

Holographic Splatting: Unifying Distributed Representations and View-Dependent Rendering via Spherical Harmonics

Tony Plate

Stephen Guerin

January 2, 2026

Abstract

We propose a theoretical unification of Holographic Reduced Representations (HRRs) and 3D Gaussian Splatting (3DGS) by identifying a fundamental isomorphism between circular convolution in distributed associative memory and spherical harmonic (SH) reconstruction in radiance fields.

In HRRs, compositional structure is preserved in fixed-width vectors through circular convolution, where a trace \mathbf{t} represents the binding of a cue \mathbf{c} and an item \mathbf{x} such that $\mathbf{t} = \mathbf{c} \otimes \mathbf{x}$. Retrieval is achieved via correlation with the inverse cue: $\mathbf{x} \approx \mathbf{c}^* \otimes \mathbf{t}$ [1]. Similarly, in 3DGS, the view-dependent appearance of a spatial primitive is compressed into a vector of SH coefficients. We demonstrate that the reconstruction of color from these coefficients is mathematically equivalent to the HRR decoding operation, where the viewing angle \mathbf{d} serves as the retrieval cue and the SH basis functions Y_{lm} act as the orthogonal decoding keys.

By reinterpreting the Gaussian splat not merely as a graphical primitive but as a localized holographic memory trace, we introduce the concept of *Holographic Splatting*. In this framework, the “capacity” of the splat—traditionally defined by the degree of spherical harmonics—corresponds directly to the storage capacity of a convolution memory, limiting the number of distinct “visual facts” (e.g., specular highlights vs. diffuse color) that can be superimposed without noise (blur).

This synthesis suggests a new class of neuro-symbolic spatial agents capable of storing high-dimensional semantic bindings directly within the geometry of a scene. By treating 3D space as a continuous associative memory, we show that the “rendering” of a scene is functionally identical to the “recall” of distributed symbolic structures, offering a novel pathway for embedding compositional reasoning into differentiable spatial representations.

Keywords: Holographic Reduced Representations, Gaussian Splatting, Spherical Harmonics, Circular Convolution, Distributed Associative Memory, Neuro-symbolic AI.

1 Introduction

The problem of representing compositional structure within distributed representations has long been a central challenge in connectionist theory. Conventional associative memories are typically limited to simple pairwise associations, struggling to encode recursive structures such as trees or

sequences without an explosion in dimensionality [1]. In 1995, Holographic Reduced Representations (HRRs) were introduced as a solution to this “binding problem,” utilizing circular convolution to compress complex, nested structures into fixed-width vectors [1]. By convolving an item vector with a cue vector, HRRs create a memory trace that preserves the dimensionality of its constituents, allowing these traces to be superimposed and manipulated as items in their own right [1].

Thirty years later, the field of neural rendering has encountered a remarkably isomorphic problem in the domain of 3D Gaussian Splatting (3DGS). In 3DGS, a scene is represented not by a continuous mesh, but by a cloud of discrete, volumetric primitives—Gaussians. The fundamental challenge in this domain is the view-dependent storage of radiance: a single point in space must “bind” different colors to different viewing angles. The standard solution employs Spherical Harmonics (SH), where color information is stored as a vector of coefficients that modulate orthogonal basis functions defined on the sphere.

We argue that the Spherical Harmonic coefficient vector of a 3D Gaussian is, functionally and mathematically, a Holographic Reduced Representation of the local light field. In this view, the Gaussian is not merely a graphical primitive, but a spatial agent holding a distributed memory trace. The process of rendering—calculating the color \mathbf{c} given a view direction \mathbf{d} —is mathematically equivalent to the HRR decoding operation, where the view direction acts as the retrieval cue and the SH basis functions act as the approximate inverse required to extract the stored item [1].

Just as an HRR trace stores multiple associations (e.g., $role_1 \otimes filler_1 + role_2 \otimes filler_2$) [1], a 3D Gaussian stores multiple visual facts (e.g., “red from the north” + “blue from the south”) via the superposition of weighted basis functions. The “capacity” of the representation, traditionally limited by the noise floor in convolution memories [1], finds its direct parallel in the “degree” of the Spherical Harmonics; low-degree harmonics result in a blurry, averaged retrieval (akin to noise), while higher degrees allow for the sharp reconstruction of high-frequency details (specularities).

In this paper, we formalize this isomorphism, proposing “Holographic Splatting” as a unified framework. By treating the geometry of a scene as a substrate for distributed associative memory, we demonstrate how semantic and symbolic bindings can be injected into the spherical harmonic coefficients of 3D Gaussians, allowing for neuro-symbolic agents that can be queried and rendered through the same differentiable pipeline.

