Towards a Photographic Representation of the Experience of Seeing: Synthetic views via Neural Radiance Fields

Submitted for the Expanded Visualities special issue.

Andrew Burrell

School of Design, Faculty of Design Architecture and Building
University of Technology Sydney

andrew.burrell@uts.edu.au

15 Broadway,
Ultimo NSW 2007
AUSTRALIA

Abstract:
This article looks at the emerging technology of the Neural Radiance Field (NeRF) and suggests that this means of digital image production, when used by a creative practitioner, produces an emergent aesthetic—a result of the affordances inherent in the digital materiality of the NeRF and its processes. The use of this machine-learning technology as a material for creating images, can lead to an entirely new way of creating photographic representations that provide us with a way of seeing and recording how we phenomenologically experience seeing. The processes and the emergent aesthetic are explored with examples drawn from the author's own practice.

Contributor notes:
Andrew Burrell is a practice-based researcher investigating the role of memory and imagination in networked virtual environments. Andrew's research explores the networked relationships between ecologies, technologies, and people and how those interactions can inform the ethical and sustainable design of more-than-human virtual spaces. Andrew is a senior lecturer in Visual Communication at the University of technology Sydney.

Keywords:
Neural radiance fields, Expanded photography, Machine learning, Phenomenology of seeing, Design Research, Creative practice

Towards a Photographic Representation of the Experience of Seeing: Synthetic views via Neural Radiance Fields

Submitted for the Expanded Visualities special issue.

Introduction

This article looks at the emerging technology of the Neural Radiance Field (NeRF). The NeRF sits as part of a larger suite of technologies that are constantly expanding visual and photographic practices but is unique for several reasons that will be discussed. NeRFs have not seen the hype recently associated with other machine-learning processes, such as Large Language Models and text-to-image Diffusion Models. This might be because their results don't seem as "magical"[1] at first. Instead of creating unique "generative" outputs (whose relationship to physical world is mediated via a massive data set of images), they create synthetic photographic representations that have a direct visual relationship with their referent in the physical world, as a view of a 3D environment constructed as a verisimilar representation of a single physical environment. In their article "Photography after AI", Palmer and Sluis (Palmer et al. 2023:20) point us to Steyerl's proposition that AI Diffusion models "shift the focus from photographic indexicality to stochastic discrimination" (Steyerl 2023: 82). What we will find with photographic images created via the NeRF process is a superposition of both an indexical and stochastically discriminated image. It will be argued that these images do something entirely new in presenting a creative practitioner with the

[1] Magical in the sense that the processes involved in generating seemingly complex and "human-like" outputs are not immediately understandable or knowable - in the sense of Arthur C. Clarke's famous line that "any sufficiently advanced technology is indistinguishable from magic."

potential to create photographic images that provide a way of representing how we experience seeing.

This article draws on practice-based research, published technical papers[2], new-media theory, and a phenomenology of seeing to form its position. It seeks to better understand this technology from the perspective of practice, and it understands the inherent nature of creative practice to be one of an ongoing dialogue with continually emerging technologies.

Neural Radiance Fields

Neural Radiance Fields first came to prominence in 2020 with the paper "NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis" (Mildenhall et al. 2022), which described the technical processes for creating a synthetic scene from a series of photographs. This scene is contained within the weights of a neural network that is trained on a series of photographs or images of an "original" scene. The resulting neural network then takes as an input a location and a viewing direction. It then output points with density and a colour[3] to reconstruct up the scene. This results in any given point in space having a potentially different colour based on the angle from which it is viewed (described as being "view-dependent" in the technical literature), and this colour, when combined with the density of that point, produces a "radiance". Ultimately, these are usually viewed through the "window" of the screen, resulting in what we can imagine to be a 5-dimensional environment

[2] This current article should not be read as a technical paper, but as an exploration of the material nature of the NeRF as a means of producing novel, camera-less photographic images informed by practice-based research.

