An Ethical Framework for Trustworthy Neural Rendering applied in Cultural Heritage and Creative Industries

Lorenzo Stacchio[1] Emanuele Balloni[2] Lucrezia Gorgolione[2] Roberto Pierdicca [2]
Adriano Mancini[2] Emanuele Frontoni[1] Benedetta Giovanola[1] Simona Tiribelli[1]
Marina Paolanti[1] Primo Zingaretti[2]
[1] Università di Macerata, [2] Università Politecnica delle Marche

Abstract

Artificial Intelligence (AI) has revolutionized various sectors, including Cultural Heritage (CH) and Creative Industries (CI), defining novel opportunities and challenges in preserving tangible and intangible human productions. In such a context, Neural Rendering (NR) paradigms play the pivotal role of 3D reconstructing objects or scenes by optimizing images depicting them. However, there is a lack of work examining the ethical concerns associated with its usage. Those are particularly relevant in scenarios where NR is applied to items protected by intellectual property rights, UNESCO-recognised heritage sites, or items critical for data-driven decisions. For this, we here outline the main ethical findings in this area and place them in a novel framework to guide stakeholders and developers through principles and risks associated with the use of NR in CH and CI. Such a framework examines AI's ethical principles supporting the definition of novel ethical guidelines.

1. Introduction

Artificial Intelligence (AI) sparked advancements across various sectors, both in industry and academia. One of the most impacted sectors corresponds to Cultural Heritage (CH) and Creative Industries (CI), often considered as a unique discipline (CCI) [24, 25, 45, 70]. Through AI paradigms, tangible and intangible CCI could be better analyzed, preserved, and promoted, with a positive social impact [45, 70]. In particular, it has facilitated digitization efforts supported by international institutions such as the EU Commission and UNESCO, democratizing accessibility, preservation, and dissemination of culture [24, 90, 91]. Monuments, sites, and intangible traditions such as crafts, and art, are just a few examples of CCI elements that are being preserved through AI.

The advent of Neural Rendering (NR) techniques has dramatically improved this digitization and preservation process considering their ability to reconstruct three-dimensional (3D) objects and scenes, being only optimized on the 2D image that depicts them. In the NR arena, Neural Radiance Fields (NeRFs) and Three-Dimensional Gaussian Splatting (3DGS), are the most adopted paradigms, which can be efficiently applied in different environments with variable illumination settings, with pictures taken in the wild and be optimized with small number of images [1, 31, 60]. This feature allows their adoption to digitize scenes and objects not only in a perfect laboratory environment but also in non-optimal ones (also for those objects that no longer exist, like lost heritage) [27, 38]. However, the application of NR in CCI raised new challenges from an ethical perspective. As a preliminary example, one could argue about the authenticity and (intellectual) property of the digital replica [24, 63]. Issues like this become even more relevant considering elements protected by UNESCO or critics for data-driven decisions [90, 91]. Nevertheless, there is a lack of work that has analyzed the ethical implications surrounding the application of NR to CCI items. This also indicates the lack of a robust framework from which new guidelines and regulations can be derived.

This paper fills this gap by reviewing the primary ethical evidence in this area. It aims to clarify the ethical principles and risks associated with the use of NR in CCI contexts. Our framework attempts to navigate the complex ethical terrain of NR integration, taking into account the well-established principles of trustworthy AI contained in the AI Act, including responsibility, reliability, fairness, sustainability, and transparency [58]. The framework has been designed with consideration of some of the world's most recognized ethical guidelines, such as the European Commission's White Paper on AI [24], the Assessment List for Trustworthy AI (ALTAI) [23], The ICOM Code of Ethics [68], UNESCO's documents on the Recommendation on the Ethics of Artificial Intelligence [90] and report on Cultural and Creative Industries [91]. Such a framework aims to support and enhance the development of NR technologies in CCIs while preserving their intrinsic values and importance. The main contributions of this paper are (i) identify the ethical pitfalls of NR paradigms for CCI; (ii) design and implement a specific ethical framework inspired by globally relevant guidelines; (iii) provide multidisciplinary guidelines for the development of NR solutions, taking into account the risk specificities of NR.

2. Related works

In this section, we provide a thorough review of the current status of the state of the art of NR as applied to CH and CI, while delineating ethical considerations. This section is therefore divided into two distinct but related parts: "Technical State of the Art" and "Ethical State of the Art".

2.1. Technical State of the Art

Traditionally, extracting three-dimensional (3D) models from 2D images has been primarily implemented through conventional geometric methods. These methods rely on established techniques such as photometric consistency and gradient-based features to extract depth cues from visual data [10, 13, 30, 34, 40, 43, 76, 79, 84]. However, recent advances in neural networks laid the path to the development of Neural Rendering (NR) techniques. These techniques are characterized by deep image or video generation methods that provide explicit or implicit control over various scene properties, including camera parameters and geometry. Such models learn complex mappings from existing images to generate new ones [86]. In such a space, two paradigms are emerging: Neural Radiance Fields (NeRFs) and Three-Dimensional Gaussian Splatting (3DGS). These have attracted considerable attention due to their power and speed of reconstruction [47, 62, 64, 86]. NeRFs are implicit neural radiance field representations via multi-layer perceptrons (MLPs), optimized via rendering reconstruction loss over 2D images to learn the complex geometry and lighting of the 3D scene they capture [31, 64]. While primarily recognized for novel view synthesis, NeRFs allow the extraction of 3D surfaces, meshes, and textures [85]. This is achieved through an internal representation as an Occupancy Field (OCF) or a Signed Distance Function (SDF), which can be easily converted into a 3D mesh using conventional algorithms such as the Marching Cubes [56]. Similarly, 3DGS aims to efficiently learn and render high-quality 3D scenes from 2D images [47]. 3DGS introduces a continuous and adaptive framework using differentiable 3D Gaussian primitives, in contrast to traditional volumetric representations such as voxel grids. These primitives parameterize the irradiance field, allowing novel views to be generated during rendering. 3DGS achieves real-time rendering through a tile-based rasterizer, unlike NeRF which relies on computationally intensive volumetric ray sampling [47, 88]. Both NeRFs and 3DGS are self-supervised and can be trained using only multi-view images and their corresponding poses, eliminating the need for 3D/depth supervision (using algorithms such as Structure from Motion to extract camera poses). In addition, they generally deliver higher photorealistic quality compared to traditional novel view synthesis methods [31]. These factors make them suitable for various applications in different domains, especially in the context of CCI, where the generation of the most faithful representation is key. NeRF has recently been considered for CH applications for different contexts and data, such as those collected with smartphones or professional cameras, in different environments [8, 18, 60]. At the same time, they have been used in the context of CI, mainly for industrial design, and various fashion applications, such as 3D object reconstruction and human generation [26, 29, 74, 93, 100]. Notwithstanding its newness, 3DGS has also been considered and applied in CH, where it was compared with NeRF for the reconstruction of real monuments, and also in CI, where it was used to efficiently generate dressed humans [1, 11]. Although not specifically applied to the CCI context, few-shot approaches amount to a variation of NR that can be optimized for the 3D representation of scenes and objects by using only a few frames (typically 1 to 10) [48, 55, 101]. Such approaches can be adopted for those CCI objects that can no longer be captured and are stored in a small number of images, but also for those objects that can only be captured from a limited set of views. In such a context, relevant works amount to PixelNeRF, which introduces an approach that preserves the spatial alignment between images and 3D representations by learning a prior over different input views [102]. In contrast, models such as DietNeRF, RegNeRF, InfoNeRF, and FreeNerf address few-shot optimization without relying on knowledge, instead employing optimization and regularisation strategies along with auxiliary semantic losses [42, 48, 101].

