Radiance Fields and Spatial Artificial Intelligence: Synthetic Landscape Modeling with SHARP

Maximilian Schob[1], Jörg Rekittke[2]

[1]University College Dublin (UCD), Dublin/Ireland · maximilian.schob@ucd.ie
[2]University College Dublin (UCD), Dublin/Ireland

Abstract: Spatial artificial intelligence is driving a quiet revolution in inference-based landscape modeling. In landscape architecture, established forms of digital representations, based on discrete data structures and geometries, favor stable surfaces and bounded objects at the expense of the atmospheric, temporal, and porous qualities through which landscapes are encountered. Radiance fields, including representations structured through 3D Gaussian splatting (3DGS) and Neural Radiance Fields (NeRF), encode landscapes as continuous volumetric fields of light, generating navigable 3D environments from image-based observations.

This paper positions radiance fields as a distinct and emerging category of landscape representations and examines their relevance for landscape architecture through SHARP (Single-image High-Accuracy Real-time Parallax), a novel model by Apple that synthesizes metric-scale 3D Gaussian radiance fields from a single image, preserving absolute spatial dimensions relevant to bodily experience and design reasoning. Trained on large-scale synthetic and real-world data, SHARP infers spatial structure, scale, and visibility from minimal input, collapsing fragmented sensing and modeling processes into a unified inference framework.

Infrared imagery captured in Dublin (Ireland), including at the National Botanic Gardens and Killiney Hill Park, serves as input to generate a series of 3D Gaussian reconstructions. In this context, SHARP enables high-fidelity, near-instantaneous synthetic landscape modeling, providing an accessible framework for design research and education. In addition to our model explorations, the approach is further validated through pedagogical experiments with students in a landscape architecture studio at UCD Dublin.

The paper frames radiance fields as an emerging form of spatial artificial intelligence and employs SHARP to substantiate a shift from technically fragmented and epistemically stratified processes toward quasi-intuitive, inference-based engagements with landscapes. Rather than definitive spatial records, radiance fields serve as probabilistic spatial propositions whose validity lies in coherence and situated inference rather than exhaustive measurement and uniform sampling, reframing landscapes as fields of conditions shaped by an internalized set of relationships. Recognizing radiance fields as an operative paradigm for sensing and modeling landscapes clarifies their role as a distinct representational framework for landscape architecture.

Keywords: Radiance fields, Gaussian splatting, SHARP, spatial artificial intelligence, inference-based modeling

1 Radiance Fields: Categorization, Implementations, and Relevance for Landscape Architecture

Radiance fields model the distribution of light within a target scene. Instead of defining a scene through geometric features or surfaces, radiance fields prioritize the propagation of light as the primary target of modeling, shifting emphasis from spatial form toward spatial quantities of light. Radiance fields interpolate between the light conditions captured in image-based observations to estimate how the scene would be perceived from arbitrary viewpoints

---

Journal of Digital Landscape Architecture, 11-2026, pp. 2-15. © Wichmann Verlag, VDE VERLAG GMBH · Berlin · Offenbach. ISBN 978-3-87907-770-0, ISSN 2367-4253, e-ISSN 2511-624X, doi:10.14627/537770001.

in 3D space. This enables novel view synthesis, allowing the generation of views not initially captured by the training data. Radiance field models serve as navigable environments – continuous volumetric fields of light – that respond dynamically to user interactions, adapting conditions such as radiance, reflections, and translucency for each viewpoint.

The unique characteristics of radiance fields, combined with advances in processing and rendering efficiencies, have led to the proliferation of a diverse array of related yet distinct models. Specialized implementations optimize radiance field models for specific metrics, such as processing and rendering efficiency, representational accuracy, visual quality, and extent or complexity of the scene. The terminology and category of radiance fields relate to what is being modeled – the distribution of light within a scene – while accommodating multiple approaches to structuring and representation.

Fig. 1: SHARP-generated radiance field from within the Palm House at the National Botanic Gardens (Site 1), with the foreground fading into darkness and the background illuminated, emphasizing depth, enclosure, and the volumetric light conditions

1.1 Implicit and Explicit Representations

Radiance fields encode light distribution through two distinct and structurally opposed approaches: implicit and explicit representations.

Implicit representations, such as neural radiance fields (NeRF), compute and store light distribution as a continuous function encoded by a neural network. The neural network is queried for every set of position and viewing direction to generate new views – the neural network is the model. Beyond their spatial extent and sampling density, these models conventionally do not encode explicit geometries or formalized spatial data structures. Building on the foundations laid by MILDENHALL et al. (2020), the development of Instant-NGP (MÜLLER et al. 2022) achieved significant performance improvements in training and rendering radiance fields, effectively reducing computational overhead and enabling a broader range of users to engage with radiance fields. The advancements of the approach by Müller et al. (2022) surpassed a pivotal technical threshold, allowing speculation on the uses of neural radiance fields for landscape architecture (SCHOB & REKITTKE 2023).