[3] Interestingly Harwood, in their poetic reflection on the pixel reminds us that this is how the human eye operates with its "rods" perceiving only brightness (or density) and "cones" perceiving only colour (Harwood [2008] 2008: 216).

being flattened down to a 2-dimensional photographic representation of a single viewing position. This is described as follows:

> This novel method takes multiple images as input and produces a compact representation of the 3D scene in the form of a deep, fully connected neural network, the weights of which can be stored in a file not much bigger than a typical compressed image. This representation can then be used to render arbitrary views of the scene with surprising accuracy and detail. (Dellaert 2022: 98)

The NeRF process, then, creates a view-dependent volumetric scene within the weights of the neural network defining a continuous volume in 3D space.

Creating a NeRF begins along the same lines as traditional photogrammetry. From a series of photographic images of a scene or object, camera positions and directions in 3D space are calculated (as estimations). There are several methods and associated tools for doing this, including commercial software and open-source tools. Many popular NeRF pipelines include a process for calculating or estimating camera positions and orientations in 3D space.[4] Figure 1 demonstrates this process of defining camera positions and directions in 3D space. It shows a view of "Agisoft Metashape" software with a data set of six images and a generated point cloud. Metashape provides a workflow for creating traditional photogrammetry but can also be used for generating camera orientations as input data for training a NeRF. This image also contains directional rays that have been added by the author for clarity.

[4] NVIDIA's Instant-NGP or Instant NeRFs (Müller et al. 2022: 102:1) is one such associated tool which provides relatively a straightforward pipeline to creating synthetic views using NeRF processes. More recently NeRF Studio (Tancik et al. 2023) provides a modular approach to using a suite of technologies and NeRF algorithms to generate novel scenes within a dedicated viewer. NeRF studio also provides a series of robust output options for converting NeRF scenes and views to other formats.

{FIGURE 1: Camera positions and view directions defining a 3D point cloud in Agisoft Metashape, 2023, Digital Image. ©Andrew Burrell}

The processes of creating a NeRF are a little more transparent than some other machine-learning processes, in that it is possible to observe some of the processes take place in real-time. This is not because the "black box" of the neural network, or the weights contained within, are visible or knowable, but because the translation of these weights, into synthetic volumetric scenes can be viewed while the training process takes place, in real-time. A scene emerges from a field of noise in a surprisingly short amount of time—a matter of seconds in most cases. This is made possible through the inbuilt viewers in these pipelines, which provide instant access to a visualization of the output. The upshot of this is that through this direct visual feedback, a creative practitioner is able to make more informed and direct actions in guiding an image outcome towards their intent, rather than the purely stochastic outcome of the machine—It allows them to see and understand how their inputs are affecting the outcome, enabling them to steer the results in a direction that aligns with their creative vision. This process can be seen demonstrated in Figure 2, which shows a novel view of a scene emerging from pixelated noise over a period of a few seconds (first three images) and a resulting view after approximately 30 seconds of training (final image in the sequence).

{FIGURE 2: Protea, 2023. Image Sequence of a NeRF scene emerging from noise.
© Andrew Burrell}

Hallucinations & Emergent Aesthetics

Using the NeRF process to produce images allows us to visually access the "window" of the built-in viewer, in which the NeRF system generates synthetic views, resulting in a continuous 3D environment. This scene emerges from what seems like random noise, resulting in a continuously changing image, showcasing what might be termed an emergent NeRF aesthetic. This aesthetic is, in part, a result of the visual representation of the extraneous noise that does not make up part of the "real world" scene but has emerged as part of the NeRF's attempt to form a particular novel view of the scene. The system attempts to visualize the data encoded in the neural network from the requested viewing direction, giving each position a colour and radiance based on its model of the scene. Where training data was perhaps sparser or misinterpreted, it continues in its task in attempting to build an image with visual continuity via predictions. Visually, this appears as mists or clouds of points with radiance and colour, depicted as pixels in space. These take on a volumetric form and encase the rendered scene in a shroud. What may have been empty space, or another part of the scene in the "ground truth" is now collapsed into abstract form through a process often described in machine learning as "hallucinations".

