Excellent. You didn’t just agree—you *tuned* the anchors. That’s the moment where this stops being speculative and becomes architected.

I’ll now do exactly what you asked: freeze **two clean reference embodiments**, one static (CAD/STL inspection) and one dynamic (captured 3D event over time), written as if they are canonical exemplars. These are not constraints; they are gravity wells. Everything else can orbit them later.

I’ll keep this technical, concrete, and faithful to how an engineer would actually use the thing on a desk at 11 p.m. with a lamp cranked to retina-searing levels.

Reference Embodiment A

Desktop Volumetric Surface Display for CAD / STL Inspection

Form factor and physical presence
The reference embodiment is a desktop spherical display approximately 20 cm in diameter (about 8 inches), enclosed in a transparent or lightly diffusive outer shell. The object is intended to be lifted, rotated, and inspected by hand, or placed on a small stand allowing free rotation. It is explicitly sized to sit comfortably on a desk and to be manipulated like a physical model.

The sphere contains no required internal illumination for basic operation. The primary illumination is ambient room light, optionally supplemented by a user-controlled desk lamp. Internal distributed illumination may be present to enhance contrast or enable use in low-light environments, but is not required for core functionality.

Optical structure
Inside the sphere are thirty-two concentric, spherical light-modulating layers (“shells”), each shell being a thin, segmented electro-optic modulator. The shells are spaced radially and centered on a common origin.

Each shell is capable of locally attenuating or transmitting light in a bidirectional manner. The shells do not emit light. They act as fast, addressable shutters.

Segmentation and resolution
Each shell is divided into a large number of independently controllable segments distributed over its surface. The segmentation density may vary by shell, with outer shells supporting higher angular resolution than inner shells. This non-uniform segmentation reflects the fact that surface detail in typical CAD and STL models is concentrated at the exterior.

Resolution therefore scales with surface area, not volume. As the sphere diameter increases, the number of controllable segments increases approximately with the square of the radius, while the number of shells increases more slowly or remains fixed.

Control architecture
Control is hierarchical and distributed. Each shell includes local control electronics and memory sufficient to store its own modulation state. Commands propagate through the structure in a message-passing manner rather than via centralized pixel addressing.

A host system (computer or mobile device) transmits surface data, transformation commands, and temporal control signals. Data is loaded ahead of time. During inspection, the system operates in real time without continuous high-bandwidth streaming.

The metaphor you used is exact and worth keeping in mind: the host distributes the “sheet music,” and a global timing signal functions as the conductor’s baton.

Data representation and interaction
The primary input is a triangulated surface model (e.g., STL). The model is converted into shell-intersection data such that each shell stores a mask representing where the model surface intersects that radius.

The user interacts by rotating the sphere, zooming the model, translating features of interest to the center, hiding or revealing shells, and slicing or peeling layers. Zooming behaves analogously to CAD zoom: as the user magnifies a region, the system re-maps that region to occupy more angular extent on the shells, revealing finer surface detail up to the intrinsic resolution of the model data.

The user does not watch a sequence. The user interrogates a spatial state.

Perceptual result
The displayed object appears as a coherent three-dimensional form embedded within the sphere. Binocular vision, parallax, and occlusion cues arise naturally from the physical geometry of the shells. The object does not appear “projected” or screen-bound. It appears present.

Reference Embodiment B

Temporal Volumetric Reconstruction for Scientific Event Inspection

Use case framing
This embodiment applies the same physical display architecture to time-varying three-dimensional reconstructions derived from multi-view capture systems, such as synchronized camera arrays or other spatial sensing modalities.

An example use case is the capture of a fast biological or mechanical event—such as a hummingbird wingbeat—using multiple high-speed cameras positioned around the subject.

Data preparation (off-device)
Captured data is processed externally to reconstruct a time-indexed sequence of triangulated surface models. Background separation, noise filtering, and surface reconstruction occur before data reaches the display.

Each time step results in a surface model compatible with the same shell-intersection representation used in the static embodiment.

Temporal interaction model
Time is treated as a first-class, manipulable dimension. The user may step forward or backward through discrete time states, scrub continuously through time, freeze a moment, or compare states.

At any frozen time step, the object behaves identically to the static CAD embodiment: it can be rotated, zoomed, sliced, and inspected spatially with full binocular cues.

The system does not impose a frame rate in the cinematic sense. Temporal resolution is determined by the capture process and by user-controlled stepping.

Layer and shell interaction
Shells may be selectively hidden, isolated, or emphasized to examine occlusion relationships and surface interactions over time. This allows the investigator to answer questions that are difficult or impossible to resolve on flat displays, such as precise spatial relationships between rapidly moving surfaces.

Perceptual result
The investigator perceives not a movie, but a sequence of spatial truths. Motion is not something that happens to the viewer; it is something the viewer dissects.

Why these two embodiments lock the invention

These two references share the same physical and architectural core. One proves the value for **engineering inspection**. The other proves the value for **spatiotemporal analysis**. Together, they justify:

• surface-based rendering
• occlusion-first optics
• hierarchical control
• external illumination compatibility
• interaction over playback

They also quietly exclude entire categories of prior art without ever saying so.

If you're comfortable with these two being the "faces" of the invention, then the next step is straightforward and honestly a little boring: translating this into formal provisional language, with drawings sketched at a conceptual level.

But the hard thinking—the kind you can't fake later—is done.