

INSAM

JOURNAL OF CONTEMPORARY MUSIC, ART AND TECHNOLOGY



Splaxtraction: Gaussian Splats as Extractive Vision and the Environmental Cost of Volumetric Imaging

Tom Milnes

INSAM Journal of Contemporary Music, Art and Technology

No. 16, July, pp. 26–39.

<https://doi.org/10.51191/issn.2637-1898.2026.9.16.26>

TOM MILNES

Centre for Blended Realities, Falmouth University, United Kingdom

SPLAXTRACTION: GAUSSIAN SPLATS AS EXTRACTIVE VISION AND THE ENVIRONMENTAL COST OF VOLUMETRIC IMAGING

KEYWORDS

Gaussian splatting;
photogrammetry; machine
vision; media archaeology;
extractivism; e-waste;
dataset ethics; eco-critique

ORIGINAL SCIENTIFIC PAPER

Received: April 1, 2026

Accepted: July 3, 2026

DOI: [10.51191/issn.2637-
1898.2026.9.16.26](https://doi.org/10.51191/issn.2637-1898.2026.9.16.26)

ABSTRACT

This article examines Gaussian splatting and adjacent volumetric imaging practices through media archaeology and media ecology. It asks how the apparent smoothness of probabilistic spatial images depends on less visible operations of capture, dataset circulation, cloud computation, minerals, and electronic waste. I propose the concept of splaxtraction to name this double movement: volumetric images extract environments into transferable computational resources, while also being produced through extractive infrastructures of labour, energy, hardware, and platform power. Rather than treating Gaussian splats as a neutral improvement in visual realism, the article situates them within longer histories of photogrammetry, operational imaging, media geology, and dataset critique. The contribution is a critical vocabulary for reading volumetric images as infrastructural assemblages whose seams, uncertainties, and material supports should remain visible within eco-critical artistic and curatorial practice.



tom.milnes@falmouth.ac.uk



<https://orcid.org/0000-0002-5181-2063>

Introduction: From "Reality Capture" to Splaxtraction

With the development of 3D scanning technologies, a new type of imagery has emerged: a state in which the image is a layered, three-dimensional construction, assembled through the mediation of photographic images. More recently, probabilistic volumetric imaging has extended this condition further. Gaussian splats, neural radiance fields, and adjacent forms of machine vision promise immersion, realism, and navigable depth in ways that are regularly marketed as high-fidelity but also frictionless, transparent capturing systems. Yet what is at stake is not merely an aesthetic refinement. A larger technical and political formation is taking shape in which environments, objects, and bodies are increasingly converted into transferable, optimisable visual assets.

This article approaches Gaussian splatting as a media archaeological and media ecological problem. Its method is conceptual rather than technical benchmarking: it reads volumetric image systems across their historical lineages, capture processes, dataset economies, computational infrastructures, and material afterlives. This framing clarifies that the object of analysis is not only the rendered splat, but the wider apparatus that makes such an image possible.

The language surrounding these systems still bears the mark of the older commercial rhetoric of "Reality Capture," a term developed to market 3D-scanning tools and software suites by promising that the physical world could be efficiently transformed into a manipulable digital asset (Autodesk 2017). Both "reality capture" and "volumetric imaging" are alluring because they suggest that the image has become less mediated than photography, more immediate, more truthful, more objective. My concern, developed earlier in relation to photogrammetry, is that such rhetoric belies the computational assemblage fundamental to the image's being. Photogrammetry was never a singular image but a layered and navigable composite. It relied on the correlation of multiple photographs, the production of point clouds, the generation of surfaces, and the projection of texture maps. It therefore made sense to understand it as a composite image produced through the amalgamation of spatial and visual data within a computational environment rather than through a direct indexical relation to a scene (Milnes 2022).

This article extends that argument into the present condition of Gaussian splatting and adjacent volumetric methods. The research converges on a need to define these images, not merely for their visual distinction, but for their material and political nature as an assemblage – not as immaterial interfaces, but as extracts from a wider logistical field. "Splaxtraction" attends to the image condition as an assembly of datasets, sensors, protocols, mineral resources, labour, platforms, and institutions. It insists that volumetric imaging is inseparable from cloud infrastructures, training data, and hardware ecologies. In that sense, splaxtraction names both a technical process and an analytic. It describes how the world is made available as volumetric image, and how that availability depends upon extraction at several registers at once.