This notion of ground truth is important in this discussion. It is a term that has a specific meaning in the machine learning space to connote something known to be a matter of fact in the original training data set. This ground truth provides the target the model aims to learn, comparing its predictions against. This is obviously fraught, given the impossibility that a piece of data can be said to be inherently truthful[5]. Still, it is useful in distinguishing between

[5] This is where many of the criticisms around the inherent bias in training data has been aimed—ground truth is often obtained via the labelling of data by a person either directly or indirectly. This is an important discussion that Kate Crawford (Crawford 2021) provides a comprehensive introduction to.

what a user/operator has deemed to be truthful, as opposed to the generated output of the system. In the case of the NeRF process, the ground truth is said to be the original photographs and their known orientations used to train the model. Yet, it might be more accurate to extend the idea of ground truth beyond these photographs, to include the actual physical scenes or objects they depict. This introduces a complex layering of references: the "ground-truth" then shifts from these photographs to the physical subject the NeRF captures. The concept of ground-truth becoming flexible and multi-dimensional, revealing a complexity in how we might interpret NeRF images.

Between these layers of ground truth and the indexical photographic representation lie the "hallucinations" of the emergent NeRF aesthetic. The use of this term "hallucination" has become popular in machine-learning circles; poetic as it may be, it is problematic in that it contributes to the personification of "AI" as having a level of intelligence or even consciousness that is not present within these systems. For example, this use of language can be found in the instructions for setting up and running Instant NeRF. The author writes: "…for natural scenes where there is a background that extends beyond this bounding box, the NeRF model will struggle and may hallucinate "floaters" at the boundaries of the box." (Instant Neural Graphics Primitives 2022) Clearly, if fewer images or lower-quality images are used to train the NeRF model, the chances of visual anomalies increase. This is because the model has less data to work with to predict what value each point should have. While to a certain extent, the amount and nature of the hallucinations present in a generated NeRF image depends on the "quality" of the training data, regardless of the training data, if you "pull back" far enough from the edges of the data, the machine's access to a finite ground truth becomes evident, and the "floaters" described in the quote above will become increasingly apparent.

{FIGURE 3: Nori, 2022, NeRF generated photographic portrait. © Andrew Burrell}

We can see the results of this process in Figure 3. “Nori” is an image in a series of NeRF portraits created by the author. In this NeRF-generated photographic portrait, a figure appears to emerge from or is encased by the “hallucinated” mist. We see fragments of a larger scene, indicating the environment the figure is sitting within, but within this single image, it feels like this environment is not fixed or solid, but dependent on our ongoing visual attention. The NeRF used to create this image was trained on 42 photographs of the subject taken from multiple angles. The view has been “pulled back” with the intent to manipulate the NeRF aesthetic to emphasise its ability to replicate a phenomenological experience of seeing. This will be discussed in detail in a following section which examines another portrait in this series in more detail.

Another form of “hallucination” common in novel views generated by NeRFs is that of the repeated object or element of the scene at different scales and positions throughout the scene. This comes about, presumably, when the training data provided to the neural network was insufficient for it to correctly form a representation of the scene with the scale and position of a single object incorrectly represented as multiple objects. What might be poetically referred to as the neural network becoming “confused” about scale and position with an understanding

that this term is as problematic as that of "hallucination"[6]. This means that when novel views of a scene are created, what might be described as ghosts, or "specters of the ground truth" appear in the image. These elements of machine hallucination can combine to produce images which can feel a long way from the "original" photographs they are trained on, and even further from the "physical world, ground truth" scene the objects these photographs were produced from. Figures 4 and 5 show this ghosting effect as it occurs in the generation of a NeRF portrait, producing repeating motifs as "specters of the ground truth." The image is one of many attempts towards understanding the NeRF aesthetic in producing the portrait series, including "Nori" (Figure 3) and "Tim" (Figure 6). The NeRF model used to create these images used an intentionally small number of training images (30) and had a relatively consistent background with which the training process had trouble, due to the repeated textures. In crafting photographic images using NeRFs, we see that it is possible to understand potential desired results in outcomes based on the images we use to train the systems and manipulate these towards desired visual outcomes.