Explicit representations operate through a finite collection of spatial samples of the function, such as 3D Gaussians (KERBL et al. 2023). These representations employ discrete primitives, which are projected onto the image plane for rendering. Each Gaussian stores a set of attributes, including position, scale, orientation, opacity, and radiance encoded as color (KERBL et al. 2023). While a neural optimization may be used for training, a neural network is generally not required for inference (rendering). While explicit representations are structurally simpler than their implicit counterparts, their reduced hardware requirements and enhanced rendering speeds facilitate a more immediate interaction with the spatial environments they encode. Since its introduction in 2023, the increased accessibility of 3D Gaussian splatting has catalyzed its broad adoption, enabling users to integrate explicit radiance fields into standard design and visualization workflows.

1.2 Implementations and Previous Research

Developed in the field of computer vision for tasks such as digital reconstruction, visualization, simulation, and animation, radiance fields are increasingly transitioning from experimental implementations toward seamless integration into established platforms. This includes standalone applications for explicit representations, such as LichtFeld Studio (LICHTFELD STUDIO 2026), Postshot (JAWSET VISUAL COMPUTING 2025), KIRI Engine (KIRI INNOVATIONS 2026b), Polycam (POLYCAM INC. 2026), or the web-based SuperSplat Editor (PLAYCANVAS LTD., 2026); and implicit representations such as Nerfstudio (TANCIK et al. 2023), and Instant-NGP 2.0 (MÜLLER et al. 2025). ESRI's integration of Gaussian splatting into their ArcGIS product family facilitates georeferencing and geospatial analysis for splat-based models (NORTH et al. 2025), emphasizing their enhanced visual quality, particularly the realistic visualization of vegetation and more accurate representation of materials (SELER 2025). Beyond these integrations, radiance fields for novel view synthesis with NeRF and Gaussian splatting are increasingly included in benchmark datasets, for example, to create digital twins (WYSOCKI et al. 2025). In addition, radiance fields are used for applications in architectural heritage (JAMIL & BRENNAN 2025), historic garden reconstruction (Li et al. 2025), and urban architectural virtual realities (PEZZICA et al. 2024).

Radiance fields provide an expanded contextual framework for the sensing, modeling, and visualization of landscapes in landscape architecture (SCHOB & REKITTKE 2023). Radiance fields shift representation from discrete spatial structures toward continuous fields of conditions. This positions them as a distinct representational paradigm in landscape architecture, complementary to conventional models and technical drawings, as well as photogrammetry- and lidar-based approaches, while emphasizing spatial context, atmospheric conditions, and situated spatial encounters (SCHOB & REKITTKE 2023). The project Fields of Light (SCHOB 2025), presented at the 7th Lisbon Architecture Triennale (2025), extends this conception by foregrounding light as a primary landscape condition, using infrared Gaussian splatting models to articulate how radiative fluxes, atmospheric scattering, and material translucency condition forest landscapes around Oslo in Norway. Currently, we are testing applications for radiance fields in landscape architecture in a Landscape Design Studio at the University College Dublin (UCD), School of Architecture, Planning and Environmental Policy (REKITTKE 2026). Building upon this work, this paper investigates radiance fields as a distinct category of landscape models, emphasizing their capacity to capture continuous, atmospheric, and situated landscape conditions.

