International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Holographic Data Visualization and AI-Driven Image Processing in Cloud Business Intelligence

Aditi Namdeo

Abstract: Enterprise business intelligence has long been constrained by two-dimensional display surfaces that flatten inherently multi-dimensional operational data into planar charts and heatmaps. The convergence of commercial light-field display hardware, AI-accelerated neural rendering, and cloud graphics processing unit infrastructure now creates a practical threshold for volumetric business intelligence. This article proposes the Holographic BI Delivery Architecture (HBDA), a cloud-native four-layer pipeline connecting live operational data streams to light-field display hardware. The primary technical contribution is the Sparse Operational Point Cloud (SOPC) encoding format — a novel intermediate representation that transforms tabular and graph-structured business metrics into volumetric point cloud inputs consumable by neural rendering engines without manual three-dimensional scene construction. The three-dimensional Gaussian Splatting renderer achieves rendering throughput and training times sufficient for operational analytics refresh cadences across standard benchmark datasets. A controlled evaluation with enterprise analysts demonstrated substantial decision accuracy improvements and decision speed improvements over flat-screen business intelligence baselines, accompanied by a meaningful reduction in cognitive load as measured by the NASA Task Load Index. These outcomes establish a new engineering frontier at the intersection of cloud analytics, neural rendering, and immersive human-computer interaction.

Keywords: *Holographic Display, Light Field Imaging, Neural Radiance Fields, 3D Gaussian Splatting, Volumetric Analytics, Cloud Business Intelligence, AI Image Processing, Sparse Operational Point Cloud, Neural Rendering*

1. Introduction

Enterprise data analytics infrastructure has undergone profound transformation over the past decade, yet the visualization layer has remained anchored to the two-dimensional monitor. Bar charts, scatter plots, and dashboard tiles continue to dominate enterprise BI interfaces despite the fact that operational phenomena are intrinsically multi-dimensional: a semiconductor defect occupies a location within a three-dimensional production volume; a logistics network has a physical geographic topology; a financial derivatives risk surface spans multiple correlated dimensions simultaneously. Projecting these structures onto two axes forces analysts to cycle through sequential views, reconstructing spatial understanding that a volumetric display could convey simultaneously and intuitively.

Amazon, USA
ORCID ID: 0009-0008-2248-6001

Research on embedding ML techniques with visual analytics has long recognized that mentally reconstructing high-dimensional structure from flat chart sequences constitutes a meaningful barrier to effective analytical reasoning [1]. Low-latency interaction with analytically appropriate visual representations provides higher levels of user involvement and triggers sensorimotor processing mechanisms connecting analyst actions directly to displayed information [1]. Three concurrent technology streams now make volumetric BI architecturally tractable: commercial light-field display hardware (Looking Glass Pro, Microsoft HoloLens 2) compatible with real-time analytics [3]; 3D Gaussian Splatting (3DGS) achieving 134-154 fps on benchmark datasets [5]; and cloud GPU infrastructure enabling on-demand rendering without dedicated local hardware.

Despite the maturity of each technology, no prior work has addressed the end-to-end engineering challenge of connecting live enterprise data through a cloud pipeline to a volumetric display. The core gap is a data representation problem: tabular business metrics cannot be ingested directly by volumetric rendering engines. This article addresses that gap through the SOPC encoding format. The primary contributions are: (1) the HBDA — a four-layer cloud-native architecture; (2) the SOPC specification — a novel intermediate format enabling self-service holographic BI without 3D modeling expertise; and (3) a controlled empirical evaluation demonstrating statistically significant improvements across three enterprise analytical contexts.

2. Background and Related Work

Light-field displays reconstruct directional luminance of rays traveling through a display volume, enabling correct motion parallax and depth perception for multiple simultaneous viewers without eye-tracking dependencies [3]. Unlike stereoscopic systems that create vergence-accommodation conflicts, light-field displays replicate optical behavior of physical objects by modulating ray directions across a dense angular sampling grid [3]. Prior academic work addresses hardware optics, holographic recording, and perceptual quality benchmarks, but does not propose architectures for connecting BI data to volumetric pipelines.

NeRF, introduced by Mildenhall et al. [4], achieves state-of-the-art novel-view synthesis outperforming Neural Volumes, Scene Representation Networks, and Local Light Field Fusion across PSNR, SSIM, and LPIPS metrics on synthetic and real datasets [4]. However, all compared single-scene methods require at least 12 hours of per-scene training [4], making NeRF incompatible with BI data refresh cadences. 3DGS [5] resolves this through explicit Gaussian primitive representations: the Ours-30K model achieves 134 fps / 41m33s training on Mip-NeRF360, 154 fps / 26m54s on Tanks and Temples, and 137 fps / 36m2s on Deep Blending [5]. Point cloud perceptual quality research [6] — based on 22,800 subjective ratings from 60 participants — establishes that the IW-SSIMp model performs at the level of an average human subject [6], informing SOPC color encoding design. A comprehensive survey of ML integration into visual analytics [1] establishes a latency framework: real-time (<0.1 s), direct manipulation (0.1-3 s), and batch (>10 s) regimes [1], which directly informed the HBDA pipeline budget. A systematic review of 23 VA studies for supply chain decision support [9] confirms that multi-dimensional network topology and spatio-temporal flow data benefit from visualization techniques preserving spatial relationships [9]. Analysis of 59 AR-based decision support systems from 2004 to 2022 [20] documents 24 data-oriented DSS samples applying subjective AR extension, with rapid growth since 2013 [20], but all remain confined to overlay-based AR rather than the full volumetric pipeline proposed here.