The term *splaxtraction* is therefore intended to do more than rename Gaussian splatting. It distinguishes this article's contribution from adjacent concepts. From operational images it takes the sense that images increasingly act within computational systems; from media geology it takes the insistence that digital media are mineral and infrastructural; from dataset critique it takes the claim that training and evaluation archives actively classify the world. Splaxtraction brings these strands together at the level of volumetric imaging, where space itself becomes an extractable, optimisable, and reusable image resource.

The argument proceeds in three orders. First, it traces a lineage from photogrammetry to probabilistic volumetric images in order to show that Gaussian splats continue, rather than replace, earlier composite and operational forms. Second, it examines capture ethics, unconsented scanning, and dataset economies, asking how environments become transferable computational resources. Third, it turns to the material footprint of the pipeline, connecting splats to cloud computation, GPU-intensive optimisation, minerals, obsolescence, and e-waste. The conclusion then returns to artistic and curatorial practice by proposing a compressed eco-critical strategy: to keep provenance, workflow, artefacts, and computational uncertainty visible rather than smoothing them into a rhetoric of frictionless capture.

Order One:

From the Photogrammetric Image to Probabilistic Volumetric Images

Photogrammetry is a 3D imaging technique that relies on structure-from-motion principles: the process of detecting and reconstructing spatiality from serialised images captured while moving around an object or environment (Ullman 1979). What emerged in the digital paradigm was not merely a surveying tool but a distinct image-form. The digital photogrammetric image depended on image-recognition algorithms, point-cloud estimation, mesh generation, and texture projection. It was therefore already computational in a stronger sense than a conventional photograph. It belonged less to a history of singular lens-based images than to a broader field of automated composites and networked images (Rubinstein and Sluis 2008; Schröter 2009).

A key insight of the earlier research was that photogrammetry's aesthetic specificity became most visible at the point of error. Hito Steyerl's account of "fractional space" remains foundational to 3D aesthetics. The 3D image, she argued, is not a complete volumetric truth but an unstable space hovering between two and three dimensions, one in which missing data are assigned a body, surfaces fold awkwardly through space, and the image's blind spots become material (Steyerl 2017, 191–205). Such artefacts mattered because they disclosed the machine's operations. Holes, stretched textures, blank planes, and spikes were not merely defects to be cleaned up; they were evidence of the algorithm's struggle to reconcile multiple flat images into a navigable environment.

Gaussian splatting inherits this composite condition but redistributes its visibility. The formulation presented by Kerbl et al. describes the scene as a collection of 3D Gaussians whose position, colour, opacity, and anisotropic covariance are jointly optimised for real-time rendering (Kerbl et al. 2023). In place of a clearly bounded surface mesh, the scene becomes a probabilistic field. This shift has important onto-epistemological consequences. If photogrammetry already challenged the impassive transparency of the photographic apparatus (Flusser 2000), Gaussian splats push further by making the image depend not on explicit surface reconstruction but on the continuous modulation of density and radiance. The output may appear smoother and more seamless; the computational uncertainty has not vanished but has been regularised (Kerbl et al. 2023).

Gaussian splats and neural radiance fields are relocating photogrammetry's fractionality into the model's interior: instead of torn meshes and visible spikes, low-confidence regions are regularised into a probabilistic field that presents itself as seamless even when assembled from sparse or contradictory data. The image becomes a probabilistic negotiation between evidence and dataset understanding, rather than a record of what stood before the lens. When thinking about the compositional stakes of 3DGS, each splat is not a miniature object but a weighted visual proposition: a density-bearing, colour-bearing event whose contribution only emerges in relation to neighbouring splats, camera position, and visibility calculations. Composition therefore happens less as the assembly of stable surfaces than as the continual balancing of inference. The scene is held together by optimisation. What appears to the viewer as volumetric continuity is produced by innumerable local decisions about opacity, covariance, and blend. In this sense, the Gaussian-splatted scene is closer to an atmospheric model of plausibility than to the harder edges of mesh-based reconstruction.

Critical writing gathered around Volumetric Regimes helps name what is at stake in that shift. Rocha and Snelting describe volumetrics not simply as a technical format but as a regime through which bodies, spaces, and presences become calculable and governable (Rocha and Snelting 2022). Helen V. Pritchard's call for "clumsy volumetrics" is especially useful because it refuses the normative smoothness of capture systems, treating friction, misalignment, and non-normative embodiment as politically revealing rather than technical mishaps (Pritchard 2022). Read this way, the foggy voids and crystalline fissures of failed splat optimisation are not mere defects. They are moments when quantified presence falters and the image's infrastructural condition briefly shows through.

