Advanced Engineering Informatics 62 (2024) 102848

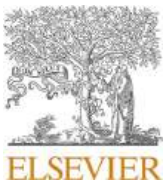

Contents lists available at ScienceDirect

Advanced Engineering Informatics

journal homepage: www.elsevier.com/locate/aei

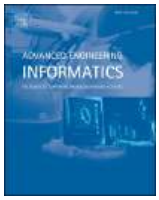

An automated CAD-to-XR framework based on generative AI and Shrinkwrap modelling for a User-Centred design approach

Riccardo Rosati [a,*], Paolo Senesi [b], Barbara Lonzi [c], Adriano Mancini [a], Marco Mandolini [b]

[a] Department of Information Engineering, Università Politecnica delle Marche, Via Brecce Bianche 12, Ancona 60131, Italy
[b] Department of Industrial Engineering and Mathematical Sciences, Università Politecnica delle Marche, Via Brecce Bianche 12, Ancona 60131, Italy
[c] Benelli Armi S.p.A., Via della Stazione, n. 50, Urbino 61029, Italy

ARTICLE INFO

Keywords:
CAD-to-XR
Extended Reality
Artificial Intelligence
Generative Adversarial Networks
Virtual Prototyping
User-Centered Design

ABSTRACT

CAD-to-XR is the workflow to generate interactive Photorealistic Virtual Prototypes (iPVPs) for Extended Reality (XR) apps from Computer-Aided Design (CAD) models. This process entails modelling, texturing, and XR programming. In the literature, no automatic CAD-to-XR frameworks simultaneously manage CAD simplification and texturing. There are no examples of their adoption for User-Centered Design (UCD). Moreover, such CAD-to-XR workflows do not seize the potentialities of generative algorithms to produce synthetic images (textures). The paper presents a framework for implementing the CAD-to-XR workflow. The solution consists of a module for texture generation based on Generative Adversarial Networks (GANs). The generated texture is then managed by another module (based on Shrinkwrap modelling) to develop the iPVP by simplifying the 3D model and UV mapping the generated texture. The geometric and material data is integrated into a graphic engine, which allows for programming an interactive experience with the iPVP in XR. The CAD-to-XR framework was validated on two components (rifle stock and forend) of a sporting rifle. The solution can automate the texturing process of different product versions in shorter times (compared to a manual procedure). After each product revision, it avoids tedious and manual activities required to generate a new iPVP. The image quality metrics highlight that images are generated in a "realistic" manner (the perceived quality of generated textures is highly comparable to real images). The quality of the iPVPs, generated through the proposed framework and visualised by users through a mixed reality head-mounted display, is equivalent to traditionally designed prototypes.

1. Introduction

In the contemporary landscape of Industry 4.0 (I4.0), it is becoming imperative for enterprises to focus on incorporating the human factor into the product design process. Accordingly, this shift towards prioritising human-centric approaches marks a significant transition from the Industry 4.0 to the Industry 5.0 paradigm [1]. This emphasis is crucial for addressing the user's physical, social, and cultural requirements, enhancing the final product's perceived value [2]. Several aspects must be considered when developing a novel product, including functionality and operability. The preliminary examination of these factors, conducted during the realisation of prototypes, can be facilitated through novel digital manufacturing tools and virtual technologies (e.g., virtual, augmented and mixed reality). These tools enable the objectification of customer requirements and their translation into technical specifications.

The paper presents a framework to help companies implement User-Centred Design (UCD). In this iterative design process, designers consider users and related needs during the design process. Thus, gathering information about users and their interactions with products is mandatory. Information can be collected through secure, comfortable, efficient, and highly esteemed physical or virtual prototypes that must be quickly generated. The interactive Photorealistic Virtual Prototype (iPVP) derived from the proposed methodology overcomes the drawbacks and limitations of traditional physical prototypes. The iPVP is developed from a product's Computer-Aided Design (CAD) model, adding aesthetic (created through generative Artificial Intelligence – AI – models) and functional parameters that can be modified and customised. Then, Extended Reality (XR) technologies allow users to interact with virtual objects and react with the real world [3]. XR will enable users to identify and define their preferred aesthetics for a new product and then make it available to the designer to translate them into design

* Corresponding author.
E-mail address: r.rosati@staff.univpm.it (R. Rosati).

https://doi.org/10.1016/j.aei.2024.102848
Received 24 January 2024; Received in revised form 31 August 2024; Accepted 23 September 2024
Available online 27 September 2024

specifications.