2.2. Ethical State of the Art

As mentioned in the introduction, while the ethical implications of AI in CCI have been explored, the ethical implications of using NR paradigms have been poorly explored. For this, works such as [33, 45, 70, 71, 73, 81] draw on important and relevant sources of knowledge regarding the ethics of AI, and in some cases, the ethics of generative AI. In particular, [45] outlined the implications of AI across sectors on a global scale, sparking debates about the ethical principles that should guide its development and use. Concerns include potential job displacement, misuse by malicious actors, accountability issues, and algorithmic bias. It also highlights efforts to engage different stakeholders, including public and private companies, questions about their motivations, and the convergence of ethical principles. Finally, it discusses the main ethical principles currently analyzed in AI ethics, while delineating guidelines to develop fair and trustworthy systems. Specifically, CH [70, 71] analyzed ethical concerns regarding the use of AI's role in activities such as creating digital replicas or providing unbiased explanations of artworks. They also developed an ethical framework for these activities, including relevant ethical principles such as shared responsibility, meaningful participation, and accountability. Their findings underscore the need to develop sector-specific ethical guidelines for AI in both tangible and intangible CH to ensure its sustainable development while preserving its values, meaning, and social impact. In the context of CI, Flick et al. [28] pointed out the urgency of defining ethical rules and exploring issues of ownership and authorship, biases in datasets, and the potential dangers of non-consensual deepfakes. In the same context, in [81], the authors analyzed the lack of ethical discussion around generative AI, particularly around biases, while exploring their implications from a sociocultural art perspective. Their findings analyzed how generative AI models showed biases towards artists' styles that were also present in the training data. We should also reflect on public primary sources of global AI ethical significance, to establish a robust AI ethical framework about NR. To this end, we rely first on the Assessment List for Trustworthy Artificial Intelligence (ALTAI) developed by the European Commission's High-Level Group on Artificial Intelligence (implemented by HLEGAI in 2019) [23]. ALTAI identifies seven requirements necessary to achieve trustworthy AI, covering aspects such as human oversight, technical robustness, privacy, transparency, fairness, societal wellbeing, and accountability. It is important to note that these ethical imperatives are regulative, not legally binding, and serve as guiding principles for the responsible development of the technology. Second, UNESCO's Recommendation on the Ethics of Artificial Intelligence and the Readiness assessment methodology provides systematic regulatory and evaluation guidance with a globally sensitive perspective to guide companies in responsibly managing the impact of AI on individuals and society [89, 90]. These recommendations emphasize bridging digital and knowledge gaps among nations throughout the AI lifecycle and precisely define the values guiding the responsible development and utilization of AI systems. In line with the EU guidelines, UNESCO emphasizes 'transparency and accountability' as key principles for trustworthy AI. Transparency guarantees that the public is informed when AI systems influence policy decisions, promoting comprehension of their significance. This transparency is essential to ensure equity and inclusivity in the outcomes of AI-based systems. Explainability refers to understanding how different algorithmic pipelines work, from the received input data to their processed outputs. . We also considered the European Commission's White Paper on AI [24], which highlights the importance of a European approach to the development of AI, based on ethical values and aimed at promoting benefits while addressing risks. In particular, it outlines the need for the trustworthiness of AI systems based on European values and fundamental rights such as human dignity and privacy. It provides a regulatory and investment-oriented approach to address the ethical risks of AI, focusing on building an ecosystem of excellence and trust throughout its lifecycle. We then considered specifically the CCI context of our research, starting with the ICOM Code of Ethics and Museums [68], which defines ethical standards on issues specific to museums and provides standards of professional practice that can serve as a normative basis for museum institutions. Such a code begins with a position statement that explains the purpose of museums and their responsibilities. It then focuses on the specific challenges faced by museums, including (i) the responsibility to safeguard both tangible and intangible natural and cultural heritage, while protecting and promoting this heritage within the human, physical, and financial resources allocated; (ii) the acquisition, conservation, and promotion of collections as a contribution to the preservation of heritage; (iii) the provision of access to, interpretation of and promotion of heritage; (iv) the definition of policies to preserve the community's heritage. (iv) to define policies for the conservation of community heritage and identity. Again in the context of the CCI, we considered the well-known artists' associations' specification of the EU AI law dedicated to the creative arts, including safeguards that require rights holders to be specifically [73, 92]. Such a document, issued by 43 unions representing creative authors, performers, and copyright holders, emphasizes the urgent need for effective regulatory measures to deal with generative AI. In particular, the document highlights how existing measures are insufficient to protect the digital ecosystem and society at large. It sets out requirements for providers of foundational models, including transparency about training materials, their accuracy and diversity, and compliance with legal frameworks for data collection and use. These proposals aim to ensure the responsible development and deployment of generative AI systems while protecting against potential harms such as misinformation, discrimination, and infringements of privacy and copyright. Finally, we have included in our analysis the UNESCO document on Cultural and Creative Industries in the COVID-19 era [91]. This document was one of the first to analyze the impact of the pandemic by exploring the use of digital technologies by audiences and cultural professionals in the CCIs, which are now becoming pervasive, particularly in the visual industries, and which can be analyzed through an ethical lens.

3. Methodology

In this section, we present a detailed approach for the analysis of ethical pitfalls within NR techniques in CCI. On top

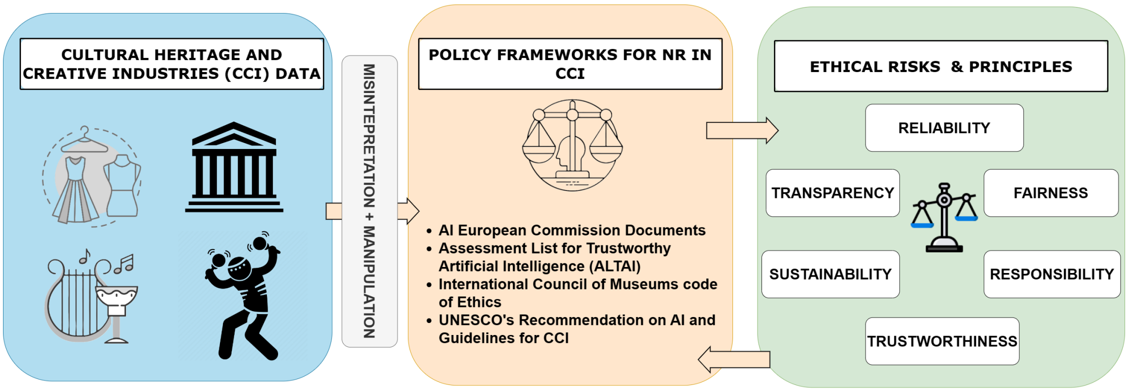

Figure 1. Ethical Workflow for NR in CCI applications. The approach starts from the selection of the NR methodology and data where they are applied. Then the ethical framework defines principles that must be evaluated to avoid risks, providing a conclusive analysis and quantification of ethical compliance scores.

of this analysis, we defined an ethical framework for assessing the trustworthiness of such techniques, given the lack of work on this topic. Our study begins with an analysis of the scientific literature on neural rendering approaches, focusing on NeRFs and 3DGS. From this research, we derived the technical challenges of their application in CCI. Then, considering these challenges, we examined ethical documents issued by public and globally relevant issuers and scientific literature. Through these, we highlight the key ethical risks that these technical challenges may pose, along with their associated and well-established principles. Following these documents and reported guidelines, we have selected those principles and risks that can be linked to specific NR challenges that could help mitigate them. The result of this process, visually illustrated in Figure 1, is a novel ethical framework that aims to build NR systems with a trustworthy approach.