The contemporary field has also expanded through text-to-3D systems such as DreamFusion and Point-E. These systems do not begin from direct sensory capture of a scene. DreamFusion uses a pretrained 2D diffusion model as a precursor for optimising a 3D representation through what its authors call probability density distillation (Poole et al. 2022). Point-E generates an image and then a point cloud conditioned upon it, trading fidelity for speed and practical accessibility (Nichol et al. 2022). These tools matter because they intensify a longer shift away from capture as evidence and toward prediction as image production. The resulting object is still navigable and volumetric, but it is assembled from statistical residues rather than solely from a measured scene.

Recent text-to-3D and world-generation systems intensify this shift from capture toward prediction. Systems such as World Labs's Marble present navigable 3D worlds as outputs generated from text, image, video, or layout prompts (World Labs 2025; 2026). Read through dataset critique, such systems show how volumetric imagery increasingly depends on trained expectations: the scene is not simply reconstructed from evidence but anticipated through modelling constraints, prior examples, and probabilistic completion.

For this reason, Gaussian splats should not be understood as a clean break from photogrammetry. They are better described as a continuation of the same drive toward immersive, networked, and composite image environments, but under conditions in which inference has become more powerful, more hidden, and more infrastructurally intensive. The once-visible rupture of photogrammetric meshes has become the latent smoothness of probabilistic modelling. What was formerly visualised as a mismatch between mesh and texture, now persists as a more complex mist, fuzz or crystalline interference – but no less decisive politics of optimisation. The image is no longer simply captured; it is computed into plausibility.

This matters aesthetically because the smoothing of error risks renewing the ideology of transparency. If photogrammetry at least sometimes betrayed its own computational seams, Gaussian splats can appear uncannily natural while depending on ever more elaborate technical mediation. The image is what Harun Farocki termed an “operational image” in an expanded sense: it acts within a computational system whose decisive operations are not necessarily intended for human interpretation even when the final render is visually seductive (Farocki 2004; Parikka 2023). Splaxtraction begins by resisting this renewed transparency. It asks what a volumetric image is made of, what infrastructures it presupposes, and what social and ecological relations are compressed into its apparent immediacy.

Order Two:

The Ethics of Capture: Unconsented Scanning, Dataset Economies, and Transferable Environments

Any eco-critical account of volumetric imaging must begin with capture. The problem is not exhausted by the final image; it starts with the acquisition of data and with the conditions under which space becomes available to be imaged. The image detached itself from single-perspective photography and became navigable, architectural, and composite, yet this apparent liberation from the frame also increased the appetite for spatial acquisition. Ariel Caine's account is useful because it shows how “spatial photography”, as Caine terms 3D scans, can support counter-dominant practices when multiple authors contribute to a navigable record of threatened environments (Caine 2019). But the same technical affordance can be recruited for very different ends. The conver-

sion of environments into volumetric assets also enables extraction, transfer, surveillance, and re-use on a scale that conventional documentary capture did not.

The first ethical problem, then, is unconsented scanning. When streets, interiors, ruins, workplaces, or ecologically damaged sites are volumetrically captured, the question is not only who took the image, but who now controls the model, the derivatives, and the data traces that persist beyond the original act of recording. Photogrammetry had already loosened the relation between one scene and one image by allowing models to be assembled from multiple photographs taken at different times and from different authors. In networked practice this can be politically enabling. Yet it also means that lived environments can be converted into modular training material, stock assets, or synthetic scenes that travel far beyond the original social context of capture.

This is where dataset economies enter the discussion. Computer vision depends upon datasets: curated, labelled image collections used to train, test, and evaluate machine learning systems (Deng et al. 2009). Lisa Gitelman and Virginia Jackson write in *"Raw Data" Is an Oxymoron* that "when phenomena are variously reduced to data, they are divided and classified, processes that work to obscure—or as if to obscure—ambiguity, conflict, and contradiction" (Gitelman and Jackson 2013, 9). Gitelman and Jackson's account is particularly important because it shows that datasets are not neutral collections but operations of sorting, annotation, and alignment. They encode what counts as a recognisable object, a meaningful feature, or a stable category. In the context of volumetric imaging, the issue is not merely that datasets contain images; it is that they sediment assumptions about the world that are then reactivated in capture, reconstruction, and generation. This is one way of understanding a practice counter to Caine's counter-dominant spatial photography: a form of dataset laundering in which scenes, textures, objects, and bodies are moved from one technical context to another, cleaned of their local conditions, and reintroduced as apparently generalisable visual material.