From the industry standpoint, the main challenges concern the high number of iPVPs to generate and made available for XR quickly. First, generating different material textures makes it possible to create many products and anticipate users' preferences. However, a system for texture generation is requested. In the context of AI, generative models can quickly make massive amounts of material textures according to the users' requirements. Second, iPVPs must be re-generated from the original CAD model of a product soon after a new texture is created. Nowadays, this is a manual, repetitive and labour-intensive process carried out through 3D modelling and rendering software tools. Generating an iPVP takes minutes to hours (depending on the complexity of the product). This procedure needs to be speeded up as much as possible by developing an automatic CAD-to-XR workflow to guarantee an affordable product development process.

From the scientific point of view, this is the first research in the literature that provides a method to create material textures for different wood classes. Moreover, based on shrinkwrap modelling, the study defines a CAD-to-XR workflow to texturing different 3D models automatically. The generated iPVPs should be a tradeoff between lightness for a real-time XR rendering and quality for good user interaction.

The paper aims to furnish a framework for implementing an automatic CAD-to-XR process. It transforms virtual models generated by 3D CAD systems into low-poly iPVPs suitable for real-time XR applications. First, material textures are generated and updated by a hyper-realistic texture model based on an evolution of Generative Adversarial Networks (GANs). This model can generate novel customised texture images integrated into the parametric virtual model. In this way, many preferences can be gathered from users before realising physical prototypes. Second, the texture is mapped on a low-poly simplified modelling cage generated around the 3D CAD model. This cage is then projected onto the high-poly CAD-imported model using a combination of Shrinkwrap, Boolean, and Subdivision Surface modifiers, resulting in a low-poly version of the identical external geometry. Through a texture baking process, high-frequency details such as engravings can be added by generating a normal map in which these details are imprinted. Blender (by Blender Foundation, Netherlands), the software tool used in this project, allows for swift, automatic updates to the virtual model in case of local changes to the component by creating a UV-mapped cage which encompasses it. The proposed method can adapt generated textures to variable geometries that are not known a priori in a flexible and fast manner.

After this introduction, the paper reviews the literature (section 2) on generative models for generating material textures and virtual prototyping for XR. Section 3 presents the proposed CAD-to-XR framework and its elements (polygon modelling, texture generation and integration with a graphic engine). The framework was used for two components of a sporting rifle (section 4) to present the texture generation, virtual modelling, and related results.

2. State of the art

This section explores the advancements in generative AI models, mainly focusing on their application in creating bitmap textures for physically based rendering (PBR) materials. This review analyses various generative approaches and their conditional variants, highlighting their capabilities and limitations in image synthesis and texture generation. Additionally, the section examines the evolving field of virtual prototyping, discusses innovative methods for CAD model conversion and texturing processes, and showcases how these technologies are integrated into modern manufacturing and design workflows. This comprehensive overview sets the stage for presenting the novel contributions of this study in the context of existing research and technological developments.

2.1. Generative AI models for creating material textures

Applying Deep Learning (DL) based generative models has shown promise in creating material textures that meet customer requirements. Generative Adversarial Networks (GANs) [4] are widely utilised and researched for their capabilities in image synthesis. The traditional GAN framework has two adversarial networks: a generator and a discriminator. The generator creates data from noise, and the discriminator distinguishes between generated and real data. They engage in a minimax game where the generator aims to fool the discriminator into thinking its data is real. In contrast, the discriminator tries to identify real and fake data correctly. Upon successful training, the generator can synthesise images from a noise vector as a starting point. The Conditional GAN (CGAN) [5] extends this concept by adding additional information, such as class labels or other auxiliary data, to the generator and discriminator. This added information makes the data generation process conditional and more directed.