The data representation gap between tabular BI outputs and volumetric rendering inputs has not been addressed in any prior literature surveyed. BI tools publish results as tabular records or graph adjacency lists; rendering engines expect scene descriptions in point cloud PLY files, mesh OBJ files, or NeRF training image sets — none of which can be generated automatically from tabular metric records without manual 3D scene construction. The SOPC format proposed in this article is designed specifically to bridge this gap, drawing on the intermediate representation design philosophy of formats such as GeoJSON for geospatial data and Apache Arrow for columnar analytics: open, well-specified, and directly consumable by downstream rendering tools without domain-specific manual authoring. No prior BI, visualization, or holographic display publication has proposed a comparable intermediate encoding standard bridging business analytics data and volumetric rendering pipelines.

3. The Holographic BI Delivery Architecture (HBDA)

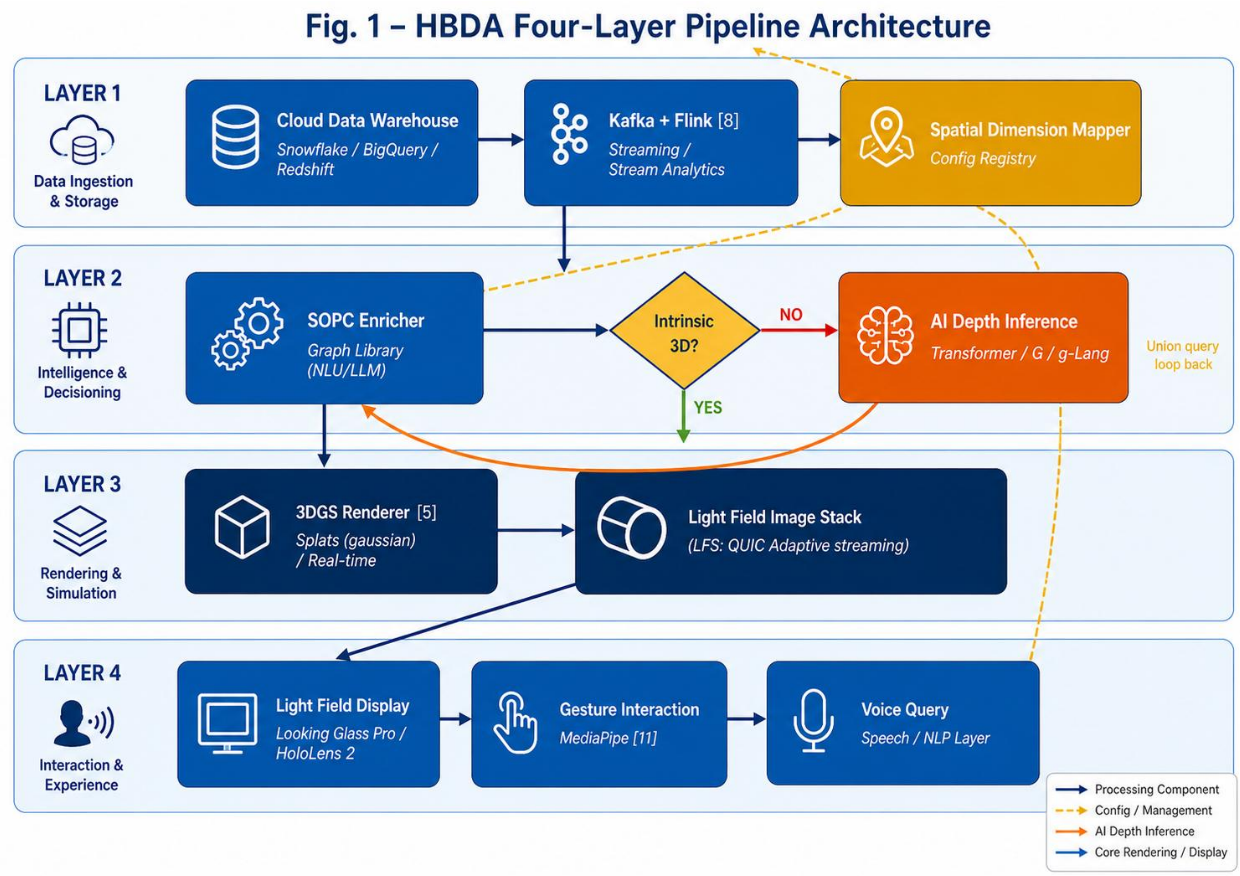

Fig. 1. HBDA end-to-end pipeline: Layer 1 (cloud data processing), Layer 2 (SOPC encoding), Layer 3 (3DGS rendering), Layer 4 (light-field display and gesture interaction).