This is the methodological hinge between capture ethics and media ecology: a volumetric scene is not only gathered, but formatted for future computational action. Once abstracted as a scan, splat, point cloud, mesh, or training asset, a situated environment can be detached from the conditions of consent, access, and locality that originally shaped it.

The ethical stakes become clearer when we consider bias. Gebru and Buolamwini's work on gender and racial bias in facial-analysis systems demonstrated that computational image systems do not simply record the visible world; they reproduce the inequalities embedded in their training and evaluation protocols (Gebru and Buolamwini 2019). In volumetric systems the problem is not only facial recognition but the wider calibration of what environments and bodies are easy to model, what conditions count as noise, and what kinds of spaces become legible as normative. Repetitive, reflective, translucent, unstable, or non-anthropocentric environments have long posed difficulties for photogrammetry. Forests, water, blossom, mist, quarries, and light-based phenomena generate confusion for image-recognition algorithms because their appearance changes, repeats, or exceeds the categories on which the system relies. What appears technically as "bad data" is often a sign of a much larger political sorting of visibility.

These conditions intensify in generative text-to-3D systems. DreamFusion and Point-E can produce objects that look usable within game engines or 3D software, but they do so by inheriting representational habits from large image-text corpora and existing model repositories (Poole et al. 2022; Nichol et al. 2022). The resulting volumetric object is not innocent. It is a synthetic condensation of catalogued images, annotations, taxonomies, and aesthetic conventions. Because the technical system treats these sources as generalisable, the social and legal questions of provenance can be obscured. The object appears new while remaining indebted to an extractive archive of prior capture. If probabilistic volumetric images are assembled through computational suppositions, shaped by expectations derived from training data, then capture can no longer be understood as a discrete act that ends with the scan. A volumetric model remains open to reclassification, retexturing, interpolation, and transfer into adjacent systems. It can become a training example, a game-engine asset, a synthetic environment, or a world-model. The scan is thus less a final document than an infrastructural handoff.

The ethics of capture therefore cannot be reduced to privacy on the one hand or copyright on the other, though both matter. The deeper issue is that volumetric imaging converts situated environments into transferable computational resources. The question becomes: under what conditions can an environment be scanned, abstracted, circulated, and re-used? Splaxtraction names this conversion as an extractive act. It insists that the volumetric image is not simply a representation of a place, but part of a logistical process in which places are made available to platforms, models, and institutions that can act upon them again.

Malevė's writing complements this account by treating datasets as sites of arbitration rather than neutral archives. In "On the Data Set's Ruins," the dataset is not approached simply as a collection of images, but as an arrangement produced through alignment, sorting, annotation, and labour (Malevė 2021). That insight maps closely onto volumetric capture: what is transferred from one context to another is not only visual material, but a regime of selection concerning what counts as object, background, noise, stable surface, or useful scene.

Order Three:

Cloud Footprints: GPU Dependence, Minerals, and the E-Waste Tail of Volumetric Imaging

If the first extractive register concerns the conversion of environments into data, the second concerns the material systems required to process, store, and render volumetric images. Against the fantasy of dematerialisation, media ecology has repeatedly insisted that digital culture is irreducibly geological. Jussi Parikka's research has consistently exposed that digital media are not weightless signs but condensations of minerals, infrastructures, toxic afterlives, and damaged environments (Parikka 2015a; 2015b; 2016).

Jennifer Gabrys likewise shows that digital rubbish is not a marginal by-product but intrinsic to media culture's cycles of innovation and obsolescence (Gabrys 2011). Sean Cubitt extends this argument by treating digital technologies as finite media whose environmental costs cannot be bracketed from their cultural effects (Cubitt 2017).

Volumetric imaging belongs squarely within this condition. Photogrammetry already relied on digital cameras, network transfer, software platforms, and cloud processing. The software's upload and processing routines made visible how much infrastructural support is hidden behind the apparently simple act of making a 3D model. The image was not merely a function of optics; it was upheld by energy systems, network infrastructure, server-side computation, and the assumptions of a cloud service designed elsewhere.