Despite this advancement, GAN and CGAN suffer from issues like mode collapse and unstable training, partially due to their reliance on the Jensen-Shannon divergence, which limits them to generating continuous data. The introduction of the Wasserstein GAN (WGAN) [6] and its improved variant, WGAN with Gradient Penalty (WGAN-GP) [7], marks a significant advancement. WGAN uses the Earth-Mover distance instead of the Jensen-Shannon divergence. This distance metric is defined as the minimum cost of transforming one distribution into another and is more effective for comparing distributions. It addresses the problems of vanishing gradients and mode collapse in GANs. WGAN-GP further improves upon WGAN by addressing the issue of weight clipping. It introduces a gradient penalty term in the loss function, which helps stabilise the training across various GAN architectures.

As documented in the literature, GANs have several applications, including generating material textures. One notable example is TextureGAN [8], a pioneering DL method for image synthesis that enables users to manipulate object textures within a fashion context. Users can apply chosen textures to sketched objects, with the network realistically implementing these textures and implicitly recognising object boundaries for practical texture synthesis.

In the manufacturing domain, a research study is related to classifying road pavement textures [9]. This example involves generating new samples using a WGAN-GP network architecture and assessing image accuracy through various ML network architectures. The WGAN-GP network is notable for producing high-quality images. However, it is limited by the relatively small size of its samples (80x80 pixels). Similarly, Lopes et al. [10] explored the creation of synthetic microscopic cross-sections of hardwood species. Their algorithm, named StyleGAN, can generate high-resolution microscope images of cross-sections that are remarkably realistic and indistinguishable from actual images. Unique to StyleGAN is its progressive resolution enhancement, which is achieved by adding layers to the network, allowing it to handle high-dimensional space images (512x512 pixels).

2.2. Virtual prototyping for XR development

XR refers to immersive technologies that merge the physical and virtual worlds. By defining the concept of the virtuality continuum, Milgram and Kishino [11] established the full spectrum of technologies between the real world and the virtual environment. Virtual Reality (VR) is an environment in which the participant is immersed. Augmented Reality (AR) means augmenting the real environment with virtual (computer graphic) objects. Mixed Reality (MR) covers the "grey area" in the centre of the virtuality continuum.

In XR, the virtual prototypes used for product visualisation in manufacturing are typically polygon-based, textured models derived from CAD models. However, each alteration to the 3D model necessitates a complex and manual process of remodelling and retexturing, which is time-consuming. This traditional method becomes impractical

for customised products that require frequent modifications. As user inputs lead to changes in the CAD model, they also necessitate adjustments to the iPVP. Consequently, this approach demands manual reworking through all stages of virtual prototyping, making it inefficient and labour-intensive, especially in scenarios where frequent updates are needed. The three critical stages of virtual prototyping for XR apps are modelling, texturing, and XR development [12]. Using 3D CAD modelling software, the virtual prototype's 3D model is first created. Transforming a CAD model into an XR-ready format is a nuanced process that begins with converting the detailed CAD geometry into a more manageable polygonal mesh. This step is essential for adapting the model to the real-time rendering requirements of XR environments. Once converted, the model is textured and UV mapped, adding realistic surface details essential for immersive visualisation.

The textured model is then imported into a 3D rendering software or a game engine, such as Unity or Unreal Engine. It undergoes further enhancements with material properties and lighting adjustments to achieve a realistic appearance. This stage is critical for rendering the model visually appealing and accurate.

Finally, the model is integrated into an XR environment, becoming part of a more extensive interactive system, combining elements like user interfaces and spatial tracking. This integration also involves optimisation to balance the model's visual quality with the performance, ensuring smooth operation in XR applications. This entire process is critical to transforming static, high-detail CAD designs into dynamic, interactive models for the immersive world of extended reality.

Typically, to utilise a CAD model in an XR system, the original model is converted into a mesh and transferred to a rendering graphics system [13]. However, a simple CAD-to-mesh conversion often proves inefficient for XR systems due to the generation of excessive triangles [14]. Meshes created through tessellation result in higher triangle density, which significantly burdens performance, particularly in real-time graphic applications within XR systems. Tang and Gu [13] proposed a novel technique for CAD model translation and simplification for a Virtual Reality system in response to these considerations. Harlan et al. [14] recently created an XR-CAD platform to connect a game engine and CAD software effectively. Massive aircraft CAD models can be rendered using a GPU-based compression technique, according to Dunming et al. [15]. Model complexity reduction, animation, and kinematic adoption were used by Lorenz et al. to create an automated CAD-to-XR conversion method [16]. Despite the advantages, the methods mentioned above cannot handle textured models. Thus, they cannot be used in the context of UCD, where textures are an essential element of the product's aesthetics.