{Figure 4 and 5: Specters of the ground truth, 2022, experimental NeRF generated images. © Andrew Burrell}

Towards a photographic representation of the experience of seeing

Much of the interest in NeRF technology to create images has been in creating virtual scenes of "real world" environments and how this may be a replacement for more expensive or cumbersome technologies such as drone cinematography and flythroughs. Yet one thing this

[6] For the developing argument against the current language being used to describe "Artificial Intelligence" and "Machine Learning" processes, see (Hunger 2023). This current article endeavors to point out where language used may be deceiving using "scare quotes".

misses, and our discussion here leads to, is that synthetic images created with NeRF technologies provide a possible visual representation of a phenomenological experience of seeing.

These novel photographic images have had quite a journey of modulation and re-modulation between data forms in their production[7]. They begin in the analogue physical environment, which is digitized as digital photographs—this physical environment acting as a ground truth for these images. These images then act as a ground truth in their own right for the modulation of the data into the internal model of the neural network through the training process. This neural network contains a model of the scene as a continuous, implicit representation – where the values are stored algorithmically and inferred rather than stored as values for any single point in space or on the screen. Given a viewing angle, this process of inference remodulates the continuous representation back to a pixel-based digital representation of the scene—presented in the form of pixels in a particular order on the screen's grid, to produce a synthetic photographic image.

This image is only one of a multitude of possible images. As an image is generated by the NeRF process, it is "manipulated into being" by the person operating the system. As a creative practitioner working with the NeRF as a material space, this is akin to a traditional photographer's control over a myriad of parameters such as framing, composition, timing, etcetera. It is a skilled and informed curation of an image (or series of images) from all the

[7] The notion of the understanding of the world around us being the result of a series of modulations and remodulations between data forms originates in "Seven thesis of the concept of post-convergence (Clemens et al., 2010). It speaks in general about the mediation of data via modulation and remodulation into a display register—in our case it can be read specifically against the photographic as a means modulation and display of visual data in image form as a result of any number of technical processes.

possibilities of an image contained within the implicit representation defining the space of possible novel views of the NeRF scene. As we have seen, a certain aesthetic can emerge in this space, which practitioners can tend towards in their curation of emergent images. What follows, then, is a technology that allows a creative practitioner to explore representing not only what they see, but their experience of seeing it.

In the introductory paragraphs to “Ways of Seeing” John Berger writes:

> It is seeing which establishes our place in the surrounding world; we explain that world with words, but words can never undo the fact that we are surrounded by it. The relation between what we see and what we know is never settled. Each evening we see the sun set. We know that the earth is turning away from it. Yet the knowledge, the explanation, never quite fits the sight. (Berger 1997)

The fragmentary nature of this emergent NeRF aesthetic presents us with an image that establishes—in the manner described by Berger—the viewer within the viewed and places them within the visually represented surroundings. The “hallucinated” peripheries of a novel view, combined with the prospect of visual “ghosts”, create a photographic image that holds within it a fixed moment that contains the potential for a shifted glance to reveal more or a refocusing of visual attention that shifts the undefined peripheries and boundaries of this now two-dimension representation of the scene.

Anyone who has ever taken a photo of a scene with a consumer-level snapshot camera or phone camera will understand the disappointment associated with the disconnect between the scene as represented in the resulting image and the phenomenological experience of having “seen” that scene. Aligned with Berger’s explanation of never entirely fitting the scene, the

modulation, in this case, has reduced an experience to the display register of the photographic image on screen or paper.

