



Bohemian melancholy

Description

I'm following on from the post #177 [Frequent feelings](#) of 4 January 2014 with an illustration. There's a private hotel-guesthouse on the outskirts of Kandy in Sri Lanka called Helga's Folly. The owner (Helga De Silva Blow Perera) describes it on Booking.com as an "Anti Hotel Residence". All the public rooms are packed with antique furniture, portraits, ornaments, stags' heads, mirrors, family memorabilia, murals and graffiti.

The hotel boasts that Elizabeth Taylor stayed here during the filming of *Elephant Walk* in 1954. Needless to say, this self-parody of a hotel lends itself to cinematic experimentation. I took this photograph in one of the rooms and invited AI video software to animate movement through the space.



I prompted the Runway platform: "The theme is Bohemian Melancholy. The camera moves into the scene. Outside the scenery is post apocalyptic, revealed in full framed by the square window at the end of the room."

The AI generated video tracked sideways through the room. Then it moved to another room of its own invention. It wasn't what I expected. So, I added a more specific prompt: "Stay in the room as per the first frame. Move the POV further into the room until we reach the end window and the view outside, which is a smouldering city in ruins after an apocalyptic event." I then edited the various outputs together.

For someone steeped in architectural geometries and their presentation using 3D modelling and CAD tools, the convincing creation of AI videos is remarkable. These videos seem to show accurate perspective, parallax and movements through space as if calculated from a digital 3D model. But there is no such model.

How does it do this? As ever, ChatGPT furnished me with some high-level observations and insights about the technology, which I've edited in what follows.

The camera makes space

It seems that AI video generation is successful at perspective and parallax because its neural network (NN) substrate has absorbed very strong statistical regularities about how space appears under camera movement.

Contemporary AI video systems are trained with vast quantities of moving-camera footage: drone flights, handheld walk-throughs, stabilised tracking shots, cinematic push-ins. Across millions of training examples the system encounters consistent correlations: nearer objects shift more rapidly across the frame than distant ones; vertical edges converge; ceilings and floors deform differently under forward motion; occluded surfaces reappear in predictable ways.

The neural models underlying these systems internalise such correlations as patterns in pixel changes over time. So, for example, when prompted with "move forward into the scene," it predicts how the succession of images in the video should transform. It does this because it has been exposed to countless similar transitions.

Time-based prediction itself inclines the neural model toward something like implicit 3D calculation. To generate the next frame convincingly it will incorporate depth, handle how one object might occlude another, and how surfaces retain their appearance and integrity from frame to frame.

If the camera moves so that part of a wall is temporarily obscured by an item of furniture, the wall texture will plausibly persist when that part of the wall is revealed again. According to my ChatGPT guide, "even without explicit geometry, the network develops *latent spatial structure* because consistent frame-to-frame prediction requires it." Quasi-3D representation emerges from the system's accommodation of camera motion. In this sense the camera "makes" the space.

Generative AI video platforms introduce other sophistications, but it's worth noting that human perception also plays a role. We audiences tolerate errors in parallax if overall motion remains smooth

and lighting is coherent.

We may not notice small geometrical inconsistencies in short clips; the human perceptual system accepts the overall flow as plausible. Most AI generated videos are only a few seconds long. So, they need to sustain their coherence only briefly. According to my AI guide, over *longer* durations, inconsistencies accumulate and the illusion weakens.

These techniques apparently favour panoramic or architecturally regular scenes because they already encode strong spatial cues: repeated verticals; receding lines; overlapping figures; consistent texture gradients. A cave temple interior, as in my earlier posts ([A place to meditate](#) and [Surprise videos](#)), contains aligned Buddha statues, railings, and ceiling motifs that provide stable anchors for inferred depth. Redundancy in the scene assists the apparent depth stratification in the neural network model.

It is worth reiterating that the video generating system is not calculating perspective analytically; it is not constructing a mathematical geometric model; it is not explicitly calculating architectural form. Instead, it performs high-dimensional probabilistic prediction constrained by absorbing spatial regularities.

According to my AI guide, in the case of AI generated video, perspective has become a statistical prior embedded in neural weights. The rules of projection persist, but as distributions rather than calculations with geometrical data.

Melancholy and the horizon

So far, I've been attempting to describe (with assistance) something of the process by which the AI video platform generates the output shown here. How do the process and the output intersect with discourses on melancholy?

One way is to think about the nature of the camera and its association with a point of view. Without explicit prompting about geometry, the Runway AI video platform landed on a final frame with a pronounced horizon crenelated with smoke and ruin. Apart from the grim content suggested by my prompt, a horizon is a visual trope of the melancholic disposition.

The German cultural critic Walter Benjamin (1892-1940) theorised the Bohemian mode of being, particularly as observed in the case of the Parisian flaneur. He also identified the melancholy entailments of the visible horizon. It's a motif repeated in countless films portraying loss, yearning, distant prospects, and even the emptiness that for some follows the achievement of hard-won goals. I've explored the perspectival aspect of melancholy elsewhere. See post: [You have reached your destination](#).

Category

1. Architecture
2. Artificial Intelligence

Tags

1. hotel
2. macabre

- 3. melancholy
- 4. mood

Date Created

March 7, 2026

Author

rcoyne99

default watermark