3.1. Challenges and Opportunities of Neural Rendering in CCI

Considering NR, in particular, possible challenges, and technical risks may arise for the specific elements in the CCI domain. These challenges include but are not limited to (i) Understanding complex AI models and validating the data collection process; (ii) Ensuring the accuracy of reconstructions; (iii) Demonstrating stability and generalization in different (social) environments; (iv) Unbiased and fair results; (v) Ethical data ownership; (vi) Minimizing environmental impact. Considering (i) significant challenges arise as the lack of interpretability of those NR models that expose knowledge prior or are being conditioned on models with prior knowledge [39, 102]. Moreover, missing descriptions of the data acquisition steps hinder accountability and a data-driven decision-making approach [78]. These challenges underline the importance of developing methods and tools to improve transparency, interpretability, and accountability in NR systems [9, 14, 39, 96]. Other technical challenges and risks related to confidence in the accuracy/fidelity of the reconstructions and the consistency of the outputs generated (ii). Inconsistent outputs could be generated due to few-shot learning approaches, in-the-wild datasets, or data corruption, requiring rigorous testing and validation procedures [59, 87, 102]. Validation of the consistency and fidelity of NR input and synthesized data in different domains is a crucial challenge to define generalized and reliable systems. Stability and generalization across different (social) environments (iii) could also be defined as an issue, considering that NR methods may lack visual generalization and inconsistent geometric representations, which are significant barriers to achieving robust performance in diverse CCI contexts [17, 18, 60]. This phenomenon could happen while optimizing an NR in a few-shot or an incomplete set of scene views. A possible solution to cope with such phenomena amounts to adopting few-shot architectures or pre-trained models [48, 101, 102]. In this particular case, however, (iv) we should consider the kind of architectural approach followed by those few-shot networks (e.g., overlook high-frequency details [101]) and the bias that those pre-trained models expose in their knowledge priors [15, 48, 102]. Such models, along with biases that could emerge within the data collection process, highlight the importance of developing methods that mitigate bias and promote equitable results [105]. It is also worth mentioning the criticalities (v) that emerge while discussing already considered challenges like misuse of input and generated data and unfaithful generation in the context of data ownership and responsibility [5, 16]. The ownership of the NR 3D-generated items entails the rightful possession of data and the responsibility to ensure usage and protection against misuse [70]. Data misuse poses a great risk, ranging from unauthorized reproduction to malicious manipulation, that could be applied in NR to generate unfaithful items [39], damaging stakeholders that have economical or emotive interest in them [70]. For example, if some views or geometric structures of the 3D models reconstructed by NR methods are inconsistent with reality, one could argue about their authenticity and also debate their intellectual property [57]. All of these aspects define the urgency of integrating social considerations into the system functionality, requiring careful human validation protocols [83]. Finally, (vi) NR lays significant risks for the environment [51, 75]. Sustainability is a critical aspect associated with the high computational demand of NR processes, and the energy used to maintain ready-to-visualize renderers [94]. Moreover, the indirect energy costs stemming from activities such as professional digital photography waste, creating photo capture settings, data transmission, and storage further contribute to the environmental footprint of NR paradigms.

3.2. Ethical Principles of Neural Rendering in CCI

Our study begins with a review of guidelines from key regulatory frameworks, including the Assessment List for Trustworthy Artificial Intelligence (ALTAI) [23], the UNESCO Recommendation on the Ethics of Artificial Intelligence [90], and the European Commission's White Paper on AI [24]. We also thoroughly analyzed [45, 73, 81], which provides a global mapping of AI regulations and robust ethical principles. Then, given the CCI context of our investigation, we considered the ICOM Code of Ethics and Museums [68] and a specification of the AI act for the creative arts [73, 92]. We have also included in our analysis the UNESCO on Cultural and Creative Industries in the face of COVID-19 [91]. Following these documents and reported guidelines, we selected specific ethical principles to develop a framework to be applied concerning the usage of NR in CCI. In the following, we highlight the ethical principles considered and how they connect to the technical challenges listed in the previous Section 3.1.

Responsibility One of the most relevant ethical principles that should be recognized for a trustworthy application of NR in CCI is responsibility. Responsibility refers to the moral obligation of individuals, organizations, and societies to ensure that AI technologies are developed, deployed, and used in ways that respect and preserve cultural heritage and promote the well-being of individuals and communities involved in creative endeavors [45]. It is worth highlighting that concerning NR, actions taken from data capturing to model training, evaluation, and deployment, rely on the different stakeholders (e.g., data generator, data owners, ML engineers). For this reason, the accountability of the action taken through NR is addressed to both engineers as well as cultural managers or creative professionals [32, 68]. For this, a multidisciplinary approach is required to ensure NR accountability, defining policies to co-create and evaluate processes and results. Such principle should also be applied to input data to NR models and those that are instead generated, providing adherence to ethical guidelines throughout the entire data lifecycle [70]. This includes transparent documentation of data sources, data usage consent, and robust security measures to safeguard against misuse [70]. Furthermore, ensuring the authenticity of generated content is essential to uphold trust and credibility in NR systems, particularly in applications where the generated output may influence decision-making or perception, such as replication of UNESCO-protected material or Digital Twins real-time monitoring [16, 21, 44, 53, 57, 82, 89]. With data ownership and compliance against ethical principles, stakeholders can mitigate the risks associated with data misuse and unfaithful generation. Responsibility towards real and generated data ownership and legal liability for unfaithful ones is essential to maintain the integrity of NR applications [45, 70].

Transparency and Explainability Transparency and Explainability are core principles in the development of NR systems, in particular, to define accountability and trustworthiness in CCI. Transparency involves the clear and open communication of processes, algorithms, data, and outcomes associated NR [45], enabling stakeholders to understand how decisions are made and assess potential biases or limitations [28, 70]. Explainability regards instead the ability of NR systems to provide understandable explanations for their synthesis and 3D model extraction [97]. For example, NR produces an incorrect visual representation of a real-world facility, and the influenced stakeholders must be able to understand the reason [12]. Considering the complexity of NR approaches for novel view synthesis and 3D object rendering and their implicit black-box structure, the adoption of explainability techniques for their analysis is required. For example, different mechanisms like visualizing the learned geometrical structure, saliency maps, interpreting network activations, or analyzing the influence of input parameters on the rendered images are all techniques that could be adopted to support NR [52, 67, 77]. In particular, such approaches could support the improvement of such systems, from both an architectural or data-centric perspective, detecting biases, but also comparing different models according to their learned features [52, 67, 77]. Such aspects are all crucial in the context of CCI, where an enormous tangible and intangible patrimony could now get digitized thanks to NR paradigms in a cheap and fast way [18]. For these reasons, is it crucial to tackle the aforementioned challenges to elucidate the inner workings of such algorithms to ensure that their decisions are understandable and accountable to guarantee NR reliability, fairness, and impact.