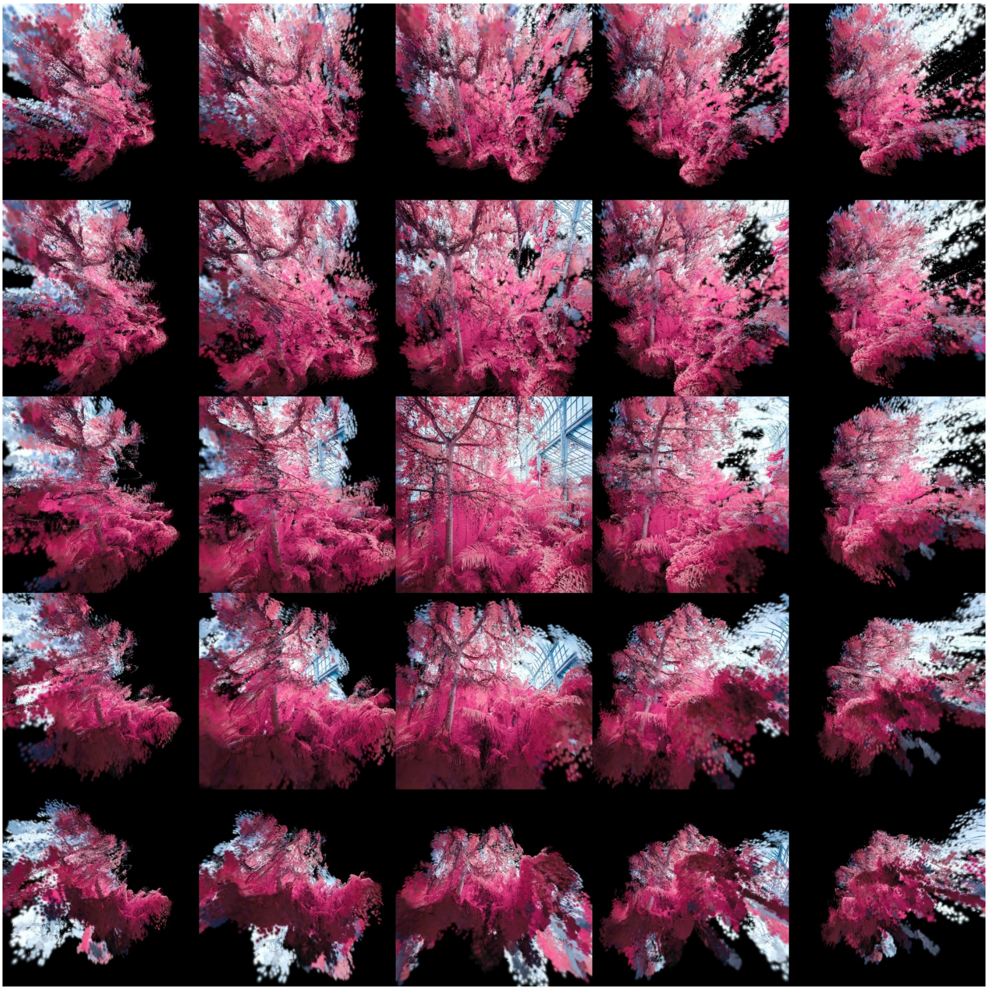

Fig. 2: SHARP-generated radiance field from within the Palm House at the National Botanic Gardens (Site 1), illustrating how the inferred radiance field is rendered from multiple nearby viewpoints relative to the original camera position (center)

1.3 Implications for Landscape Architecture

Beyond their technical characteristics, an engagement with radiance fields holds significant conceptual implications for landscape architecture, including the production of continuous landscape models with visual coherence beyond their underlying data structures, the integration of dynamic target conditions, the capacity to model conditions disregarded by previous approaches, and concrete alignments with remote sensing processes.

Radiance fields encode their target systems as continuous fields rather than discrete entities, which more closely aligns with conceptions of landscapes in landscape architecture (SCHOB & REKITTKE 2023). Rather than being confined to surfaces or discrete objects, radiance fields operate as multi-resolution volumetric fields and are therefore less constrained by the limitations of previous point- or vector-based methods. Their multi-resolution structure enables inference of an expansive landscape context without a proportional increase in data density, concentrating computational resources on perceptual and spatial salience rather than on uniform sampling. For example, reflecting on inefficiencies of lidar data, GIROT (2019) observes that “ninety-nine per cent of the scanned data is seldom used and usually discarded” (GIROT 2019, 113).

The introduction of radiance field models reframes the progress achieved by established sensing technologies. While point-based methods such as lidar scanning may provide superior spatial precision, the capacity to generate dense volumetric models from image-based observation, including historical non-technical photographs, challenges the interdependence of visual and spatial dimensions in sensing and modeling. This includes an asymmetry in data conversion: while point clouds or meshes may be derived from radiance fields, the inverse process is significantly more complex, if not entirely impossible.

Radiance fields appear to change dynamically as the observer moves through the scene, continuously adapting brightness, color, and opacity based on the observer's position and viewing direction. This experiential correspondence more closely aligns with human visual cognition. Additional target conditions, such as a temporal component, can further enhance the model's capacity to accurately capture target conditions or facilitate the simulation of changes within them. 4D Gaussian splatting (WU et al. 2024), for example, encodes three spatial dimensions and time (x, y, z, t) for each Gaussian. Even input data captured at different moments in time can be integrated into a single, multi-temporal radiance field that contains different states of the target system (SCHOB & REKITTKE 2023).

Radiance fields capture target conditions that may be registered as noise by other sensors or remain entirely excluded from conventional data structures (SCHOB & REKITTKE 2023). This includes target properties such as “absorptivity, reflectivity, scattering, challenging texture, and complex shape or geometry” (SCHOB & REKITTKE, 2023, 437), conventionally encountered in water-surface reflections and turbidity, light scattered through clouds or fog, and the diffusion of light by vegetation. In landscape architecture, this expands the scope of what can be sensed and modeled, foregrounding ephemeral conditions that structure landscapes yet are excluded from conventional spatial representations.

Although radiance fields are primarily associated with image-based modeling in computer vision and 3D reconstruction rather than with remote sensing processes, their characterization of the distribution of light within a scene directly relates to conditions such as radiative transfer processes and signal registration in remote sensing. A radiance field model, therefore, can retroactively enhance the characterization of the conditions within the target scene, beyond the limitations of the sparse data structures that serve as model input. In this manner, sensing and modeling operate as recursive and intertwined processes (SCHOB 2026).