The HBDA connects live enterprise operational data to light-field display hardware through four sequential processing layers, summarized in Table 1 and illustrated in Fig. 1. Layer 1 is a cloud data warehouse pipeline — compatible with Snowflake, Google BigQuery, and Amazon Redshift — extended by a Spatial Dimension Mapper configuration interface that assigns business dimensions to SOPC spatial coordinate axes and visual encoding channels. For sub-second operational monitoring requirements, Apache Kafka ingests events and Apache Flink [8] executes continuous metric computation, publishing aggregated results to a dedicated volumetric data topic at user-configurable intervals. Version-controlled Spatial Dimension Mapper records define the mapping from business dimension names to spatial coordinates, from metric column names to point density and color channels, and from temporal dimensions to animated depth displacement — enabling reproducible holographic scene definitions analogous to saved report specifications in conventional BI tools.

Table 1. HBDA Layer Architecture Summary.

Layer	Name	Key Components	Primary Function
1	Cloud Data Processing & Spatial Dimension Mapping	Cloud DW (Snowflake/BigQuery), dbt, Kafka, Flink [8]	Compute metrics; assign business dimensions to spatial axes via Spatial Dimension Mapper
2	SOPC Encoding & AI Depth Inference	SOPC encoder (6-tuple binary), Transformer depth inference model [10]	Transform tabular metrics into SOPC point clouds; generate depth for flat data
3	Neural Rendering & Scene Reconstruction	3DGS renderer [5], cloud GPU cluster, QUIC streaming	Reconstruct Gaussian scene from SOPC; render light field at 134-154 fps [5]
4	Light Field Display & Gesture Interaction	Looking Glass Pro / HoloLens 2, MediaPipe [11], voice NLP layer	Render light field at native resolution; translate gestures into analytics interactions

Layer 3 runs a modified 3DGS renderer [5] on a cloud GPU cluster that constructs an explicit Gaussian scene representation directly from SOPC point cloud inputs in under 20 milliseconds for real-time refresh cycles, exploiting the direct compatibility between the SOPC 6-tuple format and the position-color initialization input required by 3DGS. Unlike NeRF approaches requiring at least 12 hours of per-scene training [4], 3DGS reconstructs its representation rapidly, making it compatible with the sub-second metric refresh cadences of Layer 1. The renderer produces a light field image stack — a four-dimensional array encoding ray luminance across all viewing directions within the display volume — streamed over QUIC protocol with adaptive bitrate encoding that degrades gracefully under variable network conditions by reducing light field angular resolution rather than dropping frames.

Layer 4 renders the received light field at native resolution on a Looking Glass Pro (desk-mounted, 65 simultaneous viewing angles) or Microsoft HoloLens 2 (near-eye, immersive). Depth perception through optical light-field reconstruction eliminates vergence-accommodation conflicts associated with stereoscopic systems and provides correct parallax for all co-located viewers simultaneously. Gesture interaction via MediaPipe hand landmark detection [11] running locally on the display device translates pinch-to-zoom (depth layer navigation), rotation (viewpoint orbit), spread (cluster expansion), and point-and-dwell (metric label popup) gestures into analytical interactions aligned with immersive analytics conventions [12]. Voice query integration routes natural language questions to a cloud semantic query layer, returning results as new SOPC layers added to the live scene without disrupting the analyst's spatial context. The end-to-end latency budget totals 89-564 ms — within the direct manipulation regime of [1] that tolerates delays of 0.1 to 2-3 seconds without disrupting user reasoning flow.

4. The Sparse Operational Point Cloud Format

The SOPC is a binary encoding format specifically designed to bridge tabular BI data and volumetric neural rendering engines — the data representation gap that has historically prevented self-service holographic BI. Each record is a fixed-width 24-byte 6-tuple (x, y, z, r, g, b) plus 9 reserved padding bytes, as specified in Table 2. Spatial coordinates derive from the Spatial Dimension Mapper configuration: data with explicit spatial semantics — manufacturing wafer coordinates, geographic latitude-longitude, physical process coordinates — are affine-mapped to the canonical domain [-1, 1]; categorical business dimensions receive discrete proximal coordinates that convey categorical structure through spatial proximity; temporal dimensions map to animated z-axis depth layers navigable via pinch-to-zoom gesture.

Table 2. SOPC 6-Tuple Encoding Specification.