Reliability Reliability refers to the ability of AI applications to comply with data protection providing high accuracy and completeness considering both input datasets used to develop and train the models, and their outcomes [23, 45]. For NR to be reliable, we should first consider the completeness of the data. We should, in general, acquire around 50 and 150 pictures based on the object complexity, following a spherical omnidirectional approach to optimize NR methods [65]. Even having at our disposal such pictures, and techniques to extract 3D geometric structures from the optimized networks, like marching cubes for NeRFs and Poisson reconstruction from point clouds, which could discard several high-frequency details [36]. Moreover, such a quantity of pictures could not be available for different CCI items (due to objects that do not exist anymore, or that can't be moved to be captured from all sides [38]). Even adopting few-shot NR architectures [66, 101] we should have at our disposal, 3 to 9 sparse viewpoints to have reasonable, but noncomparable quantitative-qualitative results. To visually explore such a concept, we re-trained one of the SOTA for few-shot NeRFs, named FreeNeRF [101], using the same 3-image setting reported by the authors, depicting the results in Figure 2. As can be qualitatively appreciated, different parts of the synthesized views present artifacts and incomplete geometrical structures. It is worth highlighting that such artifacts were verified on pictures taken in a controlled laboratory setting, with fixed illumination and camera poses. This raises ethical concerns related to the missing data biases, i.e., the lack of data from underrepresented regions, cultures, and objects [70]. Such bias could negatively influence the training of NR, creating distorted geometries and textures [102]. Such bias also involves camera pose estimation, which is a necessary step for NR in case pictures were taken with classical RGB cameras [69]. In particular, this raises two ethical concerns: (i) camera pose estimation algorithms could provide inaccurate estimation or (ii) non-converge. Such situations mostly regard cases of few-shot settings with low scene coverage and in-the-wild settings [8, 20, 41, 59]. Recent methods based on Diffusion Models are emerging, with preliminary results towards a few image camera pose estimations, which however only work on fixed environmental conditions [103]. Considering these concerns, rigorous quantitative and qualitative validation of the fidelity of collected/generated data is necessary to determine the reliability of NR.

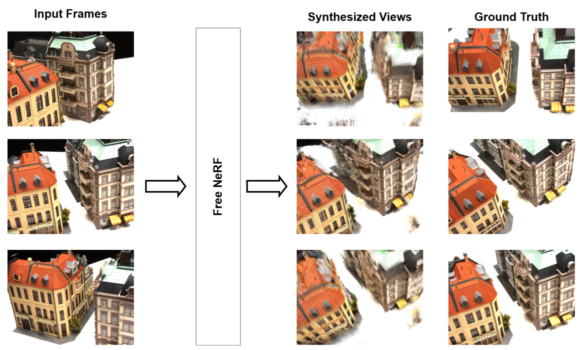

Figure 2. FreeNeRF [101] trained on 3 images from the DTU dataset with the same setting provided by the original authors and three synthesized novel views compared against their ground truths.

Trustworthiness Trustworthiness refers to the capacity of AI systems to be ethical towards transparency, accountability, and respect for human values and rights [45]. A trustworthy system not only produces accurate and consistent results but also operates in a manner that aligns with ethical principles and societal expectations [24]. Considering such a large definition, we here contextualize the trust in NR paradigms, in terms of technical robustness (the ability of the system to function reliably and effectively), and social robustness (the ability of the system to integrate and operate ethically in different social contexts) [70, 72]. Such models must demonstrate stability and reliability in their predicted generations maintaining coherent performances, most of all in use cases related to CCI, where complex objects, dresses, buildings, and variable illumination conditions would be aspects of their everyday usage [70]. Such a principle is strongly bonded and shares the same reflections of reliability and responsibility. To demonstrate trust, novel empirical frameworks should be defined to take into account the performance of such models in extreme cases (e.g., strong luminance, one-shot settings), where there is missing information about the scene or the object we want to reconstruct [19].

Sustainability The ethical dimensions of sustainability represent a critical focal factor within contemporary AI research and development [23, 90]. Central to this discourse is a comprehensive understanding of the environmental impact and optimizing resources for models' lifecycle, spanning data collection, model training, and deployment phases. NR research should so analyze the environmental footprint stemming from various computational activities integral to model development and rendering pipelines [45]. Data collection, iterative model training procedures, and model deployment exert considerable energy demands [49, 75]. Smart capture data setting and training strategies should be adopted to define computationally efficient processes to minimize energy waste [37]. For example, intelligent protocols could be adopted to reduce the number of cameras and/or GPU processing techniques for camera pose estimation [98]. Also, indirect sources of energy consumption activities like human photographer transportation, picture capture settings, digital photography, data transmission, and storage should be taken into account [7]. Considering model training and deployment, relevant efforts should involve the refinement of model architectures to optimize computational efficiency, taking into account the usage of lower-image resolutions to reduce memory and teraflops, the exploitation of optimized hardware systems, and the adoption of renewable energy sources. In particular considering model architectures, distillation or quantization techniques could be adopted to optimize NR training and deployment [35, 80]. Sustainable practices are necessary to reduce these impacts and promote environmental responsibility. This includes optimizing models, training pipelines, and infrastructure to minimize energy consumption, considering the environmental implications at every stage of the NR workflow. By prioritizing sustainability in development and deployment, stakeholders can minimize the environmental footprint of NR and contribute to a more sustainable digital ecosystem.

Fairness Fairness in AI encompasses justice, consistency, inclusion, equality, non-bias, and non-discrimination, which denotes principles and equitable treatment of individuals and communities [23, 90]. Also, NR systems must ensure their rights, dignity, and opportunities are upheld and respected [45]. Considering such principle, NR should produce consistent results that are unbiased and fair across different demographics, environments, and scenarios. Such principles are particularly at risk when considering NR methods with prior knowledge, or those that exploit regularization and optimizations for few or one-shot settings (e.g. synthetic generation from other views or ignore high-frequency details) [55, 66, 101, 106]. Ensuring unbiased and fair outcomes for NR necessitates so careful consideration of potential biases introduced during pre-training, which can influence the generation of outputs in ways that exacerbate existing inequalities or inaccuracies [55, 101]. This bias may amount to cultural, social, or historical ones, inherent in the training data or underlying assumptions embedded within the model architecture, especially for unrepresented items [3, 99]. At the same time, NR architectures that exploit strategies for few or one-shot settings (e.g., overlook high-frequency details or synthesize novel 3d views) [54, 55, 66, 101] can contribute to disparities in the representation and depiction of scenes or objects within NR outputs (creating similar phenomena to the one depicted in Figure 2). Such oversights may affect certain features or characteristics, leading to biased or unfair outcomes, mostly in contexts where high-frequency detail is essential for accurate representation (e.g., dance, fashion, art). To ameliorate these phenomena, data quality and bias analysis must be performed, along with bias examination of the pre-trained knowledge learned by the models. Moreover, technical improvements in architectures, optimization losses, regularization, and generative models should be fostered (in particular considering domain adaption paradigms [46]). This holistically includes rigorous evaluation and validation of biases, as well as the incorporation of diversity considerations into model design and development. To this date, a patch-wise level combination of quantitative metrics should be applied, like combining PSNR, LPIPS, and MSE for novel view synthesis and DICE, DMax, ASDlike for 3D meshes and Chamfer, Hausdorff, and Earth-Mover's distances for synthesized point clouds [6, 22, 61, 104]. The focus of such quantitative analysis should in particular regard cases where limited training data (few or one shot) are employed, considering that several artifacts could be generated and a small change in the input data can lead to significantly different representations [66, 101].