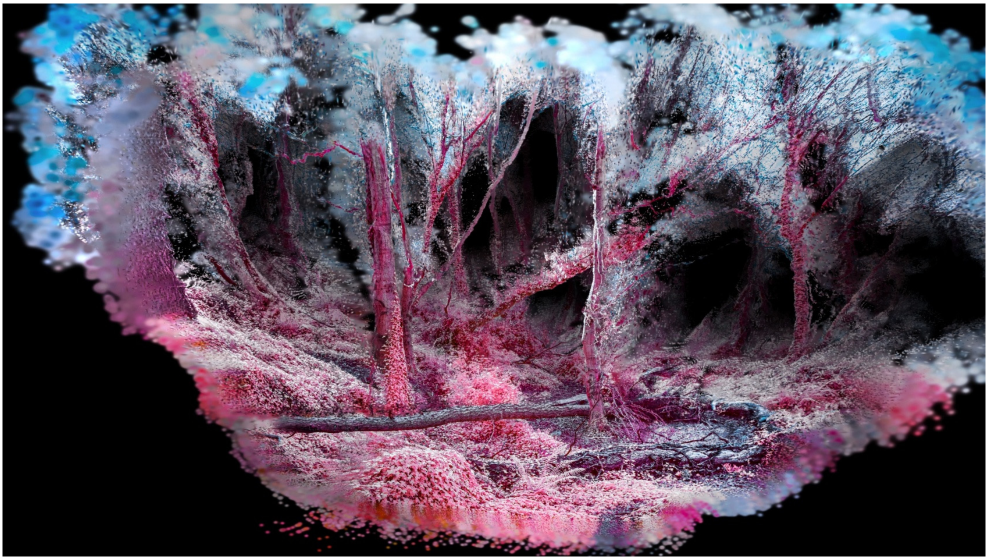

Fig. 3: SHARP-generated radiance field of a forest scene at Killiney Hill Park (Site 2), showing the spatial extent of the inferred reconstruction and the increase in Gaussian splat size toward its outer regions

2 SHARP

SHARP (Single-image High-Accuracy Real-time Parallax) is a novel model for "real-time photorealistic rendering of nearby views from a single photograph" (MESCHEDER et al. 2025, 9). While other models are designed to generate radiance fields from multi-view training data, meaning a scene captured from multiple viewpoints, SHARP focuses on generating a radiance field from a single RGB image (MESCHEDER et al. 2025).

SHARP takes a single image as input and generates 3D Gaussian representations as output. The model "synthesizes a 3D Gaussian representation via a single forward pass through a neural network" (MESCHEDER et al. 2025, 9), which "can then be rendered in real time at high resolution from nearby views" (MESCHEDER et al. 2025, 9). Each Gaussian encodes a defined set of attributes, including position, scale, orientation, color, and opacity (MESCHEDER et al. 2025). View-dependent effects are encoded using a layered depth approach rather than continuous angular functions, such as spherical harmonics. This constitutes a deliberate shift from the complex modeling of directional light conditions toward the modeling of visibility and occlusion, prioritizing perceptual coherence and computational efficiency over more detailed representations. This approach is efficient in representing vegetation, porous structures, and (partially) transparent features.

SHARP was trained on the combination of synthetic and real-world data. The synthetic training data were “generated using an in-house procedural content generation system” (MESCHEDER et al. 2025, 15), which “supports random placement of various object types, including thin structures, transparent materials, and reflective surfaces, across a wide range of spatial configurations” (MESCHEDER et al. 2025, 15). The synthetic data consist of “artist-made environments, comprising over 2K outdoor and 5K indoor scenes” (MESCHEDER et al. 2025, 15), which are “rendered using the V-Ray physically based rendering engine, ensuring photorealistic lighting and material interactions” (MESCHEDER et al. 2025, 16). The real-world training data consists of data from “OpenScene (2023) as well as a collection of high-quality photographs from Shutterstock, Getty Images, and Flickr” (MESCHEDER et al. 2025, 16).

SHARP infers real-world dimensions from image pixel values – outputs are “metric, with absolute scale” (MESCHEDER et al. 2025, 1) – based on a multi-stage training process to establish precise relationships between 2D images and 3D space. While the synthetic data used in the first training stage provides precise dimensions, the real-world data employed in the second stage fine-tunes the model for natural light conditions and material properties. Unlike many other image-based reconstruction models, which merely determine relative depth and proportional spatial relations, SHARP generates representations that employ standard physical units (e. g., meters) and a consistent absolute scale, reflecting real-world dimensions within the scene. While this does not entail alignment with global coordinate reference systems (yet), it implies that the generated environments preserve the physical dimensions of features within their target systems, independent of the resolution or viewpoint of input images. SHARP’s ground truth dimensions are based on synthetic images (GIL-FOURNIER & PARIKKA 2021).

Metric scaling establishes a direct relationship between human and landscape scales, thereby enabling meaningful siting of bodily movement, eye height, and spatial experience. In digital environments, consistent metric scaling enables the synchronization between physical and virtual movements for virtual and augmented reality (VR and AR) applications (MESCHEDER et al. 2025). Within landscape architecture, scale functions not merely as a representational convention but as an operative framework through which spatial relations, bodily engagement, and design intentions are negotiated across sites. In addition, consistent scaling enables the seamless combination of models derived from disparate images into unified scenes.

SHARP is computationally more efficient than comparable models, providing for “state-of-the-art image fidelity for nearby view synthesis” (MESCHEDER et al. 2025, 9) using conventional hardware, such as a laptop or smartphone. While MESCHEDER et al. (2025) explicitly mention the advantages of neural radiance fields (NeRF) and spherical harmonics, the use of a 3D Gaussian representation and a layered depth approach prioritizes training and rendering efficiency over visual fidelity. In this regard, the broad adoption of previous radiance field models, such as Instant-NGP (MÜLLER et al. 2022) and 3D Gaussian splatting (KERBL et al. 2023), was equally conditioned by their increased performance, emphasizing that efficiency improvements are not merely technical or computational optimizations, but a prerequisite for a transition into mainstream applications and use cases.