A creatively constructed NeRF image then not only captures a synthetic view based on a real-world scene, but also has the potential to capture this experience of seeing, an experience often lost in that disconnect between an image and its referent. This act of seeing, as perceived by the observer, is a bodily act—beyond the eye—and likewise beyond the lens in typical lens-based processes. It is described by Michael Inwood in his account of a phenomenological experience to seeing:

> Even though vision does not require noticeable change, we can of course see things change as well. But even if the scene I survey does not noticeably change, I can change the focus of my attention. My attention wanders from the street below with its apartment blocks to the mountain in the distance. The apartment blocks are still within my field of vision, only somewhere at its periphery. I can turn my head to change my perspective. I can close my eyes if the sun is too bright. (Inwood 2015: 169)

The phenomenological act of seeing is much more than the mechanics of vision. It is an act of the body and an act that is inextricably tied to both our memory and imagination. The act of seeing is not static—neither in time nor space—and as pointed out by Merleau-Ponty (in this case, talking to perception more widely) it is an interconnected relationship between the thing being seen and the thing seen, stating that "[t]he thing can never be separated from someone who perceives it" (Merleau-Ponty et al. 2012: 334).

The embodied nature of seeing is elaborated on by Sokolowski in his “Introduction to Phenomenology” and we see, that in the same way it cannot be unentwined from that which is seen, it cannot be unentwined from other senses:

> The philosophical analysis of vision, for example, will show how vision is founded upon the eye and also upon bodily mobility (on the saccadic movements of the eye, on the ability of the head to be turned, on the ability of the whole body to go from one place to another, from one viewpoint to another), how both seeing and that which is seen are moments within a whole, and how seeing is conditioned by other sensory modalities, such as touch, hearing, and kinesthesia. (Sokolowski 2000)

A creatively produced NeRF, where the practitioner takes agency in the process, guiding it towards a desired outcome, can offer an evocative visual representation of the act of seeing itself. This perspective shifts the focus from the static depiction of the external world to the act of seeing as an active, bodily experience. The 'hallucinations' present in the NeRF images, perhaps seen as imperfections by a computer scientist, are, to the creative practitioner elements of the materiality of the NeRF process that they have an amount of control over and stand as visual metaphors for the ways our perception fills gaps and constructs wholes from fragments.

{FIGURE 6: Tim, 2022, NeRF generated photographic portrait. © Andrew Burrell}

Let’s consider this in relation to the NeRF portrait “Tim” (figure 6). As with the portrait “Nori” (figure 3), this viewing of this image is reliant on our every shifting visual attention. This NeRF portrait is part of a series by the author that takes advantage of the NeRF aesthetic

to create images that capture, at least in part, the phenomenological act of seeing. This NeRF was trained on 46 images taken from positions that represent a full coverage of the subject from all angles. These images represent the indexical element of the image—as the ground truth, they remain in a direct causal relationship with what is signified – the sitter, Tim, of the title. Like all static images created from the NeRF process, it is only one of a massive number of possible images—the act of moving through the NeRf environment (navigating the continuous 3D space created via the process after training on the images has taken place) and working with cropping parameters to include or exclude parts of this 3D environment—enables creative choices to be made. Here, we find the elements of the stochastic discrimination, suggested by Steyerl, taking place, both by the algorithms involved and as importantly, the choices made by the creative practitioner. These choices include how much or how little by way of "hallucinated floaters" to include at the periphery of the image. In this image, the choice has been made to include a large amount of the periphery, where the relationship between the NeRF, the subject and the ground truth starts to break down.

If we take the image in as a whole, we notice that there is a lot of change going on. Even though we are viewing a still image, the potential for change is embedded in the confusion of variations in radiance across the image. A change of attention brings focus on a face in profile and the peripheries of the image move to the peripheries of vision. What can be discovered or divined from the noise and the machine "hallucinations"? Moreover, what it reveals to us of the original, ground truth of the images photographic and physical world origins? It does not so much replicate that scene as replicate a moment of seeing, of doing, and of being, in that moment of seeing—as an experience. In the act of viewing the static image, we are both seeing and experiencing a phenomenological act of seeing, as described by Sokolowski. Within the image is embedded the multiple viewpoints of the training data; thus, in viewing it, we experience the full spectrum of bodily involvement in that act of

seeing—the "saccadic movements of the eye, on the ability of the head to be turned, on the ability of the whole body to go from one place to another, from one viewpoint to another"(Sokolowski 2000) are all evident. The physical collection of the training data, combined with the computational processes of the NeRF's creation and, ultimately, the creative practitioner's choices, create an image that moves toward a photographic representation of the phenomenological experience of seeing.

