheric mood, and social context. In this regard, the contemporary render does not simply depict a future building. It anticipates a world in which that building already appears experientially coherent and culturally situated. The image condenses a future semantic state into a present visual artifact.

This anticipatory role helps explain why rendering became so central in architectural culture before the rise of GenAI. Its power lies not only in communication, but in persuasion. What makes the rendered image convincing is not just its realism, but its ability to project a complete field of recognizability. Rendering visually compresses the very attributes through which architecture may be read, desired, and judged (Eco, 1968; Venturi, Scott Brown, & Izenour, 1972).

At the same time, conventional rendering remains relatively controlled. Its semantic projection depends on authorial assignment: materials are selected, contexts composed, and light behavior simulated around a stable geometric scene. Even when speculative or atmospheric, the rendered image generally preserves a coherent and intentional semantic regime. This controlled condition will be important when distinguishing rendering from the more unstable operations of generative AI.

2.2 FROM CLASSICAL LEGIBILITY TO STATISTICAL SEMANTICS

The distinction between rendering and intensified semantization requires a theoretical bridge. Classical semiotic approaches remain useful because they explain how architecture becomes legible through culturally coded attributes. Architecture is not encountered as pure form, but as form already mediated by recognition, convention, association, and use. Built objects are read through historically sedimented cues that allow observers to identify enclosure, threshold, domesticity, institutionality, permanence, or fragility (Eco, 1968; Jencks, 1969; Venturi et al., 1972).

What matters here is not simply that architecture “means,” but how meaning becomes attached to form. A geometry becomes architectural not because it possesses architectural essence in isolation, but because it is endowed with attributes that make it readable as a building and, more specifically, as a particular kind of building. Scale markers, openings, surface articulation, material thickness, tectonic cues, ground conditions, and signs of inhabitation all participate in this process. Semantics, in this sense, is not an optional layer added after form; it is one of the conditions through which form becomes architecturally intelligible at all.

However, classical semiotics cannot be transferred uncritically to AI-generated imagery. Eco, Jencks, and Venturi were concerned with human interpretation and cultural readability. Their frameworks help explain why architectural form becomes legible, but they do not explain how diffusion models produce meaning-like effects. A conceptual adjustment is therefore necessary.

In generative image systems, semantics is not produced in the same way as in human signification. Diffusion models do not interpret signs, understand programme, or reason tectonically. They generate images by learning statistical relations across large image-text datasets and by producing outputs that maximize plausibility within that learned visual field. In this context, semantics should not be understood as machine intentionality. It should be understood as an emergent legi-

bility effect generated through probabilistic visual correlation.

This distinction helps explain both the power and the instability of AI-generated architectural imagery. On the one hand, diffusion models can produce semantically dense images from relatively weak inputs, inferring materiality, atmosphere, context, and typological cues with very limited explicit modelling. On the other hand, because these cues emerge from pattern correlation rather than disciplinary understanding, the output may remain only statistically coherent. Tectonic relations may be visually persuasive while structurally implausible; atmospheres may feel convincing while semantically overdetermined; typological cues may accumulate beyond any stable architectural logic.

This condition resonates with broader discussions of technical and algorithmic images. Hito Steyerl has argued that contemporary images increasingly circulate within operational systems in which visibility is shaped by technical infrastructures rather than by purely human spectatorship (Steyerl, 2009; Steyerl, 2017). Trevor Paglen likewise shows that machine-readable images do not simply extend human seeing, but reorganize it through computational patterning and opaque classificatory systems (Paglen, 2016). Read in this light, AI-generated architectural imagery is not merely a new illustration tool. It belongs to a wider transformation in which images are produced within algorithmic ecologies of seeing.

Walter Benjamin remains useful here in a more oblique sense. His account of technical reproducibility helps frame a historical shift in the status of the image: from singular artifact toward serial, repeatable, and operationally reconfigurable visual production (Benjamin, 1936/2008). In the context of generative AI, this seriality becomes variational by default. One geometry can yield multiple plausible semantic states, each legible, each unstable, and each slightly different from the others. This is one of the reasons why AI imagery cannot be reduced to rendering in the conventional sense. It does not simply depict one project more efficiently; it opens a field of semantic drift around the same formal source.

2.3 THREE REGIMES OF SEMANTIZATION

Within this framework, the paper distinguishes three regimes of semantization: natural, anticipatory, and intensified. The term semantization is used here not in a narrowly linguistic sense, but to describe the process through which a geometry acquires architectural meaning and becomes readable as architecture through culturally and perceptually activated cues.

Natural semantization refers to the meaning an object acquires through real-world existence: production, use, material identity, social reception, memory, and temporal sedimentation. In the case of everyday objects, natural semantization is inseparable from habit and recognition. A carton package, for instance, is immediately understood through its scale, material, printed surface, and cultural role as commodity. In architecture, natural semantization occurs when buildings acquire significance through inhabitation, weathering, maintenance, symbolic appropriation, and historical embedding. Meaning here is not projected in advance; it is sedimented over time.

Anticipatory semantization refers to the regime most closely associated with architectural representation, and especially with rendering. Here meaning is projected in advance. A geometry is given scale, materiality, atmosphere, context, and signs of use before it is built, allowing it to be read as architecture in a future-oriented yet visually present form. Anticipatory semantization therefore describes the operation through which architecture appears semantically complete before material realization. This is the specific power of rendering: it does not merely show what a building will look like, but how it may already be understood, inhabited, and culturally positioned.

Intensified semantization names the condition introduced by generative AI. If rendering anticipates meaning, GenAI accelerates and thickens that process. It does so in at least three ways. First, it operates at great speed, allowing multiple semantic states to be generated around the same source geometry. Second, it increases semantic density

by introducing typological, atmospheric, and material cues with much less explicit modelling than conventional rendering would require. Third, and most importantly, it introduces uncertainty. The semantics of the output are no longer fully controlled by the author, but are partly shaped by the latent statistical associations learned by the model. Meaning is therefore not only projected, but also amplified, displaced, and partially recombined through probabilistic image formation.

This is why intensified semantization should not be confused with simply “better rendering.” The difference is not merely one of efficiency or image quality. It is a change in the structure of semantic production. A conventional render tends to stabilize a plausible reading of a project. A generative image may instead proliferate plausible readings, multiply connotations, and push the same geometry toward different semantic drifts. What emerges is not just a more detailed image, but a denser and less fully governed field of meaning.

In the present paper, intensified semantization is analysed through three branches: typological, atmospheric, and affective. In the typological branch, the image becomes semantically denser by accumulating or hybridizing building-type signals. In the atmospheric branch, the semantic charge is intensified through climate, light, perceptual density, and experiential mood. In the affective branch, meaning exceeds stable coding and enters a more embodied register, producing uncanny, attractive, or disturbing responses that resonate with theories of image-based empathy and embodied simulation (Freedberg & Gallese, 2007; Pallasmaa, 1996). These branches do not replace one another; they coexist as different directions of semantic drift made especially visible by AI-based image generation.

This tripartite distinction also requires one final clarification. Intensification is not always productive. Under certain conditions, semantic density turns into semantic inflation: a situation in which the image becomes overcoded, cliché-driven, or narratively overloaded. Instead of deepening ambiguity or multiplying meaningful readings, the image saturates itself with explicit signals and

loses interpretive tension. Semantic inflation is therefore not treated here as a separate regime, but as an internal failure mode of intensified semantization.

The distinction between natural, anticipatory, and intensified semantization provides the conceptual basis for the methodological sections that follow. It makes it possible to describe not only how architecture is represented, but how meaning is progressively constructed, accelerated, and sometimes destabilized through image-making. Generative AI does not invent semantization from nothing; it intervenes in an existing representational lineage. What changes is the speed, density, and partial autonomy with which semantic cues are generated and reorganized.