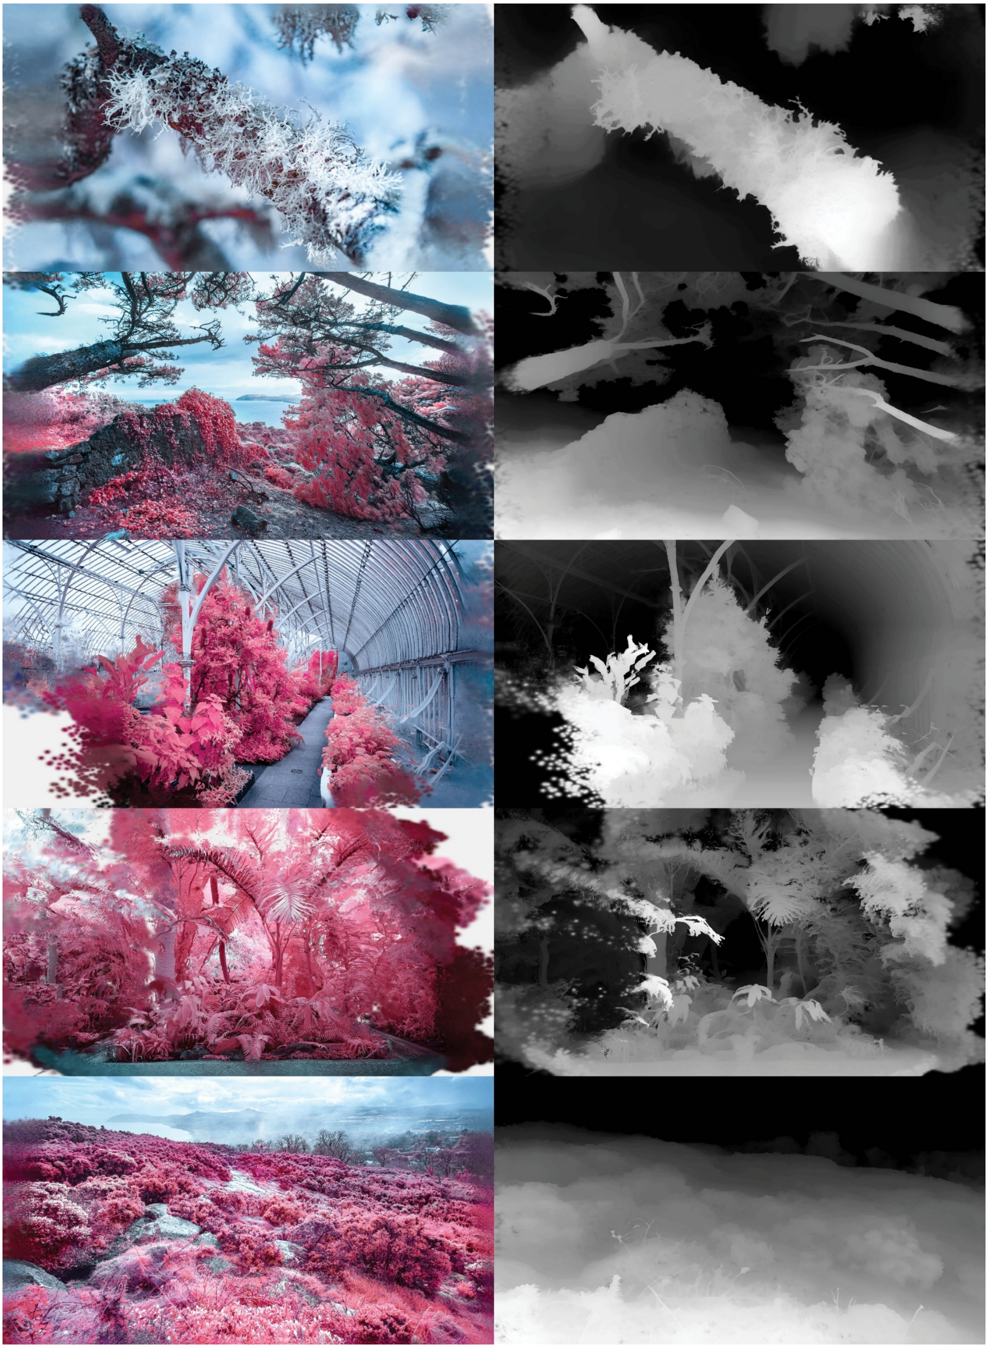

Fig. 4: SHARP-generated radiance fields (left) and corresponding greyscale depth maps (right) from the same viewpoints across Site 1 (rows 3, 4) and Site 2 (rows 1, 2, 5), illustrating the relation between perceptual coherence and inferred depth structure

3 Model Explorations with SHARP

We employ SHARP to generate volumetric 3D Gaussian radiance fields from single image inputs. To test and demonstrate the versatility of the model, rather than focusing on a single case or site, we employ a selection of images capturing varied target conditions, including differences in spatial extent, diverse landscape typologies, and varying light conditions. While other image-based reconstruction methods rely on structured capture protocols and constrained acquisition geometries, such as linear or orbital movements centered on a target, SHARP achieves high-quality reconstructions largely independent of the capture context, demonstrating "robust zero-shot generalization across datasets" (MESCHEDER et al. 2025).