4. Results and Discussions

We here summarise the key ethical principles and challenges of NR in CCI in a framework, highlighting technical risks we aim to mitigate. Table 1 schematically reports the findings produced from our investigation. The ethical documents and the scientific literature acted as mediators, bridging data related to CCI and AI ethical principles to key ethical risks of NR applied to them, providing a robust basis for defining fair regulations. In particular, considering CCI items that naturally exhibit ethical issues like bias, fairness, and responsibility and are prone to define reliability concerns. Such an ethical framework, should in principle support stakeholders in the individuation of principles and responsibilities that should be considered when designing, implementing, monitoring, and evaluating NR in CCI.

Several solutions can be implemented to reduce the identified ethical risks associated with NR. First, addressing the challenge of transparency and explainability requires comprehensive documentation of the data collection process and efforts to improve the description and interpretability of NR generative pipelines. In such a context, we could also use well-established explainability paradigms [2, 95] to describe how NR models learn from data and generate outputs. To mitigate reliability-related risks, rigorous testing and validation protocols should be established to verify the accuracy and consistency of NR reconstructions, in high-variance settings, including extreme cases (e.g., poor lighting., and occluded objects). However, reliability goes beyond technical stability and includes social aspects that should be considered developing NR model. It becomes mandatory to include social factors such as cultural sensitivity, and historical accuracy narration, collaborating with domain experts in a multi-disciplinary approach [70]. Sustainability concerns can be addressed by optimizing multi-camera hardware, and model architectures while adopting energy-efficient optimization algorithms and hardware systems [4]. In particular, adopting energy-efficient algorithms

Table 1. Neural Rendering ethical principles, challenges, and risks detailed starting from the designed ethical framework.

Ethical Principle	Challenges	Technical Risks	Detailed Explanation
Transparency and Explainability	Understanding complex AI models and validate data collection process	- Lack of interpretability - Missing description of data collection steps - Lack of controllability for erroneous reconstructions	Understanding data collection process and how NR models learn from data and produce their outputs. NR approaches require additional efforts to elucidate the inner workings of comprehensibility and accountability. Transparency is crucial to understand decision-making processes.
Reliability	Ensuring accuracy of reconstructions	- Inconsistent outputs due to few or one-shot; - Hard camera estimation due to data scarcity; - Novel view synthesis and geometrical outputs with low veridicity - Bias of pre-trained NR methods	Ensuring the accuracy of input data and generated reconstructions is crucial in CCI context. Validation frameworks applying quantitative-qualitative analysis should be designed to measure the consistency and fidelity of the generation and perform bias analysis. At the same time, novel models should be defined to reconstruct camera poses for a few shot settings, considering objects that do not exist anymore.
Trustworthiness	Demonstrating stability and generalization in different (social) environments	- Lack of visual generalization - Inconsistent Geometrical Representation - Missing social considerations into the system's functionality	Building trust in the stability and generalization capabilities of NR models. Trust depends on technical robustness and the ability to generalize. Novel empirical frameworks are needed to demonstrate reliability and build user confidence. It should also take into account the social dimension (i.e. ability to be applied in different social contexts).
Sustainability	Minimizing environmental impact	- High computational demand - Energy cost to create and maintain a capture setting	Considering the environmental impact and economical aspects of NR. Sustainable practices, such as optimizing model architectures, and green computing infrastructure are necessary to reduce environmental footprint. Also, indirect energy costs like picture capture setting, digital photography energy waste, transport, and energy consumption for data transmission and storage.
Fairness	Unbiased and fair results	- Biased NR prior knowledge - Artifacts caused by NR paradigms which exploit regularization, synthetic generation or ignore high-frequency details	Ensuring unbiased and fair outcomes for NR with prior knowledge. Addressing biases introduced during training is crucial as they can propagate through the model and affect generated outputs (considering NR methods that have prior knowledge). Fairness is especially at risk in cases of limited training data and/or integration of auxiliary networks.
Responsibility	Ethical data ownership an authenticity	- Misuse of generated data - Accountability for unfaithful generation - Intellectual property	Upholding ethical data ownership, intellectual property, usage, and authenticity. Responsible data ownership and adherence to ethical guidelines are essential to maintain the integrity and legality of applications.

involves implementing techniques such as model pruning, quantization, and compression, which reduce the computational workload with an often negligible loss in performances [50]. Ensuring fairness requires careful consideration of addressing biases in training data and model architectures. For example, imagine digitizing ancient sculptures from various civilizations for virtual museum exhibits. Biases in the training data, such as a disproportionate focus on artifacts, could led the NR model to prioritize reconstructions of artifacts from dominant cultures, neglecting others. To mitigate biases, we can curate a diverse training dataset, including artifacts from different cultures, periods, and geographical regions. The integration of auxiliary networks to detect and correct biases in the rendering process can improve the fairness of NR outputs. Finally, responsibility in NR requires ethical data ownership practices, protection against misuse of generated data, and ensuring faithful generation following ethical guidelines and legal frameworks. This could include the implementation of encryption protocols, and data anonymization techniques to protect the integrity and confidentiality of digitized objects.

5. Conclusions and future works

Our research explored the use of NR in CCI focusing on the ethical considerations and relevant legal frameworks that pertain to these domains. The output of this process is a new ethical framework that serves as a guide for addressing the potential ethical risks identified in our analysis and provides a structured approach for ethical decision-making in the context of NR applications in CCI. We have further elaborated on the specific ethical principles that should be prioritized since they are crucial to ensure the responsible use of NR. We also highlighted ethical pitfalls that require clear guidelines to protect the integrity and sustainability of CCI sectors when applying NR technologies. For future work, we will define ethical measurable standards, criteria, and metrics to quantify the ethicality level of different NR methodologies in CCI. Such a qualitative-quantitative approach will be rigorously validated across diverse contexts, assessing its adaptability and resilience.

6. Acknowledgement

This work is funded by PESARO CTE SQUARE; CUP D74J22000930008, FSC MISE 2014-2020.