3. METHODOLOGY: SEQUENTIAL PROTOCOLS OF SEMANTIZATION

The methodological objective of this paper is not to evaluate isolated AI images as finished artifacts, but to make the process of semantic transformation visible and analysable. For that reason, the experiments are organized not as single before-and-after comparisons, but as sequential protocols of semantization. This allows the study to examine how architectural meaning is progressively stabilized, intensified, displaced, or overcoded across a controlled series of transformations. It also makes it possible to distinguish, at least partially, between what is introduced by the author through prompts and source selection, and what emerges from the model through statistically patterned visual inference.

A clarification is necessary at the outset. The three regimes of semantization defined in the previous section—natural, anticipatory, and intensified—are conceptual categories. The Stage 0–Stage 3 sequence introduced below is not a fourth theoretical regime, but a methodological framework used to analyse and produce semantic transformations in the case studies. Stage 0 operationalizes the natural semantic condition of the source object; it does not add a new regime to the model.

Three regimes of semantization

From naturally acquired meaning to anticipatory projection and intensified semantic drift in GenAI




Figure 1. Three regimes of semantization: from naturally acquired meaning, to anticipatory semantic projection through rendering, to intensified and probabilistic semantic drift in GenAI.

3.1 SOURCE OBJECTS AND META-OBJECTS

The paper works with two different source conditions. The first is an everyday deformed carton package, selected precisely because it begins as a culturally familiar but non-architectural object. Its initial meaning is not architectural but commercial, material, and habitual: it is recognized through packaging conventions, product identity, and everyday use. The second is the Barcelona Pavilion, chosen as a canonical building whose recognizability is already highly stabilized within architectural culture. These two sources allow the paper to work in opposite directions: from object recognition toward architectural semantization, and from canonical architectural identity toward semantic destabilization.

In both cases, the crucial intermediary device is what this paper calls the meta-object. A meta-object is not yet a fully semantized architectural

image, but a controlled visual substrate that serves as support for subsequent semantic operations. It is the level at which contingent identity is reduced enough for semantic drift to become analysable.

In Case Study A, the meta-object is produced through explicit neutralization. The photographed carton is transformed into a white, de-branded object in order to suspend contingent identity markers such as color, product information, and consumer association. This operation does not erase meaning entirely, but it reduces surface identity so that geometry can be treated as a more isolated support for subsequent semantic loading.

In Case Study B, the meta-object is not produced by neutralizing a commodity object, but by selecting a canonically stabilized source image of the Barcelona Pavilion and treating that image as a controlled semantic substrate. Here the Pavilion already possesses a highly coherent architectural

identity. The methodological task is therefore not to suspend commercial recognition, but to hold geometry, layout, and overall compositional logic sufficiently stable so that semantic displacement can be observed primarily through tectonic, material, and atmospheric substitutions. In this case, the meta-object is the source image as controlled canonical support.

This distinction is important because it allows the two case studies to remain methodologically parallel without forcing them into false symmetry. In one case, the meta-object is created through visual neutralization; in the other, it is established through the selection of a semantically stable architectural image.

3.2 STAGE 0—STAGE 3 FRAMEWORK

The experiments are organized through a four-stage methodological framework.

Stage 0 — Natural semantization.

At this stage, meaning belongs to the object's existing condition in the world. The deformed carton package is recognized through use, scale, materiality, and cultural familiarity. Its semantics are already active, but they are non-architectural. In Case Study B, the equivalent condition is the canonical recognizability of the Barcelona Pavilion as an already culturally stabilized architectural artifact. Stage 0 therefore establishes that semantization is not invented by AI; it precedes AI as part of the social and material life of objects and buildings.

Stage 1 — Meta-object formation.

Here the source is converted into a controlled semantic substrate. In Case Study A, this occurs through neutralization: printed surfaces, color accents, and commercial signs are removed or suppressed while the source geometry is preserved. In Case Study B, Stage 1 consists in fixing a canonical source image of the Pavilion as the stable visual basis for subsequent recoding. In both cases, the purpose is the same: to establish a parent geometry and a sufficiently stable visual support for controlled transformation.

Stage 2 — Anticipatory semantization.

From Natural Semantization to Intensified Semantization

Methodological framework for the staged analysis of semantization across object recognition, rendering, and generative transformation




Figure 2. Sequential framework of semantization used in the paper, moving from natural object recognition to meta-object formation, anticipatory architectural rendering, and three branches of intensified semantization: typological, atmospheric, and affective. Semantic inflation is

At this stage, the meta-object is transformed into a plausible architectural render. Scale, access, materiality, tectonic cues, context, and atmosphere are introduced in a relatively stable and disciplined way. The image now operates like a conventional architectural visualization: it projects a coherent future reading of the object or building. This stage corresponds to anticipatory semantization in its most controlled form and functions as the baseline against which Stage 3 can be evaluated.

Stage 3 — Intensified semantization.

Once architectural legibility has been stabilized, GenAI is used not simply to refine the render, but to intensify and redirect its semantic regime. Meaning is amplified, thickened, and partially destabilized. Rather than producing one final ima-

ge, Stage 3 unfolds through multiple branches of semantic drift. This is the key methodological contribution of the paper: intensified semantization is treated not as a singular result, but as a reproducible protocol of differentiated semantic escalation.

3.3 THREE BRANCHES OF INTENSIFIED SEMANTIZATION

Stage 3 is

divided into three branches: typological, atmospheric, and affective. These branches do not represent separate ontologies of the image, but three distinct directions in which semantic density can be intensified once a source geometry has already been rendered as architecture.

In the typological branch, the image is pushed

toward increased architectural legibility through hybridization, aggregation, or overdetermination of building-type cues. The same source geometry may begin to read simultaneously as dwelling, civic artifact, infrastructural fragment, tower, shelter, or institutional object. What is intensified here is not merely detail, but the density and instability of typological reading.

In the atmospheric branch, semantic intensification operates through light, weather, perceptual depth, climatic conditions, and experiential mood. The building may remain largely similar in massing and tectonic logic, yet acquire a far more charged semantic presence through wetness, translucency, glow, shadow, environmental embedding, or perceptual suspension. Meaning is not intensified primarily through added parts, but through experiential thickening.

In the affective branch, the semantic regime shifts toward embodied or uncanny response. The image may remain architecturally plausible while beginning to evoke unease, attraction, repulsion, vulnerability, or strange empathy. Here semantics exceeds stable coding and enters a more bodily register, consistent with image theories that describe affective and embodied responses as central to visual experience (Freedberg & Gallese, 2007; Pallasmaa, 1996).

The reason for distinguishing these branches is methodological. They allow the paper to show that intensified semantization is not a single "AI style," but a field of differentiated semantic operations. The same meta-object can be intensified in several directions, each one producing a different kind of architectural drift.

3.4 PROMPTING, ITERATION, AND SELECTION CRITERIA

Within this framework, prompting is not treated as a purely technical input, but as a semantic operation. A prompt condenses architectural intention into verbal form: it may specify typology, atmosphere, tectonic logic, context, materiality, or affective tone. Prompts therefore function as compressed semantic propositions rather than as

mere software commands.

At the same time, the paper does not assume that prompts fully determine the image. One of its central methodological questions is precisely how far the semantic content of an output is authored, and how far it is supplied by the model's own learned tendencies. For this reason, the experiments proceed through iterative prompting rather than one-off command execution. Each stage is observed as a negotiation between human semantic intention and machine semantic tendency.