Conclusion

This article has examined NeRF technology and an emergent aesthetic of this technology from the point of view of the original NeRF paper (Mildenhall et al. 2020), and its resulting tools. Since the release of the first paper, a myriad of other approaches have been published[8]. Some of them working towards reducing the emergent aesthetic discussed here (as an anomaly), and others aiming to find more creative ways of using the technology outside of mere verisimilitude. There is great potential for the creative use of these emerging technologies and their approaches to the creation of photographic images. As with any creating practice, a technical knowledge of the materials and methods is vital to a deep engagement with practice. This article develops a connection with the technical as well as the philosophical and conceptual to foster this deep engagement towards developing and working with the emergent aesthetic discussed.

The potential of the NeRF to create photographic images contains within it possibilities that go beyond the representation of the seen to a representation of seeing. This potential sits at an intersection between emerging machine-learning technology, perception, and creativity, with

[8] See for example (NeRF at CVPR 2022 - Frank Dellaert) for a rundown of papers presented at the IEEE / CVF Computer Vision and Pattern Recognition Conference in 2022, covering many alternate approaches to the technology.

the emergent aesthetics of NeRF-generated images being both anchored to, and transcending of, the physical world—providing the creative practitioner with a means of visually exploring the dynamic relationship between seeing and the seen.

The creative exploration of the NeRF in the described practice-based research serves as a bridge between the technical and the experiential, between the objective capturing of scenes and the subjective experience of seeing - as well as the creative process of observing and interacting with machine-learning processes in real time. The images that result from this ongoing intervention in the process, by the practitioner, move towards a visual representation that is not only rooted in the physicality of the captured scene but also in the viewing experience itself—an experience that is extended across the process of capturing a scene in multiple photographs to the intervention in the machine learning processes. This capacity for generating continuous 3D environments from a collection of static images—which can then be remodulated and curated into 2D photographic images—offers a representation (albeit a synthetic representation) of how we might experience the act of seeing. It holds the potential to capture the ephemeral and the elusive in our visual experiences, representing not just what is seen but how it is seen—complete with the peripheral, ongoing and bodily aspects of seeing.

The technical limitations associated with this technology are equally significant, "hallucination" and "ghosts" represent a departure from the verisimilar, emerging during the process as manipulable aspects of the materiality of the NeRF, with the potential to bridge the gap between the lived experience of a seeing a scene and the scenes digital reconstruction. These “specters of the ground truth” emerge from these gaps, a visualization of the unseen’s emergence from the seen.

NeRFs play a significant role in current and future expanded visualities, offering unique opportunities for creative practitioners and image makers to explore this emerging aesthetic and expand the boundaries of photographic practice. By allowing practitioners to modulate and re-modulate data forms, and to incorporate anomalies such as so-called "machine hallucinations" and visual ghosting, they can create images that not only represent the physical environment but also capture how we experience seeing. This brings with it a potential for further exploring photography's role in representing the nature of visual perception and the phenomenological experience of seeing.

BIBLIOGRAPHY

Berger, John (1997), *Ways of Seeing: Based on the BBC Television Series with John Berger; a Book Made*, 37. pr., 1. publ. 1972 by British Broadcasting Corp. and 1977 by Penguin Books, London: British Broadcasting Corp.

Chen, Zhiqin and Zhang, Hao (2019), *Learning Implicit Fields for Generative Shape Modeling*, http://arxiv.org/abs/1812.02822. Accessed 16 September 2019.

