



Points of view

Description

Media inevitably influence the configuration, evolution, design, and experience of cities. By media I mostly mean that which is produced through the basic tools of drawing, painting, sculpture, model-making, photography, film, CAD, CGI, 3D digital renders, and camera phones (though that list expands to considerations of the mass media, news, broadcasting, social media, etc.) More recently we can add the medium of *synthetically generated and enhanced AI video*.

My article from August-September 2014 #211 [Space dehomogenised](#) addressed the contrast between Cartesian coordinate systems and Otto Bollnow's (1903-1991) account of everyday spatial experience. He emphasised the importance of human embodiment in understanding space beyond the mathematical representations on which digital media seem to depend.

His writing fits within a human-centred, embodied, phenomenological, and practical orientation to *making things*, i.e. design. Where possible I like to bring that orientation to bear on how I think about digital media in all its forms.

The article that followed my Bollnow article in the sequence was titled #212 [Eye contact](#). There I mused on how we need to feel that someone speaking to a camera is looking at us, the audience.

Hence, I presented the construction of an improvised teleprompter. One's POV (point of view) is part of embodied spatial experience. The body's vertical axis, its orientations and rotations constitute the basics of our POV.

Moving the camera

With that in mind, I decided to animate the view of a singularly axial building that I visited last month. It's the Radcliffe wing of the Bodleian Library, known as the 'camera'. That's simply Latin for a 'room' or 'chamber', though the word bears obvious connotations in the world of visuality.

I took a photograph from the tower of the adjacent church of St Mary's. Back at my computer, I instructed the [RunwayML](#) text-to-video app to move my camera position so that it looks down on the roof, thereby demonstrating the axial symmetry of the building.

I say something about the processes of such interpolations in another post. See [Surprise videos](#). Anyone who knows the area of Oxford will spot various visual discrepancies. Pullman's *His Dark Materials* and Rowling's *Harry Potter* books and films prime us to see Oxford's enchanted domes and spires through cinematic licence.

It's useful for visualisation, but is such a tool relevant to designing a building or urban space? There is a growing collection of videos demonstrating attempts to use synthetic video generation in a design context. Here's an example I assembled from my own work.


From sketch to video

I resurrected a sketch I had produced long ago while a landscape architecture student – well before the AI renaissance. I had already finished my architecture degree and had worked in architectural offices, so was attuned to the built aspects of my student landscape projects.

In this mode of landscape design, buildings were usually represented by a blank outline on a ground plan, or a vague silhouette in a picture that otherwise prioritised landforms, land-based structures and vegetation.

If we had access to generative AI then we could have filled those building-less voids with *impressions* of buildings generated from atmospheric texts and photographs.

For this recent, 2026, demonstration I started with ChatGPT: Here is a sketch of a proposal for a harbour development outside Melbourne. I drew this while a student of landscape architecture.



There's a power station in the background on the right, and a shipping terminal on the left. The proposed jetty and multi-storey mixed-use housing are clad in timber. I asked ChatGPT to turn this sketch into a photorealistic rendering.

The result was convincing, though it left out the Russian ship, moved the shipping harbour structures further away and ignored the waterway inlet on the right. It also showed a rubble wall on the land-side.

I wanted a beach, so I offered an extra prompt: The bay is fringed by a beach with people on it. Also show people on the jetty.

Amongst various anomalies, the AI showed the entry to my little harbour as too close to the beach, but I left it at that.

With a bit more time and compute at my disposal I could have prompted again for these details to be corrected. The ChatGPT rendering suffices for my demonstration.

I copied that image to RunwayML with the instruction: The camera flies to the left clockwise keeping this speculative harbour development in the centre of the frame.

As I have noted before, the app seems to have difficulty with certain prompts about the direction of travel. See [What lies around the corner](#). But the journey forward was enough for me. Besides, I was aware of my diminishing cache of credits with the software.

I went back to ChatGPT. It only deals in still images. I offered the prompt: RunwayML generated an interesting fly through for me. This is the final frame. Please re-render it to show plausible architectural construction and detailing. Keep the people and environmental conditions the same.

The details of this close up were slightly different from the prompt image. I asked for a new picture: This is a view of the east facade. Please show the same building from the south side.

Instead, it produced a view slightly to the north. But that didn't matter for my demonstration. Then I said: Now show the building in plan view, i.e. looking down as if a Google "satellite" view.

It complied easily enough with a top view. Finally, I returned to an earlier generated image: Please turn this into a pencil sketch, with water colour washes. The AI complied and I edited the stages together into this short video.

No doubt, with further prompting and regeneration, I could achieve greater consistency between the building and jetty in the various shots, and iron out some implausible configurations. In the current images, the headroom beneath the café terrace seems constricted.

The relationship between the beach, the jetty, the waterline and the building platform is not entirely convincing. I would also want to try different materials. A four-storey timber-clad structure in this exposed waterfront setting would face safety, durability, insurance and regulatory challenges.