References

[1] Rameen Abdal, Wang Yifan, Zifan Shi, Yinghao Xu, Ryan Po, Zhengfei Kuang, Qifeng Chen, Dit-Yan Yeung, and Gordon Wetzstein. Gaussian shell maps for efficient 3d human generation. *arXiv preprint arXiv:2311.17857*, 2023. 1, 2

[2] Rasheed Omobolaji Alabi, Mohammed Elmusrati, Ilmo Leivo, Alhadi Almangush, and Antti A Mäkitie. Machine learning explainability in nasopharyngeal cancer survival using lime and shap. *Scientific Reports*, 13(1):8984, 2023. 7

[3] Marcellus Amadeus, William Alberto Cruz Castañeda, André Felipe Zanella, and Felipe Rodrigues Perche Mahlow. From pampas to pixels: Fine-tuning diffusion models for ga\'ucho heritage. *arXiv preprint arXiv:2401.05520*, 2024. 7

[4] Lasse F Wolff Anthony, Benjamin Kanding, and Raghavendra Selvan. Carbontracker: Tracking and predicting the carbon footprint of training deep learning models. *arXiv preprint arXiv:2007.03051*, 2020. 7

[5] Omri Avrahami and Bar Tamir. Ownership and creativity in generative models. *arXiv preprint arXiv:2112.01516*, 2021. 4

[6] Yunpeng Bai, Chao Dong, Cairong Wang, and Chun Yuan. Ps-nerv: Patch-wise stylized neural representations for videos. In *2023 IEEE International Conference on Image Processing (ICIP)*, pages 41–45. IEEE, 2023. 7

[7] Alpha Yaya Balde, Emmanuel Bergeret, Denis Cajal, and Jean-Pierre Toumazet. Low power environmental image sensors for remote photogrammetry. *Sensors*, 22(19):7617, 2022. 7

[8] E Balloni, L Gorgoglione, M Paolanti, A Mancini, and R Pierdicca. Few shot photogrametry: A comparison between nerf and mvs-sfm for the documentation of cultural heritage. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 48:155–162, 2023. 2, 6

[9] Pablo Barceló, Mikaël Monet, Jorge Pérez, and Bernardo Subercaseaux. Model interpretability through the lens of computational complexity. *Advances in neural information processing systems*, 33:15487–15498, 2020. 4

[10] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics (TOG)*, 28(3):24, 2009. 2

[11] A Basso, F Condorelli, A Giordano, S Morena, and M Perticarini. Evolution of rendering based on radiance fields. the palermo case study for a comparison between nerf and gaussian splatting. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 48:57–64, 2024. 2

[12] Pankaj Bhambri and Alex Khang. The human-machine nexus with art-making generative ais. In *Making Art With Generative AI Tools*, pages 73–85. IGI Global, 2024. 5

[13] Michael Bleyer, Christoph Rhemann, and Carsten Rother. Patchmatch stereo-stereo matching with slanted support windows. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2011. 2

[14] Giulio Cainelli, Valentina Giannini, and Donato Iacobucci. How local geography shapes firm geography. *Entrepreneurship & Regional Development*, 34(9-10):955–976, 2022. 4

[15] Ang Cao, Chris Rockwell, and Justin Johnson. Fwd: Real-time novel view synthesis with forward warping and depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15713–15724, 2022. 4

[16] Lifeng Chen, Jia Liu, Yan Ke, Wenquan Sun, Weina Dong, and Xiaozhong Pan. Marknerf: Watermarking for neural radiance field. *arXiv preprint arXiv:2309.11747*, 2023. 4, 5

[17] Francesca Condorelli, F Rinaudo, F Salvadore, and S Tagliaventi. A comparison between 3d reconstruction using nerf neural networks and mvs algorithms on cultural heritage images. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 43: 565–570, 2021. 4

[18] V Croce, G Caroti, L De Luca, A Piemonte, and P Véron. Neural radiance fields (nerf): Review and potential applications to digital cultural heritage. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 48:453–460, 2023. 2, 4, 5

[19] Ziteng Cui, Lin Gu, Xiao Sun, Xianzheng Ma, Yu Qiao, and Tatsuya Harada. Aleth-nerf: Low-light condition view synthesis with concealing fields. *arXiv preprint arXiv:2303.05807*, 2023. 6

[20] Matteo Cutugno, Umberto Robustelli, and Giovanni Pugliano. Structure-from-motion 3d reconstruction of the historical overpass ponte della cerra: A comparison between micmac® open source software and metashape®. *Drones*, 6(9):242, 2022. 6

[21] Alexandra Dashkina, Ludmila Khalyapina, Aleksandra Kobicheva, Tatiana Lazovskaya, Galina Malykhina, and Dmitriy Tarkhov. Neural network modeling as a method for creating digital twins: From industry 4.0 to industry 4.1. In *Proceedings of the 2nd International Scientific Conference on Innovations in Digital Economy*, pages 1–5, 2020. 5

[22] Nessrine Elloumi, Habiba Loukil Hadj Kacem, Nilanjan Dey, Amira S Ashour, and Med Salim Bouhlel. Perceptual metrics quality: Comparative study for 3d static meshes. *International Journal of Service Science, Management, Engineering, and Technology (IJSSMET)*, 8(1):63–80, 2017. 7

[23] European Commission. Assessment List for Trustworthy Artificial Intelligence (ALTAI) for self-assessment. `https://digital-strategy.ec.europa.eu/it/node/806`, 2024. 1, 3, 5, 6, 7

[24] European Commission. White Paper on Artificial Intelligence: A European Approach to Excellence and

Trust. https://commission.europa.eu/publications/white-paper-artificial-intelligence-european-approach-excellence-and-trust_en, 2024. 1, 3, 5, 6

[25] European Commission. Cultural heritage and cultural and creative industries (ccis). https://research-and-innovation.ec.europa.eu/research-area/social-sciences-and-humanities/cultural-heritage-and-cultural-and-creative-industries-ccis_en, Year of access. Accessed: March 17, 2024. 1

[26] Lidia Fabra, J Ernesto Solanes, Adolfo Muñoz, Ana Martí-Testón, Alba Alabau, and Luis Gracia. Application of neural radiance fields (nerfs) for 3d model representation in the industrial metaverse. *Applied Sciences*, 14(5):1825, 2024. 2

[27] Gabriele Fangi, Wissam Wahbeh, Eva Savina Malinverni, Francesco Di Stefano, and Roberto Pierdicca. Documentation of syrian lost heritage: From 3d reconstruction to open information system. 2022. 1

[28] Catherine Flick and Kyle Worrall. The ethics of creative ai. In *The Language of Creative AI: Practices, Aesthetics and Structures*, pages 73–91. Springer, 2022. 3, 5

[29] Yonggan Fu, Zhifan Ye, Jiayi Yuan, Shunyao Zhang, Sixu Li, Haoran You, and Yingyan Lin. Gen-nerf: Efficient and generalizable neural radiance fields via algorithm-hardware co-design. In *Proceedings of the 50th Annual International Symposium on Computer Architecture*, pages 1–12, 2023. 2

[30] Yasutaka Furukawa and Jean Ponce. Accurate, dense and robust multiview stereopsis. *IEEE transactions on pattern analysis and machine intelligence*, 32(8):1362–1376, 2010. 2

[31] Kyle Gao, Yina Gao, Hongjie He, Dening Lu, Linlin Xu, and Jonathan Li. Nerf: Neural radiance field in 3d vision, a comprehensive review. *arXiv preprint arXiv:2210.00379*, 2022. 1, 2

[32] Valentina Giannini and Donato Iacobucci. The role of internal capital market in business groups. In *The Palgrave Handbook of Managing Family Business Groups*, pages 49–64. Springer, 2022. 5

[33] Valentina Giannini, Donato Iacobucci, and Francesco Perugini. Local variety and innovation performance in the eu textile and clothing industry. *Economics of Innovation and New Technology*, 28(8):841–857, 2019. 2

[34] Michael Goesele, Noah Snavely, Brian Curless, Hugues Hoppe, and Steven M Seitz. Multi-view stereo for community photo collections. In *Proceedings of the IEEE 11th International Conference on Computer Vision (ICCV)*, pages 1–8. IEEE, 2007. 2