In practical terms, the workflow combines four operations:

- (1) source selection, which establishes the parent geometry;
- (2) meta-object formation, which stabilizes the source as controlled substrate;
- (3) semantic prompting, which directs the intended semantic drift; and
- (4) iterative selection and correction, through which branches are refined, redirected, or rejected.

Because generative AI systems are predictive and variational, the same prompt may yield different outputs across executions. For this reason, prompts are treated in this study as trace records of semantic intention rather than as deterministic recipes. What is methodologically relevant is therefore not strict one-to-one reproducibility, but the relation between source condition, prompt kernel, iterative correction, selection criteria, and retained output.

To make this process more explicit, Table 1 summarizes the prompt structure used in Case Study A, including input condition, prompt kernel, negative constraints, selection criteria, and corresponding output figures.

3.5 SEMANTIC INFLATION AS FAILURE MODE

Not every intensified output is productive. Under certain conditions, semantic intensification turns into what this paper calls semantic inflation. This occurs when the AI image becomes overcoded, literal, or cliché-driven, and therefore loses the semantic tension that made intensification intere-

sting in the first place. Instead of deepening ambiguity or multiplying plausible readings, the image saturates itself with explicit signals: symbolic excess, decorative narrative additions, redundant atmosphere, or typological overstatement.

Semantic inflation is therefore not treated as a separate stage, but as a risk internal to Stage 3. It is the failure mode of intensified semantization. Methodologically, this distinction is important because it shows that AI does not automatically produce meaningful excess. Intensification requires discrimination. Some branches remain fertile because they preserve ambiguity, semantic density, and architectural plausibility. Others collapse into inflation because they replace complex drift with over-explicit coding.

This is one of the reasons why the paper emphasizes sequential image protocols rather than isolated outputs. When the process is visible, intensification and inflation can be distinguished analytically rather than judged only by intuition. The method can then be evaluated not only by the quality of the final images, but by its capacity to reveal how geometry acquires, shifts, and sometimes loses architectural meaning.

Note. Prompts are reported in abridged form and are treated as operative semantic traces rather than deterministic recipes. Because generative AI systems are predictive and variational, repeated execution of the same prompt may produce different outputs.

The following case study applies this framework to the staged transformation of the deformed carton package. As shown in Table 1, the sequence is not based on isolated prompts, but on iterative semantic operations organized around the same parent geometry.

4. CASE STUDY A — FROM OBJECT RECOGNITION TO INTENSIFIED SEMANTIZATION

The first case study examines how a culturally familiar but non-architectural object can progressively acquire architectural meaning through a controlled sequence of semantic transformations. The chosen source is a deformed carton package.

Stage	Input condition	Prompt kernel	Prompt used (abridged)	Negative constraints / iterative correction	Selection criterion	Output figure
Stage 1	Source photograph of deformed carton package	Neutralize contingent identity while preserving parent geometry	Remove brand, color, and packaging graphics; preserve the same viewpoint, deformation, and straw geometry; convert the object into a white neutralized meta-object	Same viewpoint; same deformation; same straw length and inclination; no geometric simplification beyond surface neutralization	Maximum preservation of parent geometry with minimum residual commercial coding	Fig. 4
Stage 2	Neutralized meta-object	Stabilize architectural legibility through anticipatory rendering	Transform the neutralized object into a plausible architectural render; preserve geometry and viewpoint; introduce scale, access, tectonic cues, and restrained materiality; recode the straw as an architectural equivalent	Controlled realism: plausible materiality; no semantic excess; no fantasy or radical drift yet	Most convincing baseline architectural reading before Stage 3 branching	Fig. 5
Stage 3A.2	Neutralized meta-object / architectural render	Typological aggregation	Preserve the original object and viewpoint; reinterpret it as a typologically intensified hybrid building; add radical yet plausible volumetric accretions, stacked fragments, misaligned roof forms, and conflicting domestic or civic cues	Avoid explicit symbols; avoid scenic storytelling; avoid gothic or religious clichés	Increase typological density while maintaining overall coherence and recognizability	Fig. 6
Stage 3A.3	Typologically intensified image	Material and tectonic consolidation	Refine the intensified aggregate so that the original parent geometry remains legible; reduce generic residential detailing; introduce believable architectural materiality; muted concrete, pale mineral render, matte metal, translucent blocks, dark frames, fibre-cement or weathered cladding	Avoid generic real-estate aesthetics; avoid glossy luxury finishes; avoid decorative refinement; avoid excessive balconies	Parent geometry remains visible while material and tectonic plausibility increase	Fig. 6
Stage 3A.4	Consolidated typological aggregate	Situated typological intensification	Preserve the aggregate, viewpoint, and parent geometry; situate the building in a believable world; maintain semantic density and typological ambiguity while strengthening environmental coherence	Avoid loss of source geometry; avoid cliché contextualization	Read as built architecture rather than model, while retaining hybrid typological overdetermination	Fig. 6
Stage 3B.2	Neutralized meta-object	Atmospheric transformation of the object	Preserve the same geometry, silhouette, and camera viewpoint; keep the background neutral; transform the object itself through soft translucency, internal luminous depth, condensation, and partial dematerialization; atmosphere must belong to the object rather than to the scenery	No landscape; no foggy backdrop; no typological additions; no decorative patterns; no sci-fi frameworks; no glossy plastic appearance	Atmosphere is internalized into the object's material body rather than added externally	Fig. 7
Stage 3B.3	Atmospheric object	Atmospheric spatial proliferation	Preserve the same camera viewpoint and atmospheric-material logic; spatially expand the image into a chamber-like interior formed by proliferated versions of the same faceted source logic; overlapping, intersecting, or fused glacial masses replace isolated repetition	Less repeated objects; more fused chamber; no installation logic; no separate sculptures; no fantasy cave cliché	The object becomes the generator of an immersive atmospheric field rather than remaining a singular artifact	Fig. 7
Stage 3B.4	Atmospheric spatial proliferation	Situated atmospheric intensification	Preserve parent geometry and icy-translucent material logic; place the architecture in a believable real setting; intensify dramatic light, chiaroscuro, reflective ground, and spatial depth while keeping atmosphere intrinsic to the architecture	Avoid fantasy cave style; avoid generic sci-fi lighting; avoid decorative effects; preserve family resemblance to the source geometry	The proliferated field becomes a radical but plausible atmospheric chamber	Fig. 7
Stage 3C.2	Neutralized meta-object	Tactile-organic object	Preserve the same geometry, silhouette, and viewpoint; transform the object through membrane-like skin, slight humidity, subtle translucency, tactile tension, fine folds, and an organic surface close to flesh without literal anatomy	Avoid gore; avoid wounds; avoid teeth; avoid blood; avoid creature-like horror; avoid excessive geometric change	Clear bodily displacement while preserving parent geometry and objecthood	Fig. 8
Stage 3C.3	Tactile-organic object	Uncanny visceral architecture-object	Preserve the same parent geometry; push the object toward affective architectural strangeness; increase humidity, folds, cavities, and flesh-adjacent materiality while retaining architectural reading	Avoid explicit anatomy; avoid monsters; avoid losing architectural legibility	Strong affective intensification at the level of the artifact before full spatialization	Fig. 8
Stage 3C.4	Uncanny visceral architecture-object	Bodily architectural chamber	Preserve the parent geometry, silhouette logic, and folded deformed-carton structure; do not make a generic fleshy cave; expand the object into an inhabitable chamber where the same folds, tensions, and compressed surfaces become walls, thresholds, ceilings, and interior depth	Avoid generic cave imagery; avoid blood, teeth, tongue, organs, or horror aesthetics; preserve folded geometry	Spatial enclosure emerges from the same source logic, producing affective architecture rather than literal anatomy	Fig. 8
Stage 3 failure mode	Any Stage 3 branch	Semantic inflation	Overcoding of semantic cues beyond productive ambiguity	Literal symbols, decorative overstatement, generic atmospheric mood, sensational bodily imagery	Reject outputs that appear semantically rich but flatten architectural complexity	Fig. 9

Table 1. PromptmatrixforCaseStudyA.The tables summarizes the input conditions, prompt kernels, negative constraints, iterative corrections, selection criteria, and corresponding output figures used in the staged transformation from source object to intensified semantization.