An iterative "workflow"•

Architects, designers, planners, developers and educators have already begun to ask how such processes might fit into an iterative design "workflow"•

Image generation resembles a rapid form of concept modelling, mood boarding or speculative rendering. It can help make an idea communicable before it has been resolved as a building, landscape, strategy or policy.

But there is also an interesting negative use of such outputs. Synthetic renders can be valuable because they expose issues. In the language of Phenomenology you could say they reveal and conceal aspects of a thing, task or practice. They give critics, students, clients and designers something tangible to interrogate. The image becomes a provisional object for scrutiny.

A drawing, model, image or render is something to react to. It provokes questions. It reveals contradictions. It invites others to say: that part works, that part does not, that condition has not yet been considered.

In the harbour images, for example, the apparent combination of timber buildings, café terraces, jetty decks, dunes, beach, industrial backdrop and calm water raises practical questions.

How close can a sandy beach be to a timber-decked development? What maintenance would be required to prevent sand migration, erosion or water damage?

What conditions would allow a dune landscape to survive in a heavily used public waterfront? Where are the vehicles, deliveries, refuse collection, emergency access and disabled parking? How would fire appliances reach the building? How would occupants evacuate from the upper levels? What happens during storms, high tides or sea-level rise?

Weathered timber cladding raises questions of durability, salt exposure, fire performance, maintenance cycles and detailing. A metal roof at the waterfront would need specification of fixings, flashings, gutters and drainage.

An expert would ask how the delicate roof lines and window junctions suggested by the render are to be built, how they shed water, how they age, and how much they cost.

The images show people on the beach and jetty, but there would need to be shade? How does the place work in winter, in strong wind, or at night? Are the open spaces public, or private? Is the beach for recreation or set apart for environmental restoration? How does the development relate to the industrial surroundings – as contrast, concealment, continuation or critique?

The skilled critic

This is where AI-generated media might be useful in education. Students can be asked not only to produce AI images, but to analyse them. A synthetic render can be treated as an independently produced flawed proposition. It can be marked up and questioned, as well as reverse-engineered.

What assumptions does the image make about structure, access, weather, regulation, ecology, labour, cost and use? What does it omit? What would need to be drawn in plan, section, detail or specification before the proposal could be taken seriously?

In this sense, synthetic image generation is not a substitute for design expertise. It is a prompt for architectural inquiry. It can stimulate the move from appearance to performance, from atmosphere to construction, from visual plausibility to feasibility.

The synthetic image can serve as a meeting point for different forms of judgement. A planner may see one set of problems. An architect, contractor, building control officer, developer, local resident will see the challenges through a different framing.

I could even consult an AI agent, general or specialised, for a first round of criticism: ask about access, fire safety, coastal durability, servicing, structure, environmental impact, planning issues or likely cost risks as a useful checklist.

Anthropic has released a version of its chatbot [Claude Design](#) that promises all-in-one design assistance to paid subscribers. Were this to take off, the design process could reduce to a series of agents generating and critiquing each other's propositions!

Cognitive surrender

That said, for some designers time and effort might be better spent developing their own initial proposals and subjecting these to the scrutiny and critique of other human experts and colleagues. An interesting [article](#) by David Brooks in *The Atlantic* suggests several work-related responses to AI.

Drawing on research into emerging AI-based work practices, he favours the attitude of those – mental marathoners – who enjoy hard thinking and use AI not to avoid effort but to intensify or redirect it.

They prompt for hints rather than answers, start with their own analysis, use AI to challenge rather than replicate thinking, swap AI-assisted and unaided tasks, and treat AI as a guide to thinking rather than as an authoritative source. That might be the best outcome for an AI-based design studio practice that incorporates AI video.

Reference

- Brooks, David. “The People Who Will Thrive in the AI Age.” *The Atlantic*, 28 June, 2026. Accessed 4 July 2026. <https://www.theatlantic.com/ideas/2026/06/ai-open-ai-anthropic/687689/>

Note

- Brooks draws on researches into workplace practices that identify office worker responses to AI. There are those that seek to minimise cognitive effort: “They use AI to make work easier and to think less.” Their own capacities may weaken because the task no longer keeps them in the “zone of optimal difficulty,” where tasks are hard enough to require effort, but not so hard that they overwhelm the worker. A variant response is that some workers recognise the danger of being diminished by AI, and try to resist AI incursions into the work environment. But under the pressure of deadlines, efficiency demands, multitasking, and expectations of greater output, they drift into a reluctant over-reliance on the AI tools, a kind of “cognitive surrender.”

Category

1. Architecture
2. Artificial Intelligence
3. Design
4. Podcast

Tags

1. AI
2. generative video
3. RunwayML

Date Created

July 14, 2026

Author

rcoyne99