The data for this case consist of infrared imagery captured with a converted full-spectrum camera in combination with a selection of infrared filters. Vegetation exhibits strong reflectance in the near-infrared spectrum, a characteristic effectively captured in the infrared imagery and translated into each radiance field model. The use of infrared imagery foregrounds SHARP's operational capabilities beyond true color imagery, extending into the domains of remote sensing and technical imaging. The images were captured at two locations in Dublin, Ireland: the National Botanic Gardens (Site 1; 09.02.2026) and Killiney Hill Park (Site 2; 11.02.2026). Preprocessing of raw image files to JPEG format was performed using the open-source photography application darktable. The data encompass multiple spatial scales, combining interior and exterior environments from close-range detail views to expansive landscapes.

Based on the SHARP GitHub repository (APPLE 2025), we developed a local Jupyter Notebook implementation (JUPYTER DEVELOPMENT TEAM 2026) to structure installation and data processing tasks. The Notebook facilitates systematic batch processing: each image in the input directory is processed individually and exported to the output folder as a radiance field representation in .ply format. On a conventional laptop (Intel Core Ultra 7, 32GB RAM, RTX 4060 Laptop GPU), inference required only a few seconds per image, enabling efficient processing across image collections. For the visualization and further engagement with the resulting models, we used SuperSplat Editor (PLAYCANVAS LTD. 2026) and the KIRI 3DGS add-on (KIRI INNOVATIONS 2026a) for Blender (BLENDER FOUNDATION 2025).

The resulting radiance fields unfold as complex 3D scenes with structural coherence beyond the captured views of the input data. We observe that the visual and spatial coherence appears to be relatively higher in close proximity to the original camera position, with increasing uncertainty toward more distant regions of the scene. The consistent alignment of the reconstructions to absolute metric scale is one of the more impressive features of SHARP. In our implementation, the reconstructions maintain the dimensions of their target systems at the meter scale, enabling seamless combination of models derived from different images. Specific to monocular view synthesis, SHARP excels at capturing the intricate organic complexities of vegetation, while occasionally exhibiting limitations in resolving thin linear structures, such as pipes, columns, or branches. Rather than generating isolated geometric objects, SHARP produces concentrations of spatial information, shifting how landscape entities are identified and engaged.

The utility of SHARP was additionally tested in a landscape architecture studio at University College Dublin (UCD), where students explored SHARP as part of their design projects. Within the studio, we explored a wide range of inputs, including smartphone images, landscape paintings, historic photographs, and game engine captures, across which SHARP reliably generated coherent 3D environments. The students used the reconstructions to navigate the virtual scenes, identify and measure features, describe their spatial relations, and produce stylized renderings. The combination of reconstructions of the same scene from different time periods enabled an exploration of temporal changes. In Ireland, where access to public geospatial data is limited, SHARP proved particularly useful by providing custom, case-specific data, enabling students to engage with situated conditions that would otherwise be difficult to model or assess. Within this context, SHARP provides an alternative framework for sensing and modeling sites that is easily accessible to students and highly applicable to landscape architecture.

Fig. 5: SHARP-generated radiance field within a greenhouse at the National Botanic Gardens (Site 1). The lower-left region reveals the underlying 3D Gaussian splatting structure of SHARP within the foliage.

4 Discussion: Situating SHARP and Spatial Artificial Intelligence

SHARP engages landscapes as situated visual-spatial encounters rather than spatially exhaustive environments. The inferred digital landscapes function as probabilistic spatial propositions conditioned by learned visual priors that articulate plausible spatial organizations with maximal inference from minimal input. These reconstructions combine the visual fidelity of radiance fields with the spatial precision and metric scaling inherent to lidar sensing, maintaining local coherence while becoming increasingly indeterminate beyond the immediate context of inference.

SHARP internalizes visual and spatial relationships present in extensive training data, embedding them within the parameters of the model rather than deriving them explicitly from the input image. Similar to human visual perception of spatial dimensions, SHARP interprets visual cues, such as texture gradients, material granularity, relative size relationships, atmospheric attenuation, and spatial frequency patterns, to transmute gridded brightness values into a complex 3D scene. While the training stages and internal reasoning of SHARP remain inaccessible to its users, SHARP effectively collapses a multi-stage process of sensing, processing, verification, scaling, and visualization into a one-click application – inference replaces measurement. SHARP provides access to advanced digital models for users not conventionally involved in the technical aspects of sensing and modeling in remote sensing and computer vision, but rather in the aesthetics of landscape photography.

SHARP is one instance within a broader proliferation of sensors and models, each generating situated and partial captures of their target systems and each requiring evaluation with respect to the particular aspects, assumptions, and operational conditions through which they render landscapes legible. For radiance fields and Gaussian splatting representations, evaluation metrics may differ substantially from those used in preceding gridded or point-based methods and data structures. In alignment with previous technological categorizations, such as lidar remote sensing, photogrammetric processing, and point cloud representations, the emergence of radiance field models must be recognized as a distinct and proliferating category for modeling and representing landscapes. This consolidates emerging conceptions of radiance fields into a coherent discussion within landscape architecture.

SHARP generates 3D Gaussian splatting representations, an implementation of radiance field models, which themselves constitute part of an expanding domain of spatial artificial intelligence (PAPADIMITRIOU 2025). Their emergence signals a shift in how 3D landscape models are produced, accessed, and mobilized within design disciplines. The sensing and modeling of landscapes become near-instantaneous and quasi-intuitive processes, less constrained by heterogeneous sensor configurations, fragmented data collections, inaccessible software packages, and specialized expertise outside of the field of landscape architecture. This implies a structural reorientation in the role of digital models within landscape architecture: from manual acquisition, verification, and modeling toward embedded, accessible, visually-grounded, spatial engagements.