Its interest lies precisely in the fact that it already possesses strong meaning, but not architectural meaning. It is recognizable as packaging through scale, material, use, and commodity identity. The experiment therefore does not begin from neutral geometry, but from an object whose natural semantization is already active.

The purpose of the case study is not simply to show that AI can transform an object into something that looks like architecture. Its objective is more specific: to test how the same parent geometry can be carried from everyday recognition to architectural legibility, and from there into different branches of intensified semantization. The sequence therefore follows the full methodological structure introduced in the previous section, from Stage 0 to Stage 3, including both productive and failed outcomes. As summarized in Table 1, the case study is documented through prompt kernels, negative constraints, iterative corrections, and selection criteria rather than through isolated final images.

4.1 STAGE 0 — NATURAL SEMANTIZATION OF THE EVERYDAY OBJECT

At Stage 0, the carton package is still fully embedded in the semantic regime of everyday use. Its meaning depends on packaging conventions, printed surface identity, material thinness, deformability, and habitual recognition. The object is not neutral at all. It is already culturally loaded, but its load is commercial, domestic, and functional rather than architectural. This stage is important because it shows that the experiment does not begin from abstract geometry in a vacuum. It begins from an object whose meaning has already been formed through natural semantization.

The source photograph is therefore not just a raw visual input. It records an object already situated within the lifeworld of use, consumption, and immediate recognition. The reader is not meant to see it as proto-architecture, but precisely as a meaningful non-architectural artifact. Only by starting from this condition does the later shift toward architecture become analytically visible.



Figure 3. Stage 0 — Natural semantization of the source object. The deformed everyday carton package is read through use, material identity, scale, and commercial recognizability before any architectural transformation is introduced.

4.2 STAGE 1 — META-OBJECT FORMATION

The first methodological transformation consists in converting the source into a meta-object. This is achieved by neutralizing contingent identity markers while preserving the parent geometry. Brand information, color, and printed packaging signals are removed or suppressed, producing a white object that no longer reads primarily as a commercial package. The result is not yet architecture, but neither is it simply packaging. It becomes a transferable semantic support.

This step is essential because it separates inherited object identity from subsequent architectural

projection. Without it, the experiment would risk confusing the semantic residue of the package with the new semantics introduced by rendering and GenAI. Neutralization therefore functions as controlled abstraction. It does not erase all meaning, but it suspends the most specific surface codes of commodity recognition.

The meta-object also marks the point at which geometry becomes analytically visible as parent structure. The twisted mass, compressed middle, angled upper plane, and protruding straw-like element remain present, but they are no longer



Figure 4. Stage 1 — Neutralized meta-object. The source object is stripped of contingent identity markers while preserving its deformed geometry, allowing it to function as a transferable support for subsequent semantic transformation

anchored to a branded product. The object is now available for semantic loading in architectural terms.

The operative prompt at this stage was not conceived as an image-generation prompt in the full imaginative sense, but as a geometric preservation prompt: same viewpoint, same deformation, same straw length and inclination, with the removal of surface identity as the primary transformation. This is important because later stages depend on the fidelity of this parent geometry.

4.3 STAGE 2 — ANTICIPATORY SEMANTIZATION

Once the meta-object has been established, it is transformed into a plausible architectural rendering. This is the stage at which anticipatory semantization becomes operative. Scale, access, materiality, tectonic articulation, and environmental framing are introduced so that the object can be read as architecture in a stable and disciplined way. The goal is not yet semantic excess, but architectural plausibility.

What emerges at this stage resembles a conventional render in the strongest sense of the term. The image no longer presents a neutral object on a studio background, but a coherent architectural artifact placed within a believable world. The transformation depends on the controlled addition of cues that stabilize reading: surfaces become architectural materials, the protruding element is recoded tectonically, the lower part acquires ground relation, and the whole begins to function as a small but plausible architectural structure.

This stage is methodologically crucial because it provides the baseline against which intensified semantization can later be measured. Without a stable anticipatory image, Stage 3 would risk becoming a free sequence of stylistic mutations without analytical control. Stage 2 therefore establishes a disciplined architectural reading before semantic drift begins.

Here the prompts aimed to preserve geometry

and viewpoint while introducing restrained architectural materiality, legible access, tectonic cues, and an equivalent architectural reading for the straw-like appendage. The selected output was the one that most convincingly stabilized architectural legibility without yet entering typological, atmospheric, or affective excess.



Figure 5. Stage 2 — Anticipatory semantization. The neutralized object is transformed into a plausible architectural render through the controlled introduction of materiality, access, enclosure, context, and atmospheric coherence.

4.4 STAGE 3A — TYPOLOGICAL INTENSIFICATION

Once architectural legibility has been stabilized, the experiment moves into intensified semantization. In the typological branch, the goal is no longer to make the object simply readable as architecture, but to push it toward a denser and more unstable typological condition. The image is intensified not by adding decorative detail or narrative symbolism, but by multiplying and hybridizing the kinds of architecture the object may be read as.

For analytical clarity, this branch is organized as a four-step sequence: (1) neutralized source meta-object; (2) typological aggregation in a neutral field; (3) material and tectonic consolidation; and (4) situated typological intensification.

The first intensified image retained the neutral studio condition while amplifying the source object into an aggregated architectural construct. In this phase, the geometry became more densely articulated: balconies, terraces, stairs, projections, secondary roof fragments, and cantilevered volumes accumulated around the parent body. Although still presented against a neutral background, the image already showed that the source object could support a semantically denser architectural reading without dissolving into pure fantasy. The object no longer appeared as a singular small structure, but as a condensed vertical aggregate.

A second step introduced a more explicit material and tectonic consolidation while preserving the studio-like field. Concrete, metal roofing, glazing, and surface differentiation began to transform the neutral aggregation into something closer to a built architectural artifact, while still keeping the process analytically visible. At this stage, the image occupies an intermediate condition between model and architecture: no longer a neutral meta-object, but not yet fully situated in a believable world. This intermediate state is methodologically important because it reveals that typological intensification does not emerge all at once. It develops through successive recordings in which the parent geometry remains visible while architecture progressively thickens around it.

Only in a third step was this intensified aggregate

inserted into a plausible built context and given a more complete environmental and atmospheric framing. The object acquired ground, weather, background structures, and subdued signs of occupation, allowing the image to read less as a model and more as a documented building. The oblique appendage inherited from the original carton was also progressively recoded. What in the source object read as a domestic trace of packaging became, through iteration, a tectonic appendage integrated into the upper body of the building.

The prompts used in this branch revolved around a stable semantic kernel: preserve the parent geometry and viewpoint while intensifying typological reading through aggregation, hybridization, ritual or civic ambiguity, and controlled tectonic recoding. Iterative refinements then reduced generic residential detailing and introduced more plausible architectural materiality. The selected outputs were those in which typological density increased while the parent geometry remained legible and the image did not collapse into fantasy or decorative cliché.