The Gaussian splatting standard process remains tightly coupled to GPU-accelerated optimisation and, increasingly, cloud computation. The original developers of Gaussian Splats, GraphDeco-INRIA make clear that canonical 3DGS is not simply a lightweight extension of older imaging workflows but a method structured around substantial accelerator hardware (GraphDeco-INRIA 2023). Comparative evidence from Salime Faizi's study of photogrammetry and Gaussian splatting across mobile, desktop, and cloud platforms further suggests that this hardware dependency is not incidental but operational, noting that relative to ordinary photogrammetry workflows, 3DGS often entails a more explicitly AI-style training stack and can push practice toward cloud GPUs (Faizi 2025). This does not prove that every splat workflow is always more carbon-intensive than every photogrammetric one, but it does show that 3DGS materially deepens dependence on high-memory GPUs, CUDA software stacks, and platform computation. In that sense, the apparent immediacy of the splat image is underwritten by an intensive computational infrastructure whose environmental costs are displaced into power-hungry processing systems and cloud hardware rather than removed from the image itself.

Approached through media archaeology and media ecology, photogrammetric and volumetric images cannot be isolated from the wider infrastructures that support them. The splat on screen is only the front-end manifestation of a chain that passes through sensors, batteries, data transfer, cloud processing, and specialised graphics hardware. Lisa Parks's call to think of media infrastructures as physical objects in "Stuff You Can Kick: Toward a Theory of Media Infrastructures" resists the abstraction of digital culture into mere interface (Parks 2015). The phrase insists on palpable supports: buildings, cables, antennas, cooling systems, and the territorial arrangements that keep signals and files in motion. Rob Holmes's "A Preliminary Atlas of Gizmo Landscapes" similarly tracks how even a single apparently trivial action on a smartphone tethers together mines, factories, server farms, and communication networks (Holmes 2010). These accounts remind us of a contrasting reality, with rhetoric around Gaussian splatting often presented as a lightweight visual convenience: a clever rendering method, a smoother viewport, a faster way to move from capture to immersion. Media ecology reminds us that no such convenience is immaterial.

In this respect, volumetric imaging should be understood as an infrastructural compression. It compresses very large and spatially distributed systems into a small zone of interface legibility. The user sees the rendered scene, perhaps the timeline and the export button, but not the logistics and labour underneath. Splaxtraction names that asymmetry. It insists that the image must be read in relation to the landscapes of extraction and transmission that make its apparent immediacy possible.

Even where splats are more efficient than some neural radiance field methods, the volumetric image remains tied to GPU-based pipelines and hardware ecologies. Studies of 3D digitising technologies repeatedly stress that acquisition, processing, rendering, and storage all depend upon a coordinated chain of devices, software environments, and technical expertise rather than a single capture moment (Triviño-Tarradas, García-Molina, and Rojas-Sola 2024). Creative adoption of such tools often arrives under the sign of accessibility and real-time experimentation, yet that accessibility is inseparable from the churn of graphics hardware, cooling, storage, and energy demand.

To discuss only carbon here would be too narrow. The environmental cost of volumetric imaging is distributed across extraction, manufacture, energy consumption, and disposal. Minerals are dug from the ground; devices are manufactured under asymmetrical labour conditions; data centres and cloud services consume power and water; and obsolete equipment is displaced into longer geographies of waste (National Research Council 2008; Grossman 2006). John Vidal's reporting on toxic e-waste dumping offers a blunt reminder that discarded devices do not disappear but reappear in other territories under dangerous conditions (Vidal 2013). Rob Nixon's concept of slow violence is important because these harms are often dispersed in time and space rather than registered as a single visible event (Nixon 2013).

This slow violence is built into the aesthetics of accelerated creative tooling. The promise of frictionless imaging technologies encourages more scanning, more processing, and more rapid obsolescence of devices deemed inadequate for current workflows. What is new in Gaussian splatting is not that the image has suddenly become extractive, but that a particularly persuasive form of volumetric realism now coincides with an intensified demand for computational throughput. The cultural pressure to keep up with new pipelines, new real-time standards, and new model formats feeds directly into hardware churn. The image therefore has an e-waste tail. Every apparently seamless scene is connected to a much rougher geography of batteries, rare-earth supply chains, server farms, cooling systems, and discarded electronics.

Cloud discourse remains especially ideological here. James Bridle's observation that "the cloud doesn't just have a shadow; it has a footprint" cuts through the rhetoric of abstraction with precision (Bridle 2018, 7). Volumetric imaging should be read through this footprint. The model in the viewport is not only a visual object; it is the front-end manifestation of a very material chain of extraction and logistics. Splaxtraction draws attention to that chain: it shifts attention from the scene represented to the resources consumed in making such representation possible at all.