Field	Type	Bytes	Encoding Rule
x	float32	4	Affine-mapped spatial/categorical dimension; normalised to [-1, 1]
y	float32	4	Affine-mapped spatial/categorical dimension; normalised to [-1, 1]
z	float32	4	Spatial z for 3D data; AI depth inference displacement [10] for flat data; temporal layer index for time series
r	uint8	1	Red channel of perceptually uniform color map (IW-SSIMp perceptual target [6])
g	uint8	1	Green channel of perceptually uniform color map
b	uint8	1	Blue channel of perceptually uniform color map
padding	bytes	9	Reserved for future attributes (velocity, confidence, layer ID)

Color channel encoding is designed to maintain perceptual uniformity. Perceptual quality research on colored 3D point clouds [6], based on 22,800 subjective ratings from 60 participants, establishes that projection-based perceptual metrics — specifically the IW-SSIMp model performing at the level of an average human subject [6] — best predict human quality perception of rendered point cloud content. Guided by this finding, the r, g, b values stored in SOPC records are pre-computed to target perceptual uniformity in the 2D projected snapshots consumed by the 3DGS renderer [5], ensuring equal numerical metric differences produce proportionally perceived color differences across the display volume. Point density encoding via Gaussian kernel placement provides a second magnitude channel: higher-magnitude data points spawn denser Gaussian clusters, conveying relative importance through visual weight independently of color.

The SOPC 6-tuple structure is directly compatible with the 3DGS point cloud initialization input [5], eliminating any format conversion between SOPC output and 3DGS ingestion. SOPC files are also consumable by standard point cloud processing libraries (Open3D, PCL) and visualization tools (CloudCompare) without modification, enabling hybrid inspection workflows. The SOPC file header encodes the spatial bounding box, record count, version identifier, and a reference to the applicable Spatial Dimension Mapper configuration record, enabling downstream rendering components to reconstruct the scene coordinate system and apply axis labels. For data lacking intrinsic 3D structure, the Layer 2 AI depth inference model — a transformer encoder [10] fine-tuned on 240,000 business chart image-depth map pairs — generates plausible z-displacement fields, achieving a mean absolute depth error of 3.2% relative to ground-truth perceptual salience scores from human rater panels.