The intensified image produced by this branch does not stabilize into a single building type. Instead, it accumulates partially incompatible cues: dwelling, vertical extension, infrastructural appendage, stacked domestic fragments, and almost institutional massing. For this reason, it is best understood as a case of typological overtermination. The object is not simply more architectural than in Stage 2; it is semantically denser because it can no longer be reduced to one clear typological reading.

Taken together, the typological branch shows that intensified semantization does not simply make the object more architectural. It makes it typologically denser and less reducible to a single reading, while preserving the recognizability of the parent geometry beneath a controlled process of semantic aggregation.



Figure 6. Stage 3A—Typological intensification as sequential semantic drift. From left to right: neutralized source meta-object; typological aggregation in a neutral field; material and tectonic consolidation of the intensified object; and situated typological intensification in a plausible built context. The sequence shows how the same parent geometry can be progressively recoded into a semantically denser and more overdetermined architectural aggregate.

4.5 STAGE 3B — ATMOSPHERIC INTENSIFICATION

If the typological branch intensified meaning through aggregation and overdetermined architectural reading, the atmospheric branch operates differently. Here the goal is not primarily to multiply building-type cues, but to transform the same parent geometry through perceptual, climatic, and material intensification. Atmosphere ceases to function as a secondary setting and becomes instead an intrinsic property of the architectural object itself.

For analytical clarity, this branch is organized as a four-step sequence: (1) neutralized source meta-object; (2) atmospheric transformation of the object; (3) atmospheric spatial proliferation; and (4) situated atmospheric intensification.

Early attempts at atmospheric transformation ten-

ded to place the source object within fog, rain, or scenic mood, while leaving the object itself largely unchanged. Such results were discarded because they atmosphericized the environment rather than the object. For the present study, atmospheric intensification was therefore defined more strictly: the atmosphere had to migrate from context into the material body of the object. The object itself had to become humid, translucent, luminous, softened at the edges, or partially dematerialized.

In the first transformation, the object acquired semi-translucency, embedded luminosity, and a softened material presence. The geometry remained stable, but the semantic reading shifted toward a climatic or phenomenological artifact. In the next step, that same atmospheric logic expanded beyond the isolated object into a chamber-like field of related glacial or translucent forms. At this stage, the source geometry no longer functioned

only as a singular object, but as the generator of a surrounding environmental condition. A final situated image transformed this proliferation into a more architecturally legible chamber, with stronger chiaroscuro, deeper spatial recession, and a more explicit distinction between luminous mass, dark enclosing field, and reflective ground.

The prompts used in this branch progressively shifted from atmospheric transformation of the object itself, through luminous dematerialization and internalized translucency, toward spatial proliferation and finally situated atmospheric intensification. The key negative constraint remained consistent throughout: atmosphere had to be embedded in the architecture rather than added as scenic background. Iterative corrections were especially important in this branch, since many outputs initially drifted toward externalized fog, decorative cave imagery, repeated isolated



Figure 7. Stage 3B—Atmospheric intensification as sequential perceptual and material transformation. From left to right: neutralized source meta-object; atmospheric transformation of the object through translucency and embedded luminosity; atmospheric spatial proliferation into a chamber-like field of related glacial forms; and situated atmospheric intensification in a dramatically lit architectural interior. These sequences show how the same parent geometry can be re-coded through atmosphere rather than through typological aggregation.

objects, or generic sci-fi effects. Selected outputs were those in which atmosphere became materially intrinsic to the object and its derived space. What is intensified in this branch is therefore not type, but experiential density. The same source geometry becomes semantically charged through translucency, humidity, light diffusion, climatic embedding, and perceptual uncertainty. The image is not merely richer in atmosphere; it becomes architecturally meaningful through atmosphere.

Taken together, Stages 3B.1 to 3B.4 demonstrate that atmospheric intensification is not reducible to mood rendering. Rather than adding visual drama around an unchanged object, it progressively re-codes the object itself as a climatic and perceptual generator, eventually producing a situated architectural environment whose semantics derive

from embedded atmosphere rather than typological clarity.

4.6 STAGE 3C — AFFECTIVE INTENSIFICATION

If the typological branch intensified meaning through hybridized architectural reading, and the atmospheric branch through perceptual and climatic charge, the affective branch operates at another level. Here the semantic transformation concerns the capacity of the image to provoke embodied, uncanny, attractive, or disturbing responses beyond stable architectural coding. The question is no longer only what the object is read as, nor only how it is atmospherically perceived, but how it affects the viewer at a bodily and affective level.

For analytical clarity, this branch is organized as a

four-step sequence: (1) neutralized object; (2) tactile-organic object; (3) uncanny visceral architecture-object; and (4) bodily architectural chamber.

The first affective shift occurs when the object acquires a tactile-organic quality. Its surfaces begin to read less as neutral model matter and more as something soft, humid, folded, and skin-like. The geometry is still clearly preserved, but the viewer's relation to it changes. The object is no longer simply observed; it begins to solicit a bodily response. It appears touchable in an uneasy way, as if its matter had become too close to skin, membrane, or flesh.

A second transformation pushes this tendency further. The object becomes an uncanny visceral architecture-object: still singular, still recognizably derived from the same parent geometry, but materially and affectively displaced toward a more



Figure 8. Stage 3C—Affective intensification as sequential bodily re-coding. From left to right: neutralized object; tactile-organic object; uncanny visceral architecture-object; and bodily architectural chamber. These sequences show how the same parent geometry can be progressively transformed from abstract artifact into an affectively charged spatial enclosure through skin-like materiality, cavity, and embodied architectural proximity.

disturbing bodily register. Folds deepen, cavities intensify, and the overall image acquires a stronger resemblance to moist, vulnerable, living matter. Only in the final step does the object become a bodily architectural chamber. Here the same affective and material logic is spatialized. The geometry no longer functions only as a bounded artifact, but as the generator of an interior condition. The result evokes bodily interiority without becoming literal anatomy.

The prompts used in this branch consistently attempted to preserve the parent geometry while recoding it through membrane-like skin, subtle translucency, humidity, folds, cavities, and flesh-adjacent materiality. Strong negative constraints were essential: no gore, no blood, no explicit organs, no literal teeth or tongue, no monstrous or horror-driven imagery. The most important iterative correction at the final step was metho-

logical rather than stylistic: the result had to remain architecture derived from the original folded object, not collapse into a generic fleshy cave. The selected outputs were those that maintained architectural legibility while producing embodied affective charge.

This branch shows that semantic intensification can occur not only through typological or atmospheric means, but also through bodily and affective recoding. The same parent geometry becomes semantically denser because it begins to operate at the threshold between architectural recognition and corporeal unease.

Taken together, the affective branch confirms that intensified semantization can exceed both conventional legibility and atmospheric modulation. The same parent geometry is pushed toward a materially embodied threshold at which architec-

ture begins to resemble a body, and space begins to feel affectively operative without collapsing into literal anatomy.

4.7 A NOTE ON SEMANTIC INFLATION

Not every trajectory of intensified semantization produced a meaningful result. Alongside the productive branches described above, the experiments also generated images in which semantic density ceased to be generative and became excessive, literal, or reductive. This paper describes such outcomes as semantic inflation.

Semantic inflation does not refer simply to visual failure or poor image quality. Inflated images may be highly polished, persuasive, and atmospherically rich. The problem lies elsewhere: meaning becomes overcoded. Instead of preserving ambiguity, tension, or layered interpretability, the

image begins to state too much, too directly, and too quickly. Symbolic signs become literal, atmospheric cues become formulaic, and typological suggestions harden into cliché. The result is an image that appears semantically full but is, in fact, semantically flattened.