[35] Cameron Gordon, Shin-Fang Chng, Lachlan MacDonald, and Simon Lucey. On quantizing implicit neural representations. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 341–350, 2023. 7

[36] Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. *arXiv preprint arXiv:2311.12775*, 2023. 6

[37] Puren Guler, Deniz Emeksiz, Alptekin Temizel, Mustafa Teke, and Tugba Taskaya Temizel. Real-time multi-camera video analytics system on gpu. *Journal of Real-Time Image Processing*, 11:457–472, 2016. 6

[38] Hariz Halilovich. Re-imaging and re-imagining the past after 'memoricide': intimate archives as inscribed memories of the missing. *Archival Science*, 16(1):77–92, 2016. 1, 6

[39] Ayaan Haque, Matthew Tancik, Alexei A Efros, Aleksander Holynski, and Angjoo Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19740–19750, 2023. 4, 5

[40] Heiko Hirschmuller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):328–341, 2008. 2

[41] Jakob Iglhaut, Carlos Cabo, Stefano Puliti, Livia Piermattei, James O'Connor, and Jacqueline Rosette. Structure from motion photogrammetry in forestry: A review. *Current Forestry Reports*, 5:155–168, 2019. 6

[42] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5885–5894, 2021. 2

[43] Michal Jancosek and Tomas Pajdla. Multi-view reconstruction preserving weakly-supported surfaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3121–3128, 2011. 2

[44] Anushrut Jignasu, Ethan Herron, Talukder Zaki Jubery, James Afful, Aditya Balu, Baskar Ganapathysubramanian, Soumik Sarkar, and Adarsh Krishnamurthy. Plant geometry reconstruction from field data using neural radiance fields. In *2nd AAAI Workshop on AI for Agriculture and Food Systems*, 2023. 5

[45] Anna Jobin, Marcello Ienca, and Effy Vayena. The global landscape of ai ethics guidelines. *Nature machine intelligence*, 1(9):389–399, 2019. 1, 2, 5, 6, 7

[46] Neil Joshi and Phil Burlina. Ai fairness via domain adaptation. *arXiv preprint arXiv:2104.01109*, 2021. 7

[47] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. 2

[48] Mijeong Kim, Seonguk Seo, and Bohyung Han. Infonerf: Ray entropy minimization for few-shot neural volume rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12912–12921, 2022. 2, 4

[49] Abiramy Kuganesan, Shih-yang Su, James J Little, and Helge Rhodin. Unerf: Time and memory conscious u-shaped network for training neural radiance fields. *arXiv preprint arXiv:2206.11952*, 2022. 6

[50] Jonáš Kulhánek, Erik Derner, Torsten Sattler, and Robert Babuška. Viewformer: Nerf-free neural rendering from few images using transformers. In *European Conference on Computer Vision*, pages 198–216. Springer, 2022. 8

[51] Junseo Lee, Kwanseok Choi, Jungi Lee, Seokwon Lee, Joonho Whangbo, and Jaewoong Sim. Neurex: A case for neural rendering acceleration. In *Proceedings of the 50th Annual International Symposium on Computer Architecture*, pages 1–13, 2023. 5

[52] Xiuxiu Li, Yongchen Guo, Haiyan Jin, and Jiangbin Zheng. Neural surface reconstruction with saliency-guided sampling in multi-view. *IET Image Processing*, 17(12):3411–3422, 2023. 5

[53] Yuan Li, Zhi-Hao Lin, David Forsyth, Jia-Bin Huang, and Shenlong Wang. Climatenerf: Extreme weather synthesis in neural radiance field. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3227–3238, 2023. 5

[54] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9298–9309, 2023. 7

[55] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. *arXiv preprint arXiv:2310.15008*, 2023. 2, 7

[56] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, page 163–169, New York, NY, USA, 1987. Association for Computing Machinery. 2

[57] Ziyuan Luo, Qing Guo, Ka Chun Cheung, Simon See, and Renjie Wan. Copyrnerf: Protecting the copyright of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22401–22411, 2023. 5

[58] Tambiama Madiega. Artificial intelligence act. *European Parliament: European Parliamentary Research Service*, 2021. 1

[59] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021. 4, 6

[60] G Mazzacca, A Karami, S Rigon, EM Farella, P Trybala, F Remondino, et al. Nerf for heritage 3d reconstruction. *INTERNATIONAL ARCHIVES OF THE PHOTOGRAMMETRY, REMOTE SENSING AND SPATIAL INFORMATION SCIENCES*, 48(M-2-2023):1051–1058, 2023. 1, 2, 4

[61] Aldo R Mejia-Rodriguez, Elisa Scalco, Daniele Tresoldi, Anna M Bianchi, Edgar R Arce-Santana, Martin O Mendez, and Giovanna Rizzo. Mesh-based approach for the 3d analysis of anatomical structures of interest in radiotherapy. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 6555–6558. IEEE, 2012. 7

[62] Xiaoxu Meng, Weikai Chen, and Bo Yang. Neat: Learning neural implicit surfaces with arbitrary topologies from multi-view images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 248–258, 2023. 2

[63] Alessandra Micozzi and Valentina Giannini. How does the localization of innovative start-ups near the universities influence their performance? *L'industria*, 44(1):129–152, 2023. 1

[64] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2

[65] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41 (4):102:1–102:15, 2022. 6

[66] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5480–5490, 2022. 6, 7

[67] Stavros Nousias, Gerasimos Arvanitis, Aris Lalos, and Konstantinos Moustakas. Deep saliency mapping for 3d meshes and applications. *ACM Transactions on Multimedia Computing, Communications and Applications*, 19(2): 1–22, 2023. 5

[68] International Council of Museums (ICOM). ICOM code of ethics for museums. https://icom.museum/wp-content/uploads/2018/07/ICOM-code-En-web.pdf, 2018. 1, 3, 5

[69] Jin-Si R Over, Andrew C Ritchie, Christine J Kranenburg, Jenna A Brown, Daniel D Buscombe, Tom Noble, Christopher R Sherwood, Jonathan A Warrick, and Phillipe A Wernette. Processing coastal imagery with agisoft metashape professional edition, version 1.6—structure from motion workflow documentation. Technical report, US Geological Survey, 2021. 6

[70] Sofia Pansoni, Simona Tiribelli, Marina Paolanti, Francesco Di Stefano, Emanuele Frontoni, Eva Savina Malinverni, Benedetta Giovanola, et al. Artificial intelligence and cultural heritage: Design and assessment of an ethical framework. *INTERNATIONAL ARCHIVES OF THE PHOTOGRAMMETRY, REMOTE SENSING AND SPATIAL INFORMATION SCIENCES*, 48:1149–1155, 2023. 1, 2, 3, 4, 5, 6, 7

[71] Sofia Pansoni, Simona Tiribelli, Marina Paolanti, Emanuele Frontoni, and Benedetta Giovanola. Design of an ethical framework for artificial intelligence in cultural heritage. In *2023 IEEE International Symposium on Ethics in Engineering, Science, and Technology (ETHICS)*, pages 1–5. IEEE, 2023. 2, 3

[72] Ermanno Petrocchi, Simona Tiribelli, Marina Paolanti, Benedetta Giovanola, Emanuele Frontoni, and Roberto Pierdicca. Geomethics: Ethical considerations about using artificial intelligence in geomatics. In *International Conference on Image Analysis and Processing*, pages 282–293. Springer, 2023. 6