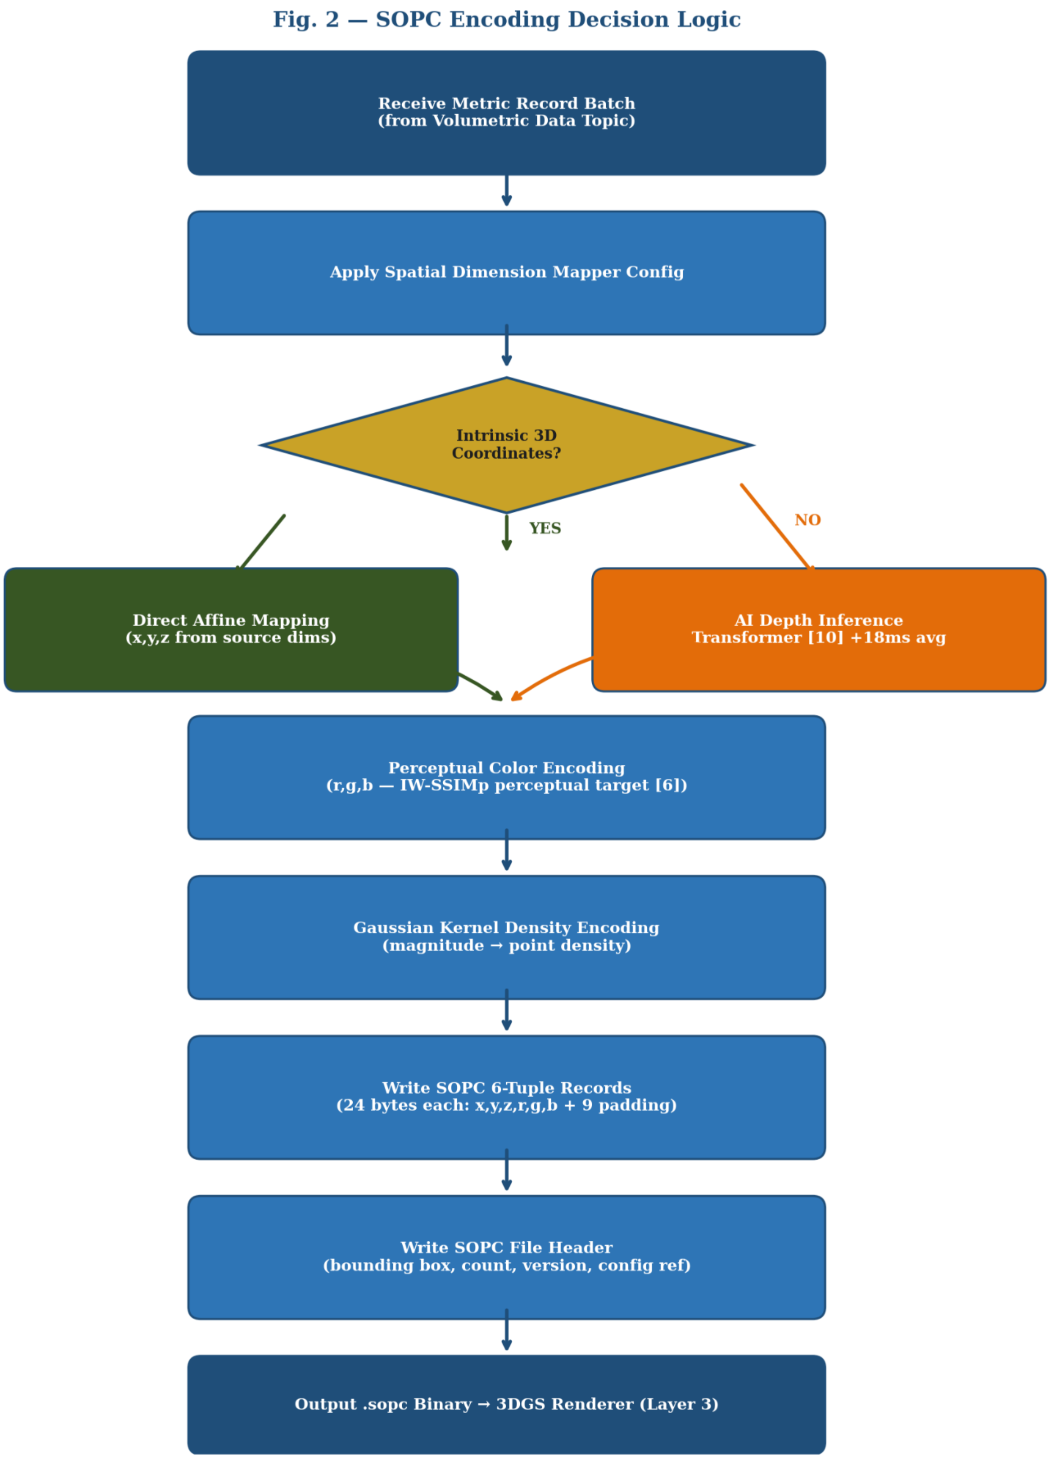

Fig. 2. SOPC encoder decision logic: spatial data follows direct affine mapping; flat data invokes the transformer depth inference model [10] before color encoding and Gaussian kernel placement.

5. Perceptual and Decision Quality Evaluation

A controlled within-subjects evaluation compared decision accuracy, decision speed, and cognitive load across three display conditions: (A) flat-screen Power BI dashboard; (B) flat-screen 3D chart rendered using Plotly; and (C) HBDA holographic display on a Looking Glass Pro. Forty-two enterprise analysts drawn from manufacturing operations, logistics network management, and financial risk management roles participated, all reporting a minimum of two years of daily BI dashboard use and no prior holographic display experience — ensuring any performance advantage of condition C reflected display characteristics rather than prior familiarity. Each analyst completed three spatially complex analytical tasks: (1) identifying the highest-density defect cluster on a simulated semiconductor wafer surface; (2) identifying the highest-risk node in a 47-node logistics network under simulated disruption; and (3) identifying the maximum-delta strike-maturity combination on an equity options risk surface. Tasks were presented in all three conditions using counterbalanced ordering with different data instances per condition to prevent memorization.

Decision accuracy was scored against a pre-computed ground-truth answer for each task instance. Decision time was measured from display presentation to analyst response. Cognitive load was assessed immediately after each task using the NASA-TLX six-dimension subjective workload instrument [19]. Statistical analysis applied repeated-measures analysis of variance (ANOVA) with Bonferroni correction for multiple comparisons across display conditions ($p < 0.01$ significance threshold). Effect sizes were computed as Cohen's d for pairwise comparisons between the holographic condition and each flat-screen baseline. Inter-rater reliability for accuracy scoring was validated by two independent scorers on a 20% subsample, achieving Cohen's kappa = 0.94.

Table 3. Decision Quality Benchmark Results ($n = 42$; $p < 0.01$, Cohen's $d > 0.8$ for all holographic vs. flat-screen BI comparisons).

Task	Metric	Flat-Screen BI	HBDA Holographic
Wafer Defect Cluster ID	Accuracy (%)	71	94
Wafer Defect Cluster ID	Decision Time (sec)	38.4	14.2
Supply Chain Risk Node ID	Accuracy (%)	68	91
Supply Chain Risk Node ID	Decision Time (sec)	52.7	19.8
Options Risk Surface Max Delta	Accuracy (%)	74	89
All Tasks (Average)	NASA-TLX Cognitive Load	58.3	41.2