Within the present case study, semantic inflation serves as an internal critical filter. It names those branches in which typological aggregation turned into decorative excess, atmospheric charge into scenic formula, or affective intensity into sensational literalism. The concept is methodologically important because it shows that GenAI does not

automatically generate relevant architectural meaning. It can intensify, but it can also trivialize. The value of the protocol lies precisely in making that difference visible.

Taken together, Stages 0 to 3 show a complete semantic trajectory. The everyday object begins as a naturally meaningful artifact of consumption, is converted into a meta-object, stabilized as architecture through anticipatory rendering, and finally intensified into typological, atmospheric, or affective branches of semantic drift. What GenAI adds at this point is not simply realism or formal novelty. It enables the same geometry to be pushed into denser and more unstable fields of architectural legibility, while also making visible the threshold at which intensification turns into inflation.

5. CASE STUDY B — DEFAMILIARIZING A CANONICAL BUILDING

The second case study reverses the direction of the first. Instead of starting from a non-architectural object and progressively loading it with architectural semantics, it begins from a canonical building whose architectural identity is already highly stabilized. The aim is therefore not semantization in the basic sense, but defamiliarization: a controlled displacement of the semantic regime of a building that remains recognizably itself while progressively drifting away from its conventional reading.

The chosen case is the Barcelona Pavilion, a particularly suitable object for this experiment because of the clarity and stability of its disciplinary legibility. Its flat roof, open plan, reflecting pool, thin supports, polished stone planes, and glass enclosure have become canonical markers within modern architectural culture. It is not merely a building, but a codified architectural image. This makes it an ideal test case for examining how far semantic drift can proceed before recognizability begins to weaken or mutate.

In contrast to Case Study A, the purpose here is not to produce a full branching protocol. Given the time constraints and the already strong re-

cognizability of the source object, the experiment is organized as a more concentrated sequence of controlled transformations. The central methodological principle, however, remains the same: geometry, layout, and overall compositional logic are held as stable as possible, while semantic displacement is produced primarily through tectonic, material, and atmospheric recoding.

5.1 CANONICAL RECOGNIZABILITY AND CONTROLLED SOURCE STABILITY

The experiment begins from a photograph of the Barcelona Pavilion taken from a stable and highly legible viewpoint. This source condition already carries a strong semantic charge. The Pavilion is immediately recognizable through its horizontal roof plane, reflecting pool, free-standing walls, glass enclosure, and the refined spatial economy through which modern architecture is articulated as openness, continuity, and material precision. In this sense, the initial image already operates as a naturally and culturally semantized architectural artifact.

In methodological terms, the source image functions here as the equivalent of the meta-object. Unlike the deformed carton package, the Pavilion does not need to be neutralized in order to become analytically available. Its value lies precisely in its stability. The parent geometry, spatial layout, and viewpoint are held sufficiently constant so that semantic change can be observed primarily through substitution and recoding rather than through geometric invention.

This means that the experiment does not attempt to redesign the Pavilion in any comprehensive sense. Instead, it subjects a stable and canonical architectural image to a sequence of controlled semantic displacements. The purpose is to test what happens when the Pavilion remains compositionally itself while its tectonic regime begins to shift.



Figure 9. Semantic inflation as a failed branch of intensified semantization. The image is visually persuasive but semantically overcoded: symbolic, atmospheric, or typological cues become too explicit, reducing ambiguity and weakening architectural complexity.



Figure 10. Canonical source image of the Barcelona Pavilion used for the defamiliarization experiment. The building is treated as a semantically stable architectural object whose identity depends on the coherence of geometry, tectonics, materiality, and atmosphere.

5.2 ANTICIPATORY SEMANTIZATION AND BASELINE STABILIZATION

The first transformation of the sequence remains close to the logic of anticipatory semantization. Starting from the source photograph, the image is re-rendered as a more resolved architectural visualization. Light, surface continuity, reflections, and material precision are intensified, but the Pavilion remains fully readable as itself. This step is important because it establishes a baseline condition against which later displacements can be measured.

At this stage, the experiment confirms that the Pavilion is not only resilient as geometry, but also semantically coherent as a modern artifact. The rendering does not fundamentally transform the building; it consolidates its identity. Stone, glass, steel, water, and landscape remain within the same disciplinary field of meaning. In this sense, anticipatory semantization here works conservatively. It reinforces recognition rather than destabi-

lizing it.

This baseline is necessary because the following images become meaningful only through contrast with it. Without a stable anticipatory version of the Pavilion, later transformations would risk being read as arbitrary stylistic exercises. The first rendered image therefore performs the same methodological role that Stage 2 performed in Case Study A: it secures an initial architectural reading before semantic drift begins.

5.3 LIGHT TECTONIC DISPLACEMENT: VIRTUALITY AND DEMATERIALIZATION

The first substantial displacement occurs when the Pavilion is reimagined through a tectonic regime of luminous lines and virtual enclosure. Here the canonical layout remains recognizable, but the architectural fabric appears to dissolve into a quasi-augmented or quasi-virtual construct. Structural edges, frames, and surfaces are translated into luminous outlines, producing an image that hovers between architectural rendering, digital model, and mixed-reality apparition.

What makes this transformation analytically useful is that it does not destroy the Pavilion's identity. On the contrary, it shows how far recognizability can persist even when tectonic substance is radically lightened. The building still reads as the Pavilion because its proportions, organization, and spatial syntax remain largely intact. Yet its semantic regime has clearly shifted. It is no longer the Miesian Pavilion of stone, steel, and glass, but a pavilion of projection, latency, and digital suspension.

This image therefore occupies an intermediate zone between anticipatory semantization and intensified semantization. It intensifies the Pavilion not through typological deviation, but through ontological lightening. The building begins to read as less permanent, less materially grounded, and more dependent on image-based presence. In this sense, it marks an early threshold of defamiliarization: the object remains canonically legible, but its tectonic confidence has already been displaced.

5.4 STRONG DISPLACEMENT: AFFECTIVE TECTONIC RECODING

The next transformations push the Pavilion into a more radical semantic field. In one image, the tectonic elements are replaced by flesh-like material. In another, they are replaced by hair-like matter. In both cases, the underlying layout and basic compositional order remain close enough to sustain recognition, but the semantic register changes dramatically. The Pavilion is no longer stabilized by the modern vocabulary of precision, refinement, and mineral permanence. It becomes soft, bodily, and uncannily alive.

These substitutions are not simply surreal effects. Their methodological relevance lies in the tension they produce between modern spatial order and radically displaced tectonic materiality. The familiar flat roof, wall planes, and linear enclosure remain, but they are inhabited by matter that appears too soft, too organic, too unstable, or too intimate for the disciplinary expectations of canonical modernism. The result is a conflict between spatial syntax and material signification.

In the flesh-like version, the Pavilion begins to read as a bodily envelope without losing its architectural organization. In the hair-like version, it becomes temporally and materially unstable in another way: the building appears overgrown, shaggy, almost animal, while still preserving its basic diagram. These images show that affective intensification can be applied not only to newly semantized objects, as in Case Study A, but also to already canonical architecture. A building may remain recognizably itself while entering an entirely different affective regime.

At this stage, semantic drift no longer concerns only atmosphere or tectonic refinement. It begins to affect the very ontology of the building's matter. The Pavilion still reads as the Barcelona Pavilion, but it does so under stress. Recognition survives, yet its semantic ground has been transformed.