[73] Anna Maria Piskopani, Alan Chamberlain, and Carolyn Ten Holter. Responsible ai and the arts: The ethical and legal implications of ai in the arts and creative industries. In *Proceedings of the First International Symposium on Trustworthy Autonomous Systems*, pages 1–5, 2023. 2, 3, 5

[74] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 2

[75] B Poole, A Jain, JT Barron, B Mildenhall, P Hedman, R Martin-Brualla, P Srinivasan, A Tewari, J Thies, et al. Rendering a better future. *A Computational Inflectio n for Scientific Discovery*, 66(8):15, 2023. 5, 6

[76] Fabio Remondino, Sabry El-Hakim, Armin Gruen, and Li Zhang. Turning images into 3d models-development and performance analysis of image matching for detailed surface reconstruction of heritage objects. *IEEE Signal Processing Magazine*, 25(4):55–65, 2008. 2

[77] Wojciech Samek, Thomas Wiegand, and Klaus-Robert Müller. Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. *arXiv preprint arXiv:1708.08296*, 2017. 5

[78] Klaus P Schneider. Methods and ethics of data collection. *Methods in pragmatics*, pages 37–93, 2018. 4

[79] Johannes L. Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *Computer Vision – ECCV 2016*, pages 501–518, Cham, 2016. Springer International Publishing. 2

[80] Mohamad Shahbazi, Evangelos Ntavelis, Alessio Tonioni, Edo Collins, Danda Pani Paudel, Martin Danelljan, and Luc Van Gool. Nerf-gan distillation for efficient 3d-aware generation with convolutions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2888–2898, 2023. 7

[81] Ramya Srinivasan and Kanji Uchino. Biases in generative art: A causal look from the lens of art history. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 41–51, 2021. 2, 3, 5

[82] Lorenzo Stacchio, Alessia Angeli, and Gustavo Marfia. Empowering digital twins with extended reality collaborations. *Virtual Reality & Intelligent Hardware*, 4(6):487–505, 2022. 5

[83] L. Stacchio, C. Scorolli, and G. Marfia. Evaluating human aesthetic and emotional aspects of 3d-generated content through extended reality. In *CEUR WORKSHOP PROCEEDINGS*. CEUR-WS, 2023. 5

[84] C. Strecha, R. Fransens, and L. Van Gool. Combined depth and outlier estimation in multi-view stereo. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, pages 2394–2401, 2006. 2

[85] Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, Abhik Ahuja, et al. Nerfstudio: A modular framework for neural radiance field development. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–12, 2023. 2

[86] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, et al. State of the art on neural rendering. In *Computer Graphics Forum*, pages 701–727. Wiley Online Library, 2020. 2

[87] Marco Toschi, Riccardo De Matteo, Riccardo Spezialetti, Daniele De Gregorio, Luigi Di Stefano, and Samuele Salti. Relight my nerf: A dataset for novel view synthesis and relighting of real-world objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20762–20772, 2023. 4

[88] Fabio Tosi, Youmin Zhang, Ziren Gong, Erik Sandström, Stefano Mattoccia, Martin R Oswald, and Matteo Poggi. How nerfs and 3d gaussian splatting are reshaping slam: a survey. *arXiv preprint arXiv:2402.13255*, 2024. 2

[89] UNESCO. Readiness assessment methodology: A tool of the recommendation on the ethics of artificial intelligence, 2023. 3, 5

[90] UNESCO. UNESCO Recommendation on the Ethics of Artificial Intelligence, 2024. 1, 3, 5, 6, 7

[91] UNESCO. UNESCO Cultural and creative industries in the face of COVID-19: an economic impact outlook, 2024. 1, 3, 5

[92] Urheber.info. Call for safeguards around generative AI. `https://urheber.info/diskurs/call-for-safeguards-around-generative-ai`, 2024. 3, 5

[93] Kangkan Wang, Guofeng Zhang, Suxu Cong, and Jian Yang. Clothed human performance capture with a double-layer neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21098–21107, 2023. 2

[94] Yuanfang Wang, Yu Li, Haoyang Zhang, Jun Yu, and Kun Wang. Moth: A hardware accelerator for neural radiance field inference on fpga. In *2023 IEEE 31st Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pages 227–227. IEEE, 2023. 5

[95] Yunqian Wen, Bo Liu, Jingyi Cao, Rong Xie, and Li Song. Divide and conquer: a two-step method for high quality face de-identification with model explainability. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5148–5157, 2023. 7

[96] Baao Xie, Bohan Li, Zequn Zhang, Junting Dong, Xin Jin, Jingyu Yang, and Wenjun Zeng. Navinerf: Nerf-based 3d representation disentanglement by latent semantic navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17992–18002, 2023. 4

[97] Feiyu Xu, Hans Uszkoreit, Yangzhou Du, Wei Fan, Dongyan Zhao, and Jun Zhu. Explainable ai: A brief survey on history, research areas, approaches and challenges. In *Natural Language Processing and Chinese Computing: 8th CCF International Conference, NLPCC 2019, Dunhuang, China, October 9–14, 2019, Proceedings, Part II 8*, pages 563–574. Springer, 2019. 5

[98] Meng Xu, Youchen Wang, Bin Xu, Jun Zhang, Jian Ren, Zhao Huang, Stefan Poslad, and Pengfei Xu. A critical analysis of image-based camera pose estimation techniques. *Neurocomputing*, 570:127125, 2024. 6

[99] Zishan Xu, Xiaofeng Zhang, Wei Chen, Jueting Liu, Tingting Xu, and Zehua Wang. Muraldiff: Diffusion for ancient murals restoration on large-scale pre-training. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2024. 7

[100] Haibo Yang, Yang Chen, Yingwei Pan, Ting Yao, Zhineng Chen, and Tao Mei. 3dstyle-diffusion: Pursuing fine-grained text-driven 3d stylization with 2d diffusion models. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 6860–6868, 2023. 2

[101] Jiawei Yang, Marco Pavone, and Yue Wang. Freenerf: Improving few-shot neural rendering with free frequency regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8254–8263, 2023. 2, 4, 6, 7

[102] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. 2, 4, 6

[103] Jason Y Zhang, Amy Lin, Moneish Kumar, Tzu-Hsuan Yang, Deva Ramanan, and Shubham Tulsiani. Cameras as rays: Pose estimation via ray diffusion. *arXiv preprint arXiv:2402.14817*, 2024. 6

[104] Zicheng Zhang, Wei Sun, Xiongkuo Min, Tao Wang, Wei Lu, Wenhan Zhu, and Guangtao Zhai. A no-reference visual quality metric for 3d color meshes. In *2021 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pages 1–6. IEEE, 2021. 7

[105] Jiali Zheng, Youngkyoon Jang, Athanasios Papaioannou, Christos Kampouris, Rolandos Alexandros Potamias, Foivos Paraperas Papantoniou, Efstathios Galanakis, Aleš Leonardis, and Stefanos Zafeiriou. Ilsh: The imperial light-stage head dataset for human head view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1112–1120, 2023. 4

[106] Zehao Zhu, Zhiwen Fan, Yifan Jiang, and Zhangyang Wang. Fsgs: Real-time few-shot view synthesis using gaussian splatting. *arXiv preprint arXiv:2312.00451*, 2023. 7