Prada et al. [17] tried to overcome this limitation by providing an intriguing method for transforming CAD models (taking textures into account) for real-time rendering. The authors overcame the issue of UV mapping intricate polygonal models exported in Standard Triangulation Language (STL) format by using the DATASMITH plugin of Unreal Engine. Unfortunately, the cited paper does not describe how to use the method to manage CAD model revisions quickly.

In virtual prototyping, CAD simplification and defeaturing can be realised through Shrinkwrap-based modelling. This feature creates a closed and simplified envelope around a more complex virtual model. The memory usage of the obtained model can be reduced by up to 90 % [18]. Several CAD systems offer this feature (e.g., *Creo Parametric* by PTC, *HyperMesh* by Altair Engineering and *Blender* by Blender Foundation). For example, the Shrinkwrap modifier in Blender [19] has been pivotal in diverse applications, demonstrating its versatility across different fields. It played a critical role in [20], which was employed for morphing between different geometric shapes, refining a model's scale based on a more intricate one.

Similarly, in [21], the Shrinkwrap modifier was adeptly used to project multiscaffold geometries onto a 3D-scanned print plane, distorting the geometry to conform to the print surface. Moreover, [22] utilised the Shrinkwrap modifier to project a flat array of elements onto a dome surface, ensuring the elements accurately adhered to the dome's curvature. The technique was also integral in [23], where it was applied to NURBS path functions, allowing catheter tunnels to follow the contours of a surface applicator closely. In this study, the Shrinkwrap modifier has been innovatively used to rapidly achieve a simplified, textured geometry from a CAD model, underscoring a novel application of this tool that enhances the efficiency and effectiveness of virtual prototyping.

In the state of the art of computer graphics, normal maps have emerged as an essential tool for rendering intricate surface details without the computational burden of high-polygon models. They have been pivotal in simulating the microscopic structure of surfaces, providing a way to add depth and complexity to textures. [24] and [25] showcase the application of normal maps to render the fine structural details of surfaces, which would otherwise require significantly more geometric detail. This method enhances the visual richness of the material while keeping the geometry lightweight and manageable. Meanwhile, [26] outlines a process that leverages detail normal maps generated through texture baking techniques. To the author's knowledge, no previous studies combine Shrinkwrap-based modelling and texture baking of detail normal maps.

2.3. Novelty of the study

The main novelty of this paper lies in two key areas. First, it introduces new methods for texture generation and automatic texturing of 3D models. Second, it integrates these results into a framework that supports the generation of iPVPs suitable for XR applications.

Concerning texture generation, to the best of our knowledge, this is the first study in the literature that presents an approach for generating synthetic images (textures) related to different (four) aesthetic quality classes of wood used to manufacture sporting rifle parts, which is the specific case study selected in this work. Using a GAN model in the random generation of material textures offers several significant manufacturing and computer graphics benefits. Firstly, GANs enable the generation of unique texture samples that extend beyond the limitations of the initially acquired dataset. This capacity for extrapolation allows for creating novel, previously unseen textures, enhancing the diversity and richness of available material representations. Secondly, GANs can be tailored to generate texture samples that align with specific user preferences or requirements, facilitating a more personalised and targeted approach to texture generation. This aspect is particularly beneficial in virtual reality and digital design applications, where user-specific customisation is paramount. Lastly, applying GANs in material texture generation contributes to data augmentation, effectively increasing the volume and variety of samples within a dataset. This augmentation is crucial for training more robust and effective DL models, as it helps mitigate overfitting issues and enhances the models' generalizability to a broader range of textures and scenarios.

The Shrinkwrap-based method is unique in the literature for automatically texturing different 3D CAD models for XR. It allows the automatic UV mapping of a texture on different variants of an original 3D model by avoiding tedious and manual activities every time. The approach, furthermore, includes CAD simplification (conversion from high to low poly), which is required when XR systems manage the iPVPs.

Integrating the above-mentioned specific methods results in a comprehensive CAD-to-XR framework, bringing substantial benefits to Engineering Informatics:

i. Combining distinct modules into a complementary system reduces the time and effort required for transitioning from design to interactive virtual prototypes, enhancing the efficiency and efficacy of product development cycles.
ii. Incorporating generative AI for material texture generation enriches the visual realism and customisation capabilities within XR environments, providing a more immersive and tailored user experience.