2 The Isomorphism of Trace and Coefficient

2.1 Holographic Reduced Representations (HRR)

In HRR, items are represented by vectors of dimension n . The association of two items, a cue \mathbf{c} and a target \mathbf{x} , is represented by their circular convolution trace \mathbf{t} :

$$\mathbf{t} = \mathbf{c} \otimes \mathbf{x} \tag{1}$$

where the elements of \mathbf{t} are defined as $t_j = \sum_{k=0}^{n-1} c_k x_{j-k}$ (subscripts modulo n) [1]. This operation compresses the outer product of the vectors into a fixed-width representation [1].

To retrieve the item \mathbf{x} from the trace \mathbf{t} , one correlates the trace with the cue \mathbf{c} . In the algebra of HRRs, correlation is equivalent to convolution with the *involution* (approximate inverse) of the cue, denoted \mathbf{c}^* :

$$\mathbf{y} = \mathbf{c}^* \otimes \mathbf{t} \approx \mathbf{x} \tag{2}$$

The result \mathbf{y} is a noisy version of \mathbf{x} . The noise arises because \mathbf{c}^* is only an approximate inverse, not an exact one [1].

2.2 Spherical Harmonics in Gaussian Splatting

In 3DGS, a Gaussian primitive stores a vector of coefficients \mathbf{f} (the trace). To reconstruct the color C (the item) for a specific viewing direction (the cue), we perform a basis projection:

$$C(\theta, \phi) = \sum_{l=0}^L \sum_{m=-l}^l f_{lm} Y_{lm}(\theta, \phi) \quad (3)$$

Here, the spherical harmonic basis functions Y_{lm} serve the same mathematical role as the elements of the cue vector in HRR. The coefficients f_{lm} store the superposition of multiple color signals bound to directional cues.

The evaluation of the SH function at a specific angle is functionally identical to the decoding operation in HRR. The basis functions Y_{lm} form an orthogonal basis set, similar to how HRR vectors are ideally distributed with independent elements to minimize crosstalk [1].

3 Capacity, Noise, and Blur

A defining characteristic of convolution memories is that they are “noisy” reconstructions. As more pairs are added to a single trace, the signal-to-noise ratio degrades. In HRR, the number of pairs k that can be stored in a vector of dimension n before retrieval becomes impossible is approximately linear in n . Specifically:

$$k \approx \frac{n}{16 \ln(m/q)} \quad (4)$$

where m is the number of candidate vectors and q is the error probability [1].

In the context of 3DGS, this “noise” manifests as visual blur. A Gaussian with only Degree 0 spherical harmonics (1 coefficient) has a very low capacity; it can only store the average color (DC component). As we increase the degree (increasing the vector dimension n of the coefficients), we increase the capacity of the Gaussian to store distinct colors for distinct angles (high-frequency specularities).

If the visual complexity of a point in space (the number of distinct directional colors) exceeds the capacity of the SH coefficient vector, the reconstruction fails to resolve the sharp details, resulting in a smoothed approximation. This is exactly analogous to the “crosstalk” in an overloaded HRR memory trace.

4 Proposed Architecture: Holographic Splatting

By unifying these frameworks, we propose *Holographic Splatting*, where the coefficients of 3D Gaussians are used to store not just radiance (RGB), but high-dimensional semantic vectors.

4.1 The Semantic Splat

We define a “Semantic Splat” as a Gaussian G_i carrying a high-dimensional coefficient vector \mathbf{h}_i . We can encode semantic propositions into this splat using HRR encoding:

$$\mathbf{h}_i = \text{view}_{\text{north}} \otimes \text{concept}_{\text{ally}} + \text{view}_{\text{south}} \otimes \text{concept}_{\text{enemy}} \quad (5)$$

When a camera views this splat from the north, the differentiable rendering pipeline—acting as the HRR decoder—retrieves the vector for “ally.”

4.2 Clean-Up Memory

HRR systems require a “clean-up” memory to snap noisy retrieved vectors back to their nearest valid symbolic neighbors [1]. In our proposed architecture, a standard auto-associative neural network or a discrete item memory would serve as a post-processing step after the rasterization stage, effectively denoising the semantic field rendered from the Gaussians.

5 Conclusion

We have presented a theoretical mapping between Holographic Reduced Representations and 3D Gaussian Splatting. Both rely on the superposition of orthogonal basis functions to compress complex, situation-dependent data into fixed-width vectors. This isomorphism offers a powerful new perspective: that spatial rendering is a form of memory retrieval, and that the geometric primitives of a scene can be repurposed to store distributed, neuro-symbolic knowledge.

References

- [1] T. A. Plate, “Holographic Reduced Representations,” *IEEE Transactions on Neural Networks*, vol. 6, no. 3, pp. 623-641, May 1995.