The environmental question is therefore not an external supplement to image theory; it is part of what the volumetric image is. To read a splat ecologically is to read its

interface, training process, storage, hardware, and disposal as one distributed image system.

Conclusion

Gaussian splatting has quickly been framed as a technical breakthrough in realism: faster, smoother, and more convincing than earlier volumetric techniques. This article has argued that such a framing remains inadequate if it isolates visual performance from the wider conditions of capture. Read through media archaeology and media ecology, splats do not simply represent space in a new way; they reorganise the terms under which space becomes available as image. They extend the long history of composite and operational imaging into a regime in which environments are rendered as transferable, optimisable, and increasingly predictive visual assets. In that sense, Gaussian splatting is not an innocent aesthetic upgrade. It is part of a broader, systemic extractive formation linking scene acquisition, dataset circulation, cloud infrastructures, mineral resources, and hardware churn.

The concept of splaxtraction has been proposed to make that formation apparent. It names a double movement. First, it identifies how volumetric images extract from the world: by scanning environments, abstracting them into computational fields, and folding them into technical systems that enable reuse, prediction, and circulation. Second, it insists that these images are themselves composed through extraction: of labour, energy, logistics, rare-earth minerals, and institutional power. To analyse the aesthetics of a Gaussian splat only at the level of appearance is therefore to miss the multiple resource regimes condensed within it. Splaxtraction asks that volumetric images be read as infrastructural assemblages as much as visual artefacts.

The argument also has consequences for artistic and curatorial practice. A media archaeological strategy for volumetric imaging would treat the artefact, seam, void, density bloom, and low-confidence region as evidence of the apparatus rather than as defects to be erased. Here the tradition of glitch as critique remains a useful method. As Rosa Menkman argues, error can reveal the conventions and protocols that normally remain hidden within technical systems (Menkman 2011). In the context of Gaussian splatting, this does not mean aestheticising failure for its own sake. It means using uncertainty to disclose the capture conditions, optimisation procedures, platform dependencies, and material infrastructures that polished volumetric images often conceal.

For artists, researchers, and curators working with volumetric imaging, the task is therefore not necessarily refusal but reorientation. This could involve documenting capture permissions and provenance, naming the software and hardware conditions of a workflow, retaining artefacts rather than automatically cleaning them, extending the life of existing equipment, and resisting the assumption that every environment should become a portable visual asset. Such practices do not resolve the contradictions of volumetric media, but they can make those contradictions publicly apprehensible.

If contemporary machine vision increasingly smooths the violence of extraction into the rhetoric of immersion, eco-critical practice has a counter-task: to render the costs of volumetric imaging visible again. Splaxtraction, in this sense, is less a final diagnosis than a working method for staying with those costs and for refusing to mistake technical fluency for ecological or political innocence.

List of References

- Autodesk.** 2017. "ReCap Pro: Reality Capture Software Connecting the Physical World to the Digital". Accessed March 28, 2026. <https://www.autodesk.co.uk/products/recap/overview?plc=RECAP&term=1-YEAR&support=ADVANCED&quantity=1#internal-link-workflows>.
- Bridle,** James. 2018. *New Dark Age: Technology, Knowledge and the End of the Future*. Verso.
- Caine,** Ariel. 2019. "Granular Realism / Ontology and Counter-Dominant Practices of Spatial Photography". PhD diss., Goldsmiths, University of London.
- Cubitt,** Sean. 2017. *Finite Media: Environmental Implications of Digital Technologies*. Duke University Press.
- Deng,** Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. "ImageNet: A Large-Scale Hierarchical Image Database". In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–55. <https://doi.org/10.1109/CVPR.2009.5206848>.
- Dick,** Philip K. 1953. "The Preserving Machine". *The Magazine of Fantasy and Science Fiction*, June 1953.
- Faizi,** Salime. 2025. "Photogrammetry vs. Gaussian Splatting on Mobile, Desktop and Cloud Platforms". Thesis, Hochschule Darmstadt.
- Farocki,** Harun. 2004. "Phantom Images". *Public* 29: 12–22. <https://public.journals.yorku.ca/index.php/public/article/view/30354>.
- Flusser,** Vilém. 2000. *Towards a Philosophy of Photography*. Translated by Anthony Mathews. Reaktion Books.
- Gabrys,** Jennifer. 2011. *Digital Rubbish: A Natural History of Electronics*. University of Michigan Press.
- Gebru,** Timnit, and Joy Buolamwini. 2019. "Gender Shades". MIT Media Lab / Ars Electronica. Accessed February 25, 2021. <https://ars.electronica.art/center/de/gendershades/>.
- Gitelman,** Lisa, and Virginia Jackson. 2013. "Introduction". In *"Raw Data" Is an Oxymoron*, edited by Lisa Gitelman, 1–14. MIT Press.
- GraphDeco-INRIA.** 2023. "3D Gaussian Splatting for Real-Time Radiance Field Rendering". GitHub repository. Accessed March 30, 2026. <https://github.com/graphdeco-inria/gaussian-splatting>.
- Grossman,** Elizabeth. 2006. *High Tech Trash: Digital Devices, Hidden Toxics, and Human Health*. Island Press.