These advancements support more rapid prototyping and iterative design, aligning with the goals of Industry 5.0 by fostering innovation and responsiveness in engineering and manufacturing domains. Employing the proposed framework makes product development processes for new and customised products faster and more effective. The preliminary research of the product's functionalities, design, and operability is conducted using these digital and virtual tools to understand the user's subjective needs in advance and transform them into technical specifications. Users can customise and interact with the product during the design process before the physical prototype is realised.

3. Material and methods

3.1. Overview

The framework developed in the study aims to rapidly obtain photorealistic models encompassing geometric and material attributes by integrating distinct techniques. The process combines two independent workflows: one for geometry modelling and another for material characterisation, exploiting a generative AI model trained on a dataset of real images (Fig. 1). The geometry of the model is reproduced by projecting an external UV-mapped cage, denoted as *Shrinkwrap Cage* [19], onto the imported model. This solution facilitates the transference of external geometric attributes to a UV-mapped polygonal object to which a PBR material, developed in the graphics engine utilising albedo textures generated via GAN, can be applied. An albedo texture is an image texture without any shadows or highlights. The material characterisation of the object requires training the GAN model and subsequent testing with validation metrics. Then, the GAN can generate unlimited images of albedo textures used in the graphics engine to configure materials for various components. To better exemplify the framework to simplify comprehension, the figures in the chapter will be related to the specific case study (i.e., the wooden stock of the rifle, described in §4.1).

3.2. Polygon modelling workflow

The fast geometry modelling method was developed in Blender by leveraging modifiers [27], which are computational operations that dynamically alter an object's geometry or other attributes. Modifiers allow for non-destructive changes, providing a more flexible and efficient workflow in the 3D design process. To implement this method, it is necessary to obtain the Shrinkwrap Cage for the specific component that requires to be modelled, as well as other 3D cutting tools, which will be discussed further in the subsequent chapter. The seams for UV mapping must be marked on the cage to unwrap the mesh onto the UV texture plane, allowing bitmap textures to be mapped onto its surface. The manual modeling of the cage and the cutting tools, and the UV mapping of the elements are performed only once for each component. After this initial step, they can be integrated into an automatic cycle to obtain a 3D low-poly textured model from CAD geometry quickly. Special attention must be given to how modifiers affect the Shrinkwrap Cage during the modelling phase to ensure their correct behaviour on the mesh.

The principal actor in this method is the Shrinkwrap modifier, which allows a polygonal geometry to deform so that it adheres to the external surface of a specified 3D object (Fig. 2). When applied to the Shrinkwrap Cage, the mesh can deform and conform to the external geometry of the