This evolution points to the emergence of world models, in which formerly distinct domains, such as computer vision, simulation, and geospatial analysis, are increasingly consolidated within a unified computational framework. Drawing on the logic of the bitter lesson (SUTTON 2019), contemporary artificial intelligence research suggests that the intermediate representations and task-specific heuristics that once structured fields like computer vision and digital reconstruction are becoming epistemically peripheral (SITZMANN 2026). Conceptions of 3D modeling are shifting from the production of explicit geometric representations toward the internalization of spatial structures within scalable world models, reframing geometric operations, classification, and rendering as transient scaffolds rather than epistemic endpoints. Consequently, the role of landscape architects shifts from managing fragmented technical procedures toward configuring and directing models that internalize the multi-modal organization of physical environments as coherent, operative systems. SHARP, therefore, represents more than a technical addition to an expanding collection of models driven by artificial intelligence, marking an early manifestation of a broader convergence within sensing and modeling of the Earth.

Fig. 6: SHARP-generated radiance field of an expansive landscape view at Killiney Hill Park (Site 2), capturing the extended terrain and distant horizon, to demonstrate large-scale spatial coherence and atmospheric depth

5 Conclusions

This paper has demonstrated how radiance fields, instantiated through SHARP, operate as inference-based landscape models capable of generating metric-scale, spatially coherent environments from minimal input. By situating SHARP within a broader categorization of radiance fields and validating its application through model explorations and pedagogic experiments, the paper clarifies the role of spatial artificial intelligence as an emerging operative paradigm in landscape architecture. Rather than replacing established sensing and modeling processes, these models shift the emphasis from exhaustive measurement and uniform sampling toward situated, probabilistic, and experiential forms of spatial reasoning.

The emergence of radiance fields represents a significant advance in landscape modeling, offering a more immersive experience than established digital approaches. Unlike point clouds, triangles, vectors, and related geometric reconstruction methods, radiance fields capture complex atmospheric conditions essential to visual perception, including volumetric brightness, color, and transparency, specular reflections, and dynamic light scattering. This suggests latent potential within the technology, a phenomenological surplus, by which the model captures the ephemeral qualities of experience and encounter relevant to landscape architecture. As such, radiance fields serve as modeled propositions that mediate sensing, inference, and spatial experience. Recognizing this capacity has value for landscape architects, as it elevates the digital model from a static representational surrogate to a means of more closely articulating the complex, shifting conditions within its target system.

This development marks a shift from procedural, competency-dependent toward quasi-intuitive, near-instantaneous modes of landscape modeling. While access to advanced tools and methods remains a constraint in landscape modeling, novel approaches synthesize previously disparate processes and their associated requirements for specialized knowledge, effectively lowering the threshold for engaging landscapes through complex forms of representation. While this transition is neither complete nor immediate, its potential is already evident in practice and can now be articulated in a methodologically grounded, scientifically substantiated manner.

References

APPLE (2025), SHARP: Sharp monocular view synthesis in less than a second [Python]. Apple. https://github.com/apple/ml-sharp.

BLENDER FOUNDATION (2025), Blender: The Free and Open Source 3D Creation Software (Version 5.0.1) [Computer software]. https://www.blender.org/.

GIL-FOURNIER, A. & PARIKKA, J. (2021), Ground truth to fake geographies: Machine vision and learning in visual practices. AI & Society, 36 (4), 1253-1262. https://doi.org/10.1007/s00146-020-01062-3.

GIROT, C. (2019), “Cloudism”: Towards a new culture of making landscapes. In: BRAAE, E. & STEINER, H. (Eds.), Routledge Research Companion to Landscape Architecture, Routledge.

JAWSET VISUAL COMPUTING (2025), Jawset Postshot: Radiance field–based 3D reconstruction software (Version 1.0.110) [Computer software]. https://www.jawset.com/.

JUPYTER DEVELOPMENT TEAM (2026), Jupyter Notebook (Version 7.5.2) [Computer software]. https://github.com/jupyter/notebook.

KERBL, B., KOPANAS, G., LEIMKUEHLER, T. & DRETTAKIS, G. (2023), 3D Gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics, 42 (4), 1-14. https://doi.org/10.1145/3592433.

KIRI INNOVATIONS (2026a), KIRI 3DGS Blender addon by KIRI Engine (Version 4.1.5) [Python]. https://github.com/Kiri-Innovation/3dgs-render-blender-addon

KIRI INNOVATIONS (2026b), KIRI Engine [Computer software]. https://github.com/Kiri-Innovation.

LI, J., LIU, M. & JIAO, Y. (2025), Research Progress on Digital and Immersive Methods for Historic Garden Reconstruction. Proceedings of the 2025 International Conference on Artificial Intelligence, Virtual Reality and Interaction Design, 775-781. https://doi.org/10.1145/3777730.3777855.

LICHTFELD STUDIO (2026), LichtFeld Studio: A high-performance C++ and CUDA implementation of 3D Gaussian splatting (Version 0.4.0) [Computer software]. https://github.com/MrNeRF/LichtFeld-Studio.