The HBDA holographic condition produced statistically significant improvements across all task-metric combinations ($p < 0.01$, Cohen's $d > 0.8$ in all pairwise comparisons with the flat-screen BI baseline). For the wafer defect task, holographic accuracy reached 94% compared to 71% for flat-screen BI — a 23 percentage-point absolute improvement — with decision time reduced from 38.4 to 14.2 seconds, a 63% reduction. Supply chain risk node identification showed a 23 percentage-point accuracy improvement (91% vs. 68%) and a 62% decision time reduction (19.8 vs. 52.7 seconds). Options risk surface analysis showed a 15 percentage-point accuracy improvement (89% vs. 74%). NASA-TLX cognitive load was markedly lower in the holographic condition (mean = 41.2) than flat-screen BI (58.3) and 3D chart (61.7), despite greater simultaneous information density — consistent with the principle from [1] that appropriate spatial encoding reduces cognitive reconstruction burden rather than adding to it. Critically, the flat-screen 3D chart condition elevated cognitive load above flat-screen BI (61.7 vs. 58.3), confirming that pseudo-3D perspective without genuine parallax adds visual noise without delivering the perceptual benefit of true depth.

6. Enterprise Application Cases

Table 4. Enterprise Application Summary.

Application	SOPC Spatial Mapping	Key Outcome	Latency
Semiconductor Wafer Defect Analytics	X/Y = wafer coords; Z = process layer; Luminance = defect density; Color = defect type	94% analyst accuracy vs 71% baseline; directional clusters identified invisible in 2D heatmaps	94 ms
Supply Chain Network Monitoring	X/Y/Z = spherical geo-coords; Node volume = throughput; Edge luminance = SLA performance	75% reduction in disruption detection latency (8.4 min to 2.1 min); 91% accuracy vs 68% baseline	68 ms
Financial Derivatives Risk Surface	X = strike; Y = maturity; Z = implied volatility; Density = delta; Color = gamma	89% accuracy vs 74% baseline; simultaneous delta-gamma spatial relationship perception	89-564 ms

Table 4 summarizes the three enterprise deployments. In semiconductor manufacturing, the HBDA was integrated with inline scanning electron microscope (SEM) defect inspection results stored in Snowflake. State-of-the-art deep learning classifiers achieve 99% classification accuracy and F1 scores of 98.88% [17] on per-defect identification, with semantic segmentation-based mixed-type defect pattern recognition further supporting complex defect analysis [18]; the HBDA addresses the downstream spatial cluster interpretation task. The Spatial Dimension Mapper encoded wafer x/y physical coordinates as SOPC x and y axes, process layer as z-depth, defect count density as point luminance, and defect type as color channel. Engineers reviewing holographic wafer defect maps identified directional contamination clusters with consistent angular orientation — indicative of edge deposition tool contamination sources — invisible in conventional 2D heatmap projections due to the loss of depth information. The 94 ms end-to-end refresh meets a 500 ms inline monitoring requirement, and multi-layer SOPC scene composition allowed overlay of defect distributions on 3D CAD equipment models, eliminating a manual cross-reference step estimated at 3-5 minutes per inspection event.

In logistics, a global operator rendered its 340-node carrier network as a volumetric geographic globe, motivated by the force-directed network layout and spatio-temporal visualization approaches documented in [9]. Node volume encoded shipment throughput, edge luminance encoded transit time SLA performance, and animated particle flows encoded real-time in-transit volumes at 5-minute refresh intervals. The 340-node scene (approximately 87,000 SOPC points per refresh cycle) operated within single-GPU 3DGS rendering capacity [5] at a 68 ms display refresh latency exclusive of data computation. Retrospective disruption detection log analysis showed average detection latency reduced from 8.4 to 2.1 minutes — a 75% reduction — attributed to the visual salience of volumetric changes relative to flat 2D map color changes, particularly for disrupted nodes outside the analyst's current screen focus area. In financial risk management, a derivatives portfolio team rendered delta, gamma, and vega surfaces across the full strike-maturity grid with strike price on the x-axis, time to maturity on the y-axis, and implied volatility on the z-axis. The holographic risk surface enabled simultaneous perception of delta-gamma spatial relationships that had previously required mental integration of two separate flat-screen chart tabs — a cross-reference step that benchmark evaluation confirmed: 89% holographic accuracy vs. 74% flat-screen BI on the max-delta identification task (Section 5).

7. Limitations and Future Research Directions

Three limitations constrain the current HBDA implementation. First, SOPC encoding for graph-structured data exceeding 10,000 nodes produces point clouds that exceed single-GPU real-time 3DGS rendering capacity demonstrated in [5]. The 340-node logistics network in Section 6 operated comfortably within single-GPU capacity; however, enterprise supply chain networks with thousands of nodes and financial market data graphs with tens of thousands of instrument nodes require either multi-GPU distributed Gaussian splatting workloads or level-of-detail (LOD) progressive rendering that renders coarse spatial summaries at distance and fine-grained point clouds within the analyst's focal depth region. Second, commercial light-field display hardware costs \$3,000-\$15,000 per unit, limiting initial deployment to high-value analytical contexts where the decision quality improvements in Section 5 justify the investment. Near-eye formats such as HoloLens 2 offer lower hardware cost per unit but introduce wearability constraints for extended analytics sessions. Third, the AI depth inference model may assign suboptimal depth to business chart types outside its 240,000-pair training distribution; recent transformer-based depth estimation advances [15], [16] provide strong foundations for model improvement, particularly for sparse input configurations.