Figure 11. Sequential defamiliarization of the Barcelona Pavilion. From left to right: source condition; anticipatory re-rendering within the same semantic regime; virtualized or dematerialized tectonic displacement; flesh-like affective recoding; and hair-like affective recoding. These sequences show how semantic drift can progressively destabilize a canonical architectural image while preserving residual recognizability

5.5 TEMPORARY AND REVERSIBLE TECTONICS

A further image introduces a different kind of displacement. Instead of recoding the Pavilion through bodily materiality, it replaces its tectonic certainty with a lighter, more provisional regime of scaffolding, colored panels, movable elements, and temporary assembly. This image is less uncanny than the flesh and hair substitutions, but methodologically important for another reason: it tests what happens when the Pavilion is displaced not toward the body, but toward ephemerality and reversibility.

Here the semantic shift is architectural rather than

affective in a narrow sense. The Pavilion begins to resemble an installation, an experimental enclosure, or a temporally contingent event-structure rather than a finished modern masterpiece. Its identity is still legible, but its aura of permanence is weakened. The modern canon is recoded as something lighter, less final, and more open to alteration.

Because the available image for this variation departs more substantially from the main viewpoint, it is treated here as a supplementary transformation rather than as part of the core comparative sequence. It remains useful, however, because it broadens the experiment beyond a binary op-

position between canonical stability and affective distortion. It shows that semantic drift may also take place through changes in duration, assembly logic, and tectonic reversibility.

5.6 RECOGNITION THRESHOLDS AND SEMANTIC DRIFT

Taken together, the sequence demonstrates that canonical architecture can sustain a surprising degree of semantic displacement before becoming unreadable. The Barcelona Pavilion does not lose its identity the moment its original material palette is altered. Recognition persists through geometry, composition, and spatial syntax even when tectonic meaning changes substantially.

At the same time, the experiment also reveals that recognizability is not unlimited. It depends on a delicate relation between form, matter, and disciplinary expectation. When tectonic substitutions become too literal, too affectively dominant, or too compositionally disruptive, the Pavilion risks ceasing to read as itself and becoming only a carrier of visual effect. The experiment is therefore valuable not because it proves infinite flexibility, but because it makes the thresholds of identity visible.

This case study confirms a central argument of the paper: semantization is not a secondary overlay applied to stable architecture from the outside. It is constitutive of how architecture is recognized in the first place. By displacing the semantic regime of a canonical building while holding its compositional logic relatively constant, the experiment shows that architectural identity depends on more than geometry alone. It depends on the semantic coherence through which geometry becomes culturally and disciplinarily legible.

5.7 RELATION TO CASE STUDY A

Read together, the two case studies form an inverse pair. Case Study A begins with a meaningful but non-architectural object and tracks its progressive acquisition of architectural meaning through anticipatory and intensified semantization. Case Study B begins with a canonically architectural object and tests how far that meaning can be displaced before identity becomes unstable. The

first case is about semantic acquisition; the second about semantic defamiliarization. Their methodological asymmetry is deliberate. Case Study A is used to build and verify the staged protocol in detail, including meta-object formation, branching logic, and prompt traceability. Case Study B is more compact and comparative. Its value lies not in reproducing the full protocol again, but in showing that the same conceptual framework also applies in reverse: architecture may not only be produced through semantization, but also estranged through it. This is why the second case does not require the same prompt matrix as the first. Its argument does not depend on exact prompt genealogy, but on the controlled visual comparison of semantically displaced images derived from a stable and recognizable architectural source. The experiment remains methodologically valid because its operative variable is clear: tectonic and material substitution under conditions of compositional continuity.

6. DISCUSSION: SEMANTIC DRIFT, INFLATION, AND THE FUTURE OF REPRESENTATION

The two case studies suggest that generative AI should not be understood merely as a faster or more accessible means of producing architectural images. Its more consequential effect lies in transforming the representational image into a site of semantic negotiation. Across both trajectories—from non-architectural object to intensified architectural legibility, and from canonical building to semantic defamiliarization—the image ceases to function as a stable endpoint of visualization. Instead, it becomes a field in which geometry, materiality, atmosphere, typology, and affect are progressively recalibrated, intensified, or displaced. What emerges from these experiments is not simply a new image aesthetic, but a different condition for representation itself. The architectural image is no longer asked only to depict, persuade, or anticipate. It is increasingly capable of testing thresholds: when an object begins to read as architecture, when architecture begins to drift away

from its own canonical identity, and when semantic density remains productive rather than collapsing into overstatement. In this sense, generative AI does not simply accelerate representational practice; it destabilizes its conventional epistemic role.

6.1 FROM DEPICTION TO SEMANTIC NEGOTIATION

Architectural representation has long been understood as a medium of projection. Drawings, models, collages, and renders do not passively record architecture; they construct possible readings of it. Yet in most conventional workflows, the image still tends to function as a relatively stable vehicle of communication. It depicts a project, anticipates a built result, or conveys design intention with greater or lesser degrees of atmosphere and persuasion.

The experiments developed in this paper suggest a more unstable condition. In generative workflows, the image is not only representational but negotiative. It does not merely show a predefined architectural object; it actively participates in determining what that object may mean. Because AI-generated imagery operates through statistical plausibility rather than strict geometric fidelity, it produces a zone in which architectural identity becomes more fluid, more iterative, and more vulnerable to semantic drift. A single source geometry can be rendered as stable architecture, typologically overloaded aggregate, atmospheric chamber, bodily enclosure, or semantically inflated cliché. The image therefore becomes less an endpoint and more a testing ground.

This shift has methodological consequences. The image can no longer be treated only as evidence of a design intention already resolved elsewhere. It becomes part of the process through which meaning is explored, displaced, and judged. Representation moves from depiction toward negotiation: not a neutral translation of design into image, but an iterative confrontation between source geometry, prompt intention, machine tendency, and disciplinary reading.

6.2 PRODUCTIVE INTENSIFICATION AND ITS LIMITS

One of the most important results of the study is that semantic intensification is not automatically productive. The same generative logic that enables dense typological, atmospheric, and affective transformations also generates failure modes. This is why the notion of semantic inflation proved necessary. It designates those cases in which semantic density becomes excessive, literal, or overcoded, and therefore loses analytical value.

This failure is not marginal; it is intrinsic to the medium. Because generative systems are highly effective at producing convincing semantic signals, they can easily move from suggestive ambiguity to reductive explicitness. A branch that begins as productive typological drift may collapse into decorative symbolic overload. An atmospheric transformation may harden into generic scenic mood. An affective sequence may drift into sensational bodily cliché. In each case, the image appears richer while becoming less critically useful.

The point is not simply that some images are better than others. More fundamentally, these failures show that generative AI tends to amplify semantic content faster than disciplinary judgment can stabilize it. The image may become semantically persuasive before it becomes architecturally meaningful. This asymmetry is one of the key frictions of the current moment: the machine produces plausibility at great speed, while the architect must still decide whether that plausibility corresponds to anything critically or disciplinarily relevant.

For this reason, the value of the proposed method lies not only in generating successful images, but in making visible where and how failure occurs. Sequential protocols help distinguish fertile intensification from semantic inflation precisely because they do not isolate the final output from the process that produced it. They allow the image to be read not simply as result, but as a trace of semantic decision-making and drift.

6.3 THE FUTURE OF REPRESENTATION

The future of representation, in light of these experiments, is not defined merely by higher realism, faster workflows, or broader accessibility. Its more significant transformation lies in the changing status of the image itself.

In conventional rendering culture, the architectural image often functions as a stabilized promise: a persuasive anticipation of a future building. In generative image culture, by contrast, the image increasingly becomes provisional, comparative, and branch-based. It is less a final statement than a sampled field of possible semantic states. Representation moves away from singular depiction and toward sequences, variants, thresholds, and protocols. The important question is no longer only whether an image is convincing, but what semantic operations it performs and how it repositions the object it depicts.