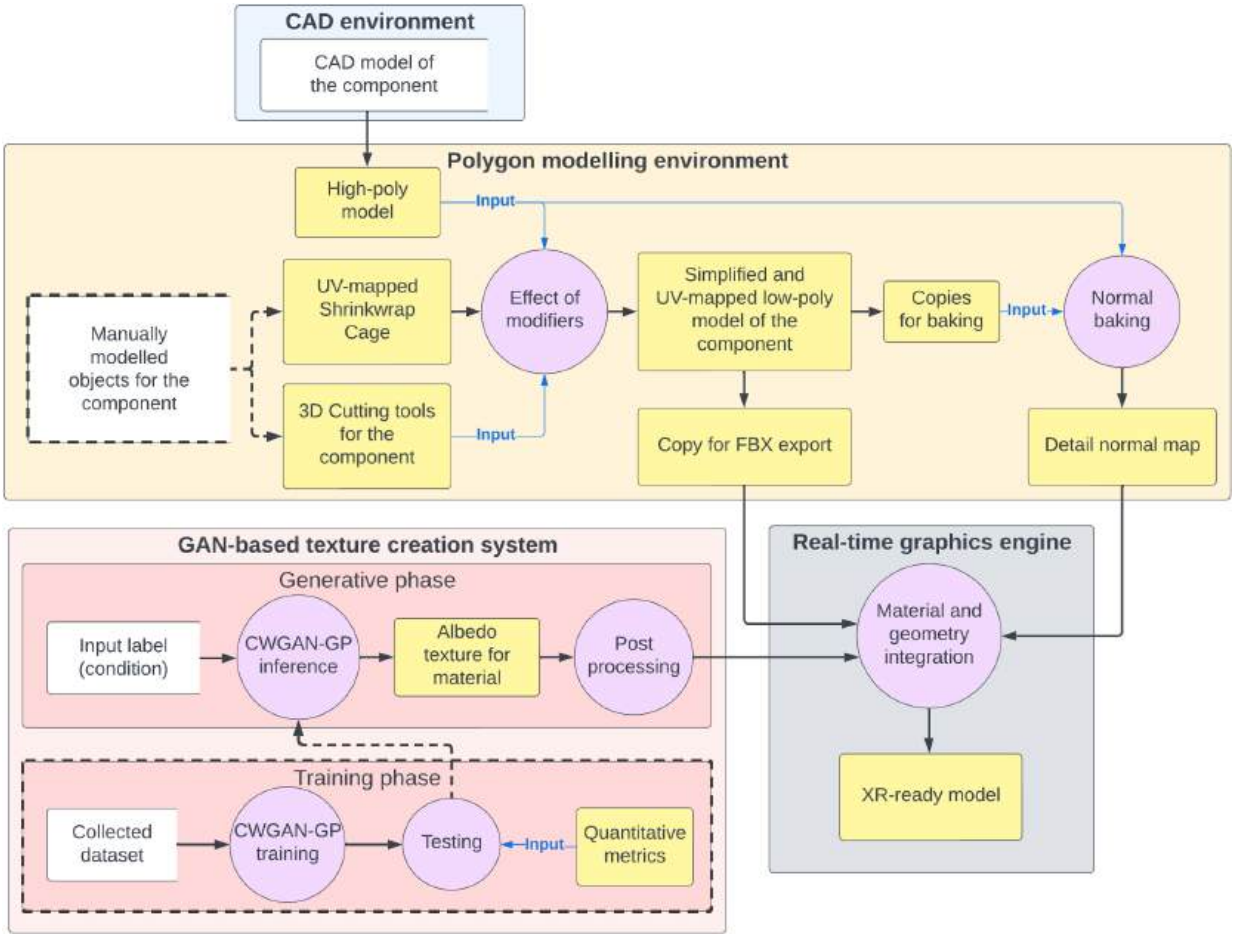

Fig. 1. CAD-to-XR Data Flow Diagram in Yourdon and Coad DFD notation. Solid lines represent a continuous update cycle flow, while dashed-lined flows and elements are one-time setup operations.

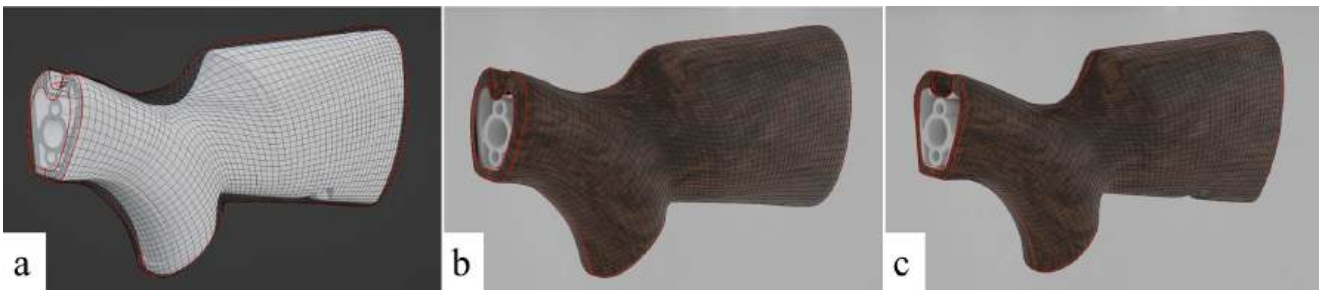

Fig. 2. Effect of the Shrinkwrap modifier acting on the UV-mapped cage, with the high-poly model selected as the target object. (a) shows the geometry and relative position of high poly, displayed as a solid white object, and the textured cage which encloses it, displayed in wireframe view. (b) shows the same scene, but the cage is displayed with a shaded rendering, illustrating the cage's texture. (c) shows how the external cage deforms, thanks to the Shrinkwrap modifier, obtaining the UV-mapped low-poly geometry.

high-poly model imported from CAD software.