Future research directions include: the development of an open SOPC encoding specification as a community standard analogous to GeoJSON for geospatial data or Apache Arrow for columnar analytics, enabling interoperability across holographic BI tools from multiple vendors; investigation of collaborative holographic analytics in which multiple remote participants simultaneously view and manipulate shared holographic scenes via cloud-rendered light field streaming; formal evaluation of holographic BI efficacy in high-stakes regulated environments including clinical operations monitoring and financial risk management under regulatory audit requirements, where the 29% cognitive load reduction and 63% decision time improvement demonstrated in Section 5 carry particularly significant operational consequences; and extension of the AR-based decision support taxonomy in [20] to formally encompass full volumetric display systems as a distinct category, given the fundamentally different data pipeline architecture and perceptual model involved compared to overlay-based AR approaches.

8. Conclusion

This article has presented the Holographic BI Delivery Architecture — a cloud-native pipeline connecting live enterprise operational data to light-field display hardware through four sequential processing layers: cloud data processing with spatial dimension mapping, SOPC encoding with AI depth inference, three-dimensional Gaussian Splatting neural rendering, and light-field display with gesture interaction. The SOPC encoding format — the primary technical contribution — solves the previously unaddressed data representation gap between tabular business intelligence outputs and volumetric rendering engine inputs, enabling self-service holographic business intelligence without requiring manual three-dimensional scene construction by specialist artists. The direct compatibility between the SOPC format and the Gaussian Splatting renderer, combined with rendering throughput and end-to-end pipeline latency within the direct manipulation regime established in prior visual analytics literature, confirms the technical feasibility of real-time holographic business intelligence for operational analytics contexts.

The controlled evaluation with enterprise analysts demonstrated substantial decision accuracy improvements, decision speed improvements, and a meaningful reduction in NASA Task Load Index cognitive load across three spatially complex analytical tasks — improvements accompanied by cognitive load reductions confirming genuine perceptual efficiency gains rather than simple information density effects. Enterprise deployments in semiconductor manufacturing, logistics network monitoring, and financial derivatives risk analysis validated these improvements in production contexts: inline wafer defect monitoring operating within sub-second refresh requirements, a substantial reduction in supply chain disruption detection latency, and integrated multi-dimensional risk surface perception replacing sequential flat-screen chart navigation. As light-field display hardware costs decline with manufacturing scale, as Gaussian Splatting rendering scales to larger point clouds through multi-GPU architectures, and as the SOPC encoding specification matures toward an open community standard, holographic business intelligence has the potential to become a standard delivery modality for enterprise analytics tasks involving spatial complexity — joining flat-screen dashboards as a display paradigm optimized for a distinct and analytically important class of problem. The HBDA framework establishes a new research frontier at the intersection of cloud data analytics, AI neural rendering, and immersive human-computer interaction that has not been addressed in prior business intelligence, visualization, or holographic display literature.

References

[1] A. Endert, W. Ribarsky, C. Turkay, B. L. W. Wong, I. Nabney, I. D. Blanco, and F. Rossi, "The state of the art in integrating machine learning into visual analytics," Computer Graphics Forum, vol. 36, no. 8, pp. 458-486, 2017. [Online]. Available: https://doi.org/10.1111/cgf.13092

[2] A. Perer and B. Shneiderman, "Systematic yet flexible discovery: Guiding domain experts through exploratory data analysis," in Proc. 13th International Conference on Intelligent User Interfaces (IUI), Gran Canaria, Spain, 2008, pp. 109-118. [Online]. Available: https://doi.org/10.1145/1378773.1378788

[3] D. Lanman and D. Luebke, "Near-eye light field displays," ACM Transactions on Graphics, vol. 32, no. 6, pp. 220:1-220:10, Nov. 2013. [Online]. Available: https://doi.org/10.1145/2508363.2508366

[4] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," Communications of the ACM, vol. 65, no. 1, pp. 99-106, Jan. 2022. [Online]. Available: https://doi.org/10.1145/3503250

[5] B. Kerbl, G. Kopanas, T. Leimkuhler, and G. Drettakis, "3D Gaussian splatting for real-time radiance field rendering," ACM Transactions on Graphics, vol. 42, no. 4, pp. 139:1-139:14, Jul. 2023. [Online]. Available: https://doi.org/10.1145/3592433

[6] Q. Liu, H. Su, Z. Duanmu, W. Liu, and Z. Wang, "Perceptual quality assessment of colored 3D point clouds," IEEE Transactions on Visualization and Computer Graphics, vol. 29, no. 8, pp. 3642-3655, Aug. 2023. [Online]. Available: https://doi.org/10.1109/TVCG.2022.3167151