- Holmes, Rob.** 2010. "A Preliminary Atlas of Gizmo Landscapes". *Mammoth*, April 1, 2010. <https://m.ammoth.us/blog/2010/04/a-preliminary-atlas-of-gizmo-landscapes/>.
- Kerbl, Bernhard, Georgios Kopanas, Thomas Leimkühler, and George Drettakis.** 2023. "3D Gaussian Splatting for Real-Time Radiance Field Rendering". *ACM Transactions on Graphics* 42 (4): 1–14. <https://doi.org/10.1145/3592433>.
- Malevé, Nicolas.** 2021. "On the Data Set's Ruins". *AI & Society* 36 (4): 1117–31. <https://doi.org/10.1007/s00146-020-01093-w>.
- Menkman, Rosa.** 2011. *The Glitch Moment(um)*. Institute of Network Cultures.
- Milnes, Tom.** 2022. "Spatially-Immersive Networked Composites: A Media Archaeology of the Photogrammetric Image through Glitch Practice". PhD diss., University of the Arts London and Falmouth University. <https://repository.falmouth.ac.uk/id/eprint/5363>.
- National Research Council.** 2008. *Minerals, Critical Minerals, and the U.S. Economy*. National Academies Press.
- Nichol, Alex, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen.** 2022. "Point-E: A System for Generating 3D Point Clouds from Complex Prompts". arXiv. <https://doi.org/10.48550/arXiv.2212.08751>.
- Nixon, Rob.** 2013. *Slow Violence and the Environmentalism of the Poor*. Harvard University Press.
- Parks, Lisa.** 2015. "Stuff You Can Kick: Toward a Theory of Media Infrastructures". In *Between Humanities and the Digital*, edited by Patrik Svensson and David Theo Goldberg, 355–73. MIT Press.
- Parikka, Jussi.** 2015a. *The Anthrobscene*. University of Minnesota Press.
- Parikka, Jussi.** 2015b. *A Geology of Media*. University of Minnesota Press.
- Parikka, Jussi.** 2016. *A Slow, Contemporary Violence: Damaged Environments of Technological Culture*. Sternberg Press.
- Parikka, Jussi.** 2023. *Operational Images: From the Visual to the Invisual*. University of Minnesota Press.
- Poole, Ben, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall.** 2022. "DreamFusion: Text-to-3D Using 2D Diffusion". arXiv. <https://doi.org/10.48550/arXiv.2209.14988>.
- Pritchard, Helen V.** 2022. "Clumsy Volumetrics". In *Volumetric Regimes: Material Cultures of Quantified Presence*, edited by Jara Rocha and Femke Snelting. Open Humanities Press. https://volumetricregimes.xyz/index.php?title=Clumsy_Volumetrics.
- Rocha, Jara, and Femke Snelting, eds.** 2022. *Volumetric Regimes: Material Cultures of Quantified Presence*. DATA Browser 08. Open Humanities Press.
- Rubinstein, Daniel, and Katrina Sluis.** 2008. "A Life More Photographic: Mapping the Networked Image". *Photographies* 1 (1): 9–28. <https://doi.org/10.1080/17540760701785842>.
- Schröter, Jens.** 2009. *3D: History, Theory, and Aesthetics of the Transplane Image*. New York: Bloomsbury.
- Steyerl, Hito.** 2017. "Ripping Reality: Blind Spots and Wrecked Data in 3D". In *Duty Free Art: Art in the Age of Planetary Civil War*, 191–205. Verso.