3.2.1. Integration with auxiliary modifiers

Fig. 2 shows that the Shrinkwrap modifier is adequate for capturing the global, low-frequency geometric attributes of the imported 3D model. However, it fails in shapes such as sharp edges, holes, or grooves, as shown in Fig. 3 (b). To remedy this, a combination of other modifiers is necessary. Specifically:

- *Subdivision surface modifier*: this improves the mesh resolution, ensuring that the Shrinkwrap process is more precise.
- *Boolean modifier*: this helps capture intermediate-level features by removing specific geometric shapes from the model.
- *Solidify modifier*: this modifier is needed when the wrapped mesh does not constitute a closed volume but has openings, as in the case of Fig. 2. In this instance, it ensures that the exported polygonal object has a thickness.

The Boolean modifier is used by cutting the Shrinkwrap Cage with 3D objects that locally refine the texture, enabling it to wrap around the mid-frequency features properly. As is done for the Shrinkwrap Cage, these cutting tools need to be created through manual modelling, paying attention to the impact of their subtracted geometry on the behaviour of the Shrinkwrap modifier acting on the cage. These cutting tools need to replicate the negative form of the feature that needs to be wrapped by the cage to increase the local resolution of the mesh so that it can wrap around the feature, avoiding geometry defects. This modifier ensures that the Shrinkwrap modifier can more accurately conform to complex, high-poly structures, even in areas with medium-frequency details, as shown in Fig. 3 (c).

The output of this process is a lightweight low-poly model devoid of the high-frequency details of the original 3D model. This low-poly model will be used with a normal map exported as a 16-bit raster image. This normal map must be applied to the material to simulate the high-frequency details on the model's external surface. The normal map is generated through a process known as texture baking.

3.2.2. Texture baking process for detail normal map

Texture baking is a process in 3D graphics that involves pre-computing and transferring various surface details—such as lighting, shadows, and color information—from a high-fidelity model or scene onto a texture map. This technique significantly enhances rendering efficiency in real-time applications by reducing the need for dynamic calculations during runtime. In [28], various aspects related to texture baking are thoroughly explored, highlighting how Blender employs the raycasting technique for this task. In Blender's baking approach, ray-casting involves casting rays from the surface faces of a low-resolution model along their normal direction, intersecting and capturing attributes from a high-resolution model. One of the most common uses of this baking process is the reproduction of bump maps, also known as normal maps, where the texture details, such as bumps and grooves, are simulated on a lower-resolution model to give the illusion of a more complex surface without the computational load of a high-poly model. In Blender, the texture baking of the normal map involves three key elements:

- *High-poly model*: this is the source for the fine surface details captured in the normal map.
- *Low-poly model*: this is the base model for generating the normal map.
- *Baking cage*: this sets the volume boundaries for the baking operation (optional).

The baking cage is unnecessary, but its adoption results in computational and qualitative outcome optimisation. This tool must be created by duplicating the low-poly model and then upscaling it along the face normal direction. By shifting all the faces in the direction of their outer normal directions, a uniform gap between all the faces between the original and up-scaled mesh is obtained. This results in a larger cage encompassing both the high and low-poly models. The baking process uses the UV mapping of the low-poly model (as seen in Fig. 4) to produce a detail normal map that, when correctly combined with the object's material, allows for visualising the desired fine surface details.

After generating the geometry and the corresponding normal map,

Fig. 3. Elements relative to the employment of the Boolean modifier. Figure (a) shows the 3D-modelled cutting tool used for the sling hole on the rifle stock. Without this cutter, the Shrinkwap modifier produces artefacts (b). Figure (c) shows the Shrinkwrap modifier action corrected thanks to the cutting tool showed in (a).