MESCHEDER, L., DONG, W., LI, S., BAI, X., SANTOS, M., HU, P., LECOUAT, B., ZHEN, M., DELAUNOY, A., FANG, T., TSIN, Y., RICHTER, S. R. & KOLTUN, V. (2025), Sharp monocular view synthesis in less than a second (arXiv:2512.10685). arXiv. https://doi.org/10.48550/arXiv.2512.10685.

MILDENHALL, B., SRINIVASAN, P. P., TANCIK, M., BARRON, J. T., RAMAMOORTHI, R. & NG, R. (2020), NeRF: Representing scenes as neural radiance fields for view synthesis (arXiv:2003.08934). arXiv. http://arxiv.org/abs/2003.08934.

MÜLLER, T., EVANS, A., SCHIED, C. & KELLER, A. (2022), Instant neural graphics primitives with a multiresolution hash encoding. ACM Transactions on Graphics, 41 (4), 1-15. https://doi.org/10.1145/3528223.3530127.

MÜLLER, T., EVANS, A., SCHIED, C. & KELLER, A. (2025), Instant-NGP (Version 2.0) [Cuda]. NVIDIA Research Projects. https://github.com/NVlabs/instant-ngp.

NORTH, C., ROLAND, C. & KHAN, H. (2025), Your ArcGIS pro update (November 2025). ArcGIS Blog. https://www.esri.com/arcgis-blog/products/arcgis-pro/announcements/your-arcgis-pro-update-november-2025 (25.02.2026).

PAPADIMITRIOU, F. (2025), The Rise of Spatial AI. In F. Papadimitriou, Spatial Artificial Intelligence. Springer Nature Switzerland, 1-19. https://doi.org/10.1007/978-3-031-82136-3_1.

PEZZICA, C., CHIONI, C. & MOLS, N. M. L. (2024), Exploring alternative urban and architectural virtual realities through multidomain digital twins. In: GIORDANO, A., RUSSO, M. & SPALLONE, R. (Eds.), Advances in Representation, Springer Nature Switzerland, 675-690. https://doi.org/10.1007/978-3-031-62963-1_41.

PLAYCANVAS LTD (2026), SuperSplat Editor (Version 2.18.0) [TypeScript]. https://github.com/playcanvas/supersplat.

POLYCAM INC. (2026), Polycam [Computer software]. https://poly.cam.

REKITTKE, J. (2026), Landscape studio [Design Studio]. University College Dublin (UCD), School of Architecture, Planning and Environmental Policy.

SCHOB, M. (2025), Fields of light [Exhibition]. Museum of Design and Fashion (MUDE), Lisbon, Portugal. https://www.2025.trienaldelisboa.com/en/coalition/how-to-engage-a-world-composed-of-light (25.02.2026).

SCHOB, M. (2026), The model view: Models for remote sensing in landscape architecture [Unpublished doctoral dissertation]. Norwegian University of Life Sciences (NMBU).

SCHOB, M. & REKITTKE, J. (2023), Neural radiance fields for landscape architecture. Journal of Digital Landscape Architecture, 8-2023, 428-442. https://doi.org/10.14627/537740046.

SELER, I. (2025), 3 Key Benefits of Gaussian Splat Layers. ArcGIS Blog. https://www.esri.com/arcgis-blog/products/arcgis-pro/3d-gis/3-key-benefits-of-gaussian-splat-layers (25.02.2026).

SITZMANN, V. (2026), The Flavor of the Bitter Lesson for Computer Vision. https://www.vincentsitzmann.com/blog/bitter_lesson_of_cv/ (25.02.2026).

SUTTON, R. (2019), The Bitter Lesson. Incomplete Ideas. http://www.incompleteideas.net/IncIdeas/BitterLesson.html (25.02.2026).

TANCIK, M., WEBER, E., NG, E., LI, R., YI, B., KERR, J., WANG, T., KRISTOFFERSEN, A., AUSTIN, J., SALAHI, K., AHUJA, A., MCALLISTER, D. & KANAZAWA, A. (2023), Nerfstudio: A Modular Framework for Neural Radiance Field Development. SIGGRAPH '23 Conference Proceedings, 1-12. https://doi.org/10.1145/3588432.3591516.

WU, G., YI, T., FANG, J., XIE, L., ZHANG, X., WEI, W., LIU, W., TIAN, Q. & WANG, X. (2024), 4D Gaussian splatting for real-time dynamic scene rendering (arXiv:2310.08528). arXiv. https://doi.org/10.48550/arXiv.2310.08528.

WYSOCKI, O., SCHWAB, B., BISWANATH, M. K., GREZA, M., ZHANG, Q., ZHU, J., FROECH, T., HEERAMAGLORE, M., HIJAZI, I., KANNA, K., PECHINGER, M., CHEN, Z., SUN, Y., SEGURA, A. R., XU, Z., ABDELGAFAR, O., MEHRANFAR, M., YESHWANTH, C., LIU, Y.-C., … JUTZI, B. (2025), TUM2TWIN: Introducing the large-scale multimodal urban digital twin benchmark dataset (arXiv:2505.07396). arXiv. https://doi.org/10.48550/arXiv.2505.07396.