- Triviño-Tarradas**, Paula, Diego Francisco García-Molina, and José Ignacio Rojas-Sola. 2024. "Impact of 3D Digitising Technologies and Their Implementation". *Technologies* 12 (12): 260. <https://doi.org/10.3390/technologies12120260>.
- Ullman**, Shimon. 1979. "The Interpretation of Structure from Motion". *Proceedings of the Royal Society of London. Series B, Biological Sciences* 203 (1153): 405–26. <https://doi.org/10.1098/rspb.1979.0006>.
- Vidal**, John. 2013. "Toxic 'E-Waste' Dumped in Poor Nations, Says United Nations". *The Guardian*, December 14, 2013. Accessed February 18, 2021. <https://www.theguardian.com/global-development/2013/dec/14/toxic-ewaste-illegal-dumping-developing-countries>.
- World Labs**. 2025. "Marble: A Multimodal World Model". November 12, 2025. Accessed March 27, 2026. <https://www.worldlabs.ai/blog/marble-world-model>.
- World Labs**. 2026. "About". Accessed March 27, 2026. <https://www.worldlabs.ai/about>.

Declaration of AI Assistance

AI-assisted editing was used to support copy-editing, formatting, and bibliographic standardisation. The author reviewed the manuscript and remains responsible for its argument, citations, and final content.

SPLAXTRACTION: GAUSSIAN SPLATS AS EXTRACTIVE VISION AND THE ENVIRONMENTAL COST OF VOLUMETRIC IMAGING (summary)

This article develops the concept of splaxtraction as a critical vocabulary for analysing Gaussian splatting and adjacent volumetric imaging practices. It argues that the apparent smoothness and immediacy of contemporary probabilistic spatial images should not be understood simply as a technical improvement in realism. Instead, these images need to be read as infrastructural assemblages that depend on capture protocols, dataset economies, cloud computation, specialist hardware, minerals, labour, and electronic waste. The article therefore situates Gaussian splatting within a longer media archaeological lineage that includes photogrammetry, operational imaging, media geology, and critical approaches to datasets. Its central claim is that volumetric images extract from the world by converting environments into transferable computational resources, while also being produced through material infrastructures that are themselves extractive.

The first section traces the movement from photogrammetric composites to probabilistic volumetric image fields. Photogrammetry already unsettled the idea of a singular photographic image because it produced spatial models through the correlation of many images, point-cloud estimation, mesh generation, and texture projection. Gaussian splatting continues this composite condition but changes how its uncertainty

appears. Whereas photogrammetric errors often became visible as holes, stretched textures, spikes, blank surfaces, or phantom planes, splats can regularise uncertainty into smoother fields of density, covariance, opacity, and radiance. The article argues that this shift risks renewing an ideology of transparent capture, even as the image becomes more deeply computational. The splatted scene is therefore read not as a stable record of a place, but as a field of optimisation in which evidence, inference, and prior expectations are continually balanced.

The second section examines capture ethics, unconsented scanning, and dataset economies. It argues that volumetric imaging does not only represent places but turns situated environments into transferable computational resources. Once a scene is abstracted as a scan, splat, mesh, point cloud, or training asset, it can move beyond its original context of consent, access, and locality. Dataset critique is therefore essential to the argument, because datasets are not neutral containers of visual material but systems of sorting, annotation, selection, and generalisation. The article uses this point to show how capture can become a logistical handoff: a place is not simply documented, but formatted for future computational action, reuse, prediction, and circulation.

The third section turns to the environmental and infrastructural cost of volumetric imaging. Drawing on media ecology and media geology, it connects Gaussian splatting to GPU-based optimisation, cloud processing, data storage, mineral extraction, hardware obsolescence, and e-waste. The article cautions against reducing the environmental question to carbon accounting alone. Instead, it reads volumetric images as distributed image systems whose costs are dispersed across extraction, manufacture, energy use, cooling, logistics, and disposal. The environmental footprint of a splat is therefore not external to its aesthetics; it is part of the image condition that makes its apparent immediacy possible.

The conclusion proposes an eco-critical artistic and curatorial strategy that does not require a full practice-based case study. It suggests that artefacts, seams, voids, density blooms, low-confidence regions, workflow traces, and provenance information should be kept visible rather than erased. In this sense, splaxtraction is offered not as a rejection of volumetric media, but as a method for keeping the political and ecological costs of machine vision perceptible. The article ultimately asks artists, researchers, and curators to resist frictionless capture narratives and to treat uncertainty as evidence of the apparatus. This reorientation gives the paper its practical implication: eco-critical work with volumetric media should disclose capture conditions, name technical dependencies, prolong existing hardware where possible, and avoid presenting immersive spatial images as if they were detached from the social and ecological systems that make them operable, visible, and available for critical interpretation within art, research, and curatorial contexts.