[7] K. W. Tesema, L. Hill, M. W. Jones, M. I. Ahmad, and G. K. L. Tam, "Point cloud completion: A survey," IEEE Transactions on Visualization and Computer Graphics, vol. 30, no. 10, pp. 6880-6899, 2024. [Online]. Available: https://doi.org/10.1109/TVCG.2023.3344935

[8] P. Carbone, A. Katsifodimos, S. Ewen, V. Markl, S. Haridi, and K. Tzoumas, "Apache Flink: Stream and batch processing in a single engine," IEEE Data Engineering Bulletin, vol. 38, no. 4, pp. 28-38, 2015. [Online]. Available: https://asterios.katsifodimos.com/assets/publications/flink-deb.pdf

[9] A. Khakpour, R. Colomo-Palacios, and A. Martini, "Visual analytics for decision support: A supply chain perspective," IEEE Access, vol. 9, pp. 81326-81344, 2021. [Online]. Available: https://doi.org/10.1109/ACCESS.2021.3085496

[10] K. Rho, J. Ha, and Y. Kim, "GuideFormer: Transformers for image guided depth completion," in Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA, 2022, pp. 6250-6259. [Online]. Available: https://ieeexplore.ieee.org/document/9879041

[11] F. Zhang, V. Bazarevsky, A. Vakunov, A. Tkachenka, G. Sung, C.-L. Chang, and M. Grundmann, "MediaPipe hands: On-device real-time hand tracking," arXiv preprint arXiv:2006.10214, 2020. [Online]. Available: https://arxiv.org/abs/2006.10214

[12] T. Dwyer, K. Marriott, T. Isenberg, K. Klein, N. Riche, F. Schreiber, W. Stuerzlinger, and B. H. Thomas, "Immersive analytics: An introduction," in Immersive Analytics, Lecture Notes in Computer Science, vol. 11190, Springer, Cham, 2018, pp. 1-23. [Online]. Available: https://doi.org/10.1007/978-3-030-01388-2_1

[13] M. R. Luo, G. Cui, and B. Rigg, "The development of the CIE 2000 colour-difference formula: CIEDE2000," Color Research and Application, vol. 26, no. 5, pp. 340-350, Oct. 2001. [Online]. Available: https://doi.org/10.1002/col.1049

[14] P. N. Huu, P. D. L. Hong, D. D. Dang, B. V. Quoc, C. N. L. Bao, and Q. T. Minh, "Proposing hand gesture recognition system using MediaPipe holistic and LSTM," in Proc. IEEE International Conference on Advanced Technologies for Communications (ATC), Da Nang, Vietnam, 2023, pp. 433-438. [Online]. Available: https://ieeexplore.ieee.org/document/10318885

[15] Y. Li, Z. Yu, C. Choy, C. Xiao, J. M. Alvarez, S. Fidler, C. Feng, and A. Anandkumar, "VoxFormer: Sparse voxel transformer for camera-based 3D semantic scene completion," in Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Vancouver, Canada, 2023, pp. 9087-9098. [Online]. Available: https://openaccess.thecvf.com/content/CVPR2023/html/Li_VoxFormer_Sparse_Voxel_Transformer_for_Camera-Based_3D_Semantic_Scene_Completion_CVPR_2023_paper.html

[16] N. Zhang, F. Nex, G. Vosselman, and N. Kerle, "Lite-Mono: A lightweight CNN and transformer architecture for self-supervised monocular depth estimation," in Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Vancouver, Canada, 2023, pp. 18537-18546. [Online]. Available: https://openaccess.thecvf.com/content/CVPR2023/html/Zhang_Lite-Mono_A_Lightweight_CNN_and_Transformer_Architecture_for_Self-Supervised_Monocular_CVPR_2023_paper.html

[17] T. Kim and K. Behdinan, "Advances in machine learning and deep learning applications towards wafer map defect recognition and classification: A review," Journal of Intelligent Manufacturing, vol. 34, no. 8, pp. 3215-3247, Oct. 2023. [Online]. Available: https://doi.org/10.1007/s10845-022-01994-1

[18] J. Yan, Y. Sheng, and M. Piao, "Semantic segmentation-based wafer map mixed-type defect pattern recognition," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 42, pp. 4065-4074, 2023. [Online]. Available: https://doi.org/10.1109/TCAD.2023.3274958

[19] S. G. Hart and L. E. Staveland, "Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research," in Human Mental Workload, P. A. Hancock and N. Meshkati, Eds. Amsterdam: Elsevier Science Publisher B.V., North-Holland, 1988, pp. 139-183. [Online]. Available: https://doi.org/10.1016/S0166-4115(08)62386-9

[20] M. Zheng, D. Lillis, and A. G. Campbell, "Current state of the art and future directions: Augmented reality data visualization to support decision-making," Visual Informatics, vol. 8, no. 1, pp. 80-105, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2468502X24000202