Clemens, Justin and Nash, Adam (2010), *Seven Theses on the Concept of 'Post-Convergence'*, https://www.academia.edu/27057333/Seven_theses_on_the_concept_of_post_convergence.

Crawford, Kate (2021), *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*, New Haven: Yale University Press.

Dellaert, Frank (2022), 'Technical perspective: Neural radiance fields explode on the scene', *Communications of the ACM*, 65:1, pp. 98–98.

Drucker, Johanna (2021), *The Digital Humanities Coursebook: An Introduction to Digital Methods for Research and Scholarship*, Abingdon, Oxon; Routledge.

Harwood, Graham ([2008] 2008), 'Pixel', in M. Fuller (ed), *Software Studies: A Lexicon*, Cambridge, Mass: MIT Press, pp. 213–7.

Hunger, Francis (2023), Unhype Artificial 'Intelligence'! A Proposal to Replace the Deceiving Terminology of AI., Zenodo, https://doi.org/10.5281/zenodo.7524493 . Accessed 12 April 2023.

Instant Neural Graphics Primitives (2022), https://github.com/NVlabs/instant-ngp/blob/e45134b9bcf50d0c04f27bc3ab3cde57c27f5bc8/docs/nerf_dataset_tips.md. Accessed 17 April 2023.

Inwood, Michael ([2015] 2015), ‘The Use and Abuse of Vision’, in A. Cimino and P. Kontos (eds), *Phenomenology and the Metaphysics of Sight*, vol. 13, United States: BRILL, pp. 165–83.

Merleau-Ponty, Maurice and Landes, Donald A. (2012), Phenomenology of Perception, Abingdon, Oxon; New York: Routledge.

Mildenhall, Ben, Srinivasan, Pratul P., Tancik, Matthew, Barron, Jonathan T., Ramamoorthi, Ravi and Ng, Ren (2020), ‘NeRF: representing scenes as neural radiance fields for view synthesis’, *Communications of the ACM*, 65:1, pp. 99–106.

Müller, Thomas, Evans, Alex, Schied, Christoph and Keller, Alexander (2022), ‘Instant neural graphics primitives with a multiresolution hash encoding’, *ACM Transactions on Graphics*, 41:4, p. 102:1-102:15.

NeRF at CVPR 2022 - Frank Dellaert (n.d.), https://dellaert.github.io/NeRF22/. Accessed 10 April 2023.

Palmer, Daniel and Sluis, Katrina (2023), ‘Photography after AI’, *Artlink*, 43:2, pp. 18–27.

Rudin, Cynthia (2019), ‘Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead’, *Nature Machine Intelligence*, 1:5, pp. 206–215.

Shiffman, Daniel (2017), 3.4: Linear Regression with Gradient Descent - Intelligence and Learning https://www.youtube.com/watch?v=L-Lsfu4ab74. Accessed 31 May 2017.

Sokolowski, Robert (2000), *Introduction to Phenomenology*, Cambridge, UK ; New York: Cambridge University Press.

Somaini, Antonio (2022), ‘On the photographic status of images produced by generative adversarial networks (GANs)’, *Philosophy of Photography*, 13:1, pp. 153–164.

Steyerl, Hito (2023), ‘MEAN IMAGES’, *New Left Review*,140/141, pp. 82–97.

Tancik, Matthew, Weber, Ethan, Ng, Evonne, Li, Ruilong, Yi, Brent, Kerr, Justin, Wang, Terrance, Kristoffersen, Alexander, Austin, Jake, Salahi, Kamyar, Ahuja, Abhik, McAllister, David and Kanazawa, Angjoo (2023), *Nerfstudio: A Modular Framework for Neural Radiance Field Development*, http://arxiv.org/abs/2302.04264. Accessed 8 February 2023.

Yiu, Sheung (2022), ‘Excerpts from Everything Is a Projection (2020–present): Digital photography and 3D photogrammetry’, *Philosophy of Photography*, 13:1, pp. 149–60.