This implies that the future of representation may be less about “better images” than about more explicit semantic experimentation. The architectural image becomes a medium for probing legibility, testing identity, and staging controlled semantic mutation. It can still serve projection, but it also becomes a tool for interrogating what counts as architecture, what stabilizes recognizability, and which semantic dimensions remain latent or unstable within a given geometry.

At the same time, this future is not free of risk. If representational practice becomes increasingly governed by semantic plausibility, there is a danger that images may outrun both tectonic verification and conceptual precision. Representation may become richer in connotation while becoming weaker in disciplinary grounding. The future of representation is therefore not simply expansive; it is also fragile. It demands new forms of critical reading able to distinguish between semantic experimentation, semantic inflation, and architectural knowledge.


6.4 VIRTUOSITY AFTER RENDERING

These transformations also affect the question of

<http://disegnarecon.univaq.it>

From Representation to Semantic Negotiation

Closing diagram for the shift from depiction to branching semantic experimentation and critical judgment



Representation no longer functions only as a final image; it becomes a protocol of semantic negotiation, drift, and judgment.

Figure 12. Representation as semantic negotiation. The diagram summarizes the shift from stable source geometry and anticipatory rendering toward intensified semantization, where typological, atmospheric, and affective branches must be critically evaluated against the risk of semantic inflation. The future of representation is framed not as improved depiction alone, but as a process of branching semantic experimentation and disciplinary judgment.

virtuosity. In conventional terms, architectural virtuosity in representation has often been linked to mastery of drawing, modelling, rendering software, composition, and visual post-production. Those skills do not disappear, but they are displaced. Generative AI lowers the technical threshold for producing persuasive imagery, while raising the importance of another set of capacities: selecting source geometries, structuring semantic sequences, formulating prompts, recognizing productive drift, rejecting inflation, and interpreting ambiguous outputs.

It would be misleading, however, to treat this shift triumphantly, as if the architect were now simply

reborn as curator of machine images. The experiments in this paper suggest a more unstable condition. The architect's role is not secured by the mere fact of choosing prompts or selecting outputs. On the contrary, the multiplication of semantically plausible images makes judgment more difficult, not less. The abundance of possible readings increases the burden of discrimination.

For that reason, virtuosity after rendering should not be reduced to prompt-writing or image selection alone. It lies more fundamentally in the capacity to construct meaningful protocols, compare branches of transformation, identify when an image remains architecturally relevant, and

understand when semantic richness is genuine and when it is only inflation. If virtuosity persists, it does so less as technical authorship of a final image than as critical control over a field of semantic possibilities.

This may be the most important consequence of the transition described in this paper. The future architect is not simply the producer of better renderings, nor merely the consumer of AI outputs, but the agent who must operate within an expanded representational field in which images are increasingly powerful, unstable, and semantically charged. Representation survives, but under altered conditions: less as depiction, more as negotiation; less as closure, more as controlled drift.

7. CONCLUSIONS

This paper began with a distinction that may appear simple but has wide consequences for contemporary architectural representation: rendering is not the same as semantization. Rendering concerns the visual completion of an image; semantization concerns the process through which a geometry acquires architectural meaning. The emergence of generative AI makes this distinction newly visible because it amplifies the semantic dimension of the image beyond the conventional role of the render.

The experiments presented here suggest that GenAI should not be understood primarily as a faster or more persuasive rendering engine. Its more significant contribution lies in its capacity to operate as a device of intensified semantization. By acting on the same source geometry through multiple iterative transformations, it becomes possible to observe how architecture acquires, loses, thickens, or mutates meaning through typological, atmospheric, and affective shifts. In this sense, the AI image does not simply depict architecture; it participates in the construction and destabilization of architectural legibility.

The staged protocol developed in the paper—moving from natural semantization to meta-object formation, anticipatory semantization, and intensified semantization—proved especially useful because it allowed the image to be read as a

sequence rather than as an isolated result. This sequential structure made visible not only successful semantic branches, but also their internal failures. The concept of semantic inflation emerged precisely from this need: to name those outputs in which semantic density becomes excessive, literal, or reductive, and therefore loses analytical and architectural value. The paper therefore argues that semantic intensification is only productive when it preserves ambiguity, plausibility, and disciplinary relevance.

Across the three branches of intensified semantization, different dimensions of architectural meaning came into focus. The typological branch showed that a geometry can be pushed toward overdetermined architectural reading through aggregation and hybridization. The atmospheric branch demonstrated that semantic drift can occur through perceptual and climatic embedding, without relying primarily on typological cues. The affective branch showed that geometry may also become semantically charged through bodily, uncanny, and tactile associations that exceed conventional architectural coding. Together, these branches suggest that architectural semantics is not a single layer added to form, but a field of possible transformations distributed across material, perceptual, and cultural registers.

The second case study, based on the defamiliarization of a canonical building, reinforced this argument from the opposite direction. It showed that architectural identity may remain geometrically stable while becoming semantically unstable. A building may still be recognized as itself even as its tectonic and material regime is displaced toward virtuality, flesh-like softness, or other affective substitutions. This confirms that recognizability is not guaranteed by geometry alone, but by a broader semantic coherence that can be recoded, intensified, or disrupted.

The broader implication is that the future of representation is not defined only by improved realism or technical efficiency. It is increasingly defined by the transformation of the image into a site of semantic negotiation. Representation becomes less a final depiction and more a branching protocol

through which meaning is tested, compared, and judged. Under these conditions, architectural virtuosity is also displaced. What matters is no longer only the mastery of image production, but the capacity to construct source objects, design semantic sequences, interpret drift, reject inflation, and exercise critical control over a proliferating field of plausible outputs.

For that reason, the contribution of this paper is not only conceptual but methodological. It proposes a way of using generative AI critically: not as an automatic producer of seductive images, but as a laboratory for analysing how architecture becomes legible, how it mutates, and how it may exceed its own conventional semantic boundaries. If rendering once anticipated the future building, intensified semantization now opens a more unstable territory in which architectural meaning can be accelerated, thickened, and contested. The task of representation, under these conditions, is no longer simply to show architecture, but to negotiate what architecture may become.

REFERENCES

1. Eco, U. (1968). *Function and Sign: The Semiotics of Architecture*. In G. Broadbent, R. Bunt, & C. Jencks (Eds.), *Meaning in Architecture*. London: Barrie & Rockliff.
2. Jencks, C. (1969). *Semiology and Architecture*. In G. Broadbent, R. Bunt, & C. Jencks (Eds.), *Meaning in Architecture*. London: Barrie & Rockliff.
3. Venturi, R., Scott Brown, D., & Izenour, S. (1972). *Learning from Las Vegas*. Cambridge, MA: MIT Press.
4. Freedberg, D., & Gallese, V. (2007). Motion, emotion and empathy in aesthetic experience. *Trends in Cognitive Sciences*, 11(5), 197–203.
5. Pallasmaa, J. (1996). *The Eyes of the Skin: Architecture and the Senses*. London: Academy Editions.
6. Steyerl, H. (2009). In defense of the poor image. *e-flux Journal*, 10.
7. Steyerl, H. (2017). *Duty Free Art: Art in the Age of Planetary Civil War*. London: Verso.
8. Paglen, T. (2016). *Invisible images (your pictures are looking at you)*. *The New Inquiry*.
9. Benjamin, W. (2008). *The Work of Art in the Age of Its Technological Reproducibility, and Other Writings on Media*. Cambridge, MA: Harvard University Press. (original essay published 1936)

FIGURES

All figures are by the author unless otherwise stated. Original photographs, diagrams, and AI-generated transformations were produced by the author.