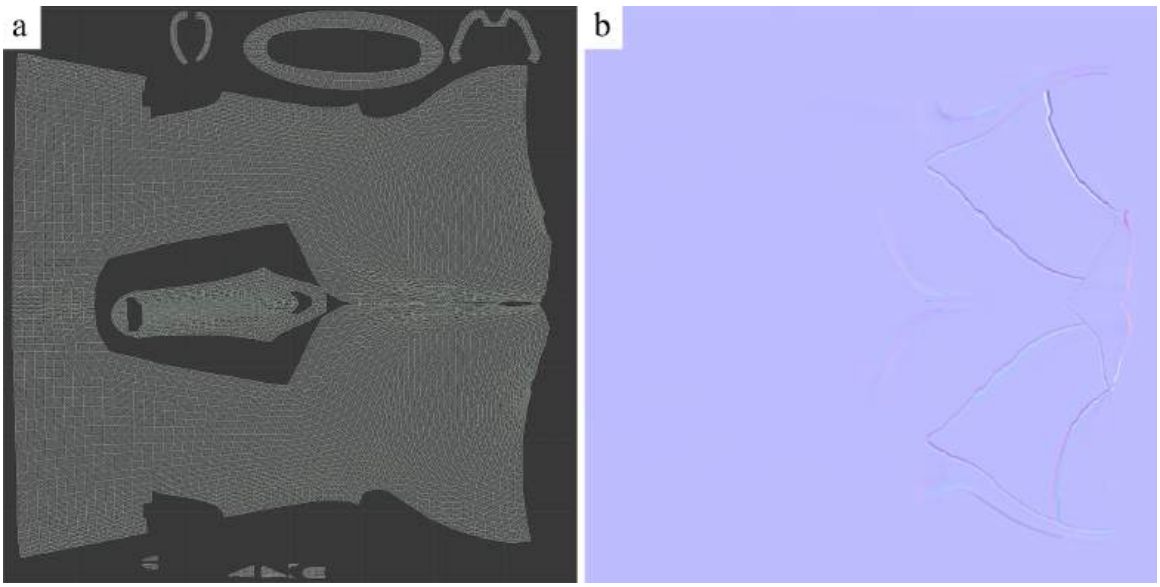

Fig. 4. UV mapping of the rifle stock low-poly exported model (a) and its relative detail normal map (b) obtained by the baking process.

these two data must be exported from the modelling software to be integrated into the graphics engine, with the textures related to the material for which the GAN image generation system is exploited.

3.3. Gan-based texture creation process

3.3.1. Conditional Wasserstein GAN model

The employed generative model is a Conditional Wasserstein GAN with Gradient Penalty (CWGAN-GP) [29]. This generative model is a variant of traditional GAN architecture that addresses stability and convergence issues in the conventional GAN training process. In our context, the CWGAN-GP is suitable because it combines the conditional aspect of CGANs with the stable training features of WGAN-GP, as described in §2.1. It conditions the generator and discriminator on auxiliary information, i.e. class labels. This integration enables the CWGAN-GP to generate high-quality data closely aligned with the specified conditions, allowing it to create specific required textures. The objective function of CWGAN-GP is a minimax equation that incorporates both the conditional information and the gradient penalty. The discriminator and generator's loss functions are structured to minimise their respective objectives while considering the conditional information and the gradient norm. In our case, the condition is related

Table 1
Dataset description. Aggregated classes were used for the generative task. It is worth noting the class imbalance, which represents an additional challenge in the CWGAN-GP training process.

Quality Classes	# Images	Aggregated Classes	# Images	Examples
1	165	1	165	
2^- 2 2^+	148 212 177	2	537	
3^- 3 3^+	179 306 344	3	829	
4^- 4 4^+	208 275 106	4	589	

to the quality class of the wood the user wants to implement, as illustrated below.

3.3.2. Dataset

The CWGAN-GP network was trained using a dataset comprising 2120 images, each measuring 470x270 pixels. These images were captured through a specially designed quality control bench with multiple cameras, a lighting system, and a handling/hooking system [30]. The description and preprocessing of the dataset are reported in [31]. This setup facilitated the quick collection of images focusing primarily on the wooden parts of rifles. The dataset used and the related use case have been widely explored in the literature [31–33]. It presents ten different wood aesthetic quality classes, ranging from class 1 to class 4^{+}. For simplicity in training the generative model and considering the small number of images available for each class, in this case, only four classes were considered, aggregating the minority classes, as shown in Table 1. Nevertheless, this rationale met the company's demands, primarily because the choice had no significant economic or productive impact, given the substantial similarity between the aggregated classes.

3.3.3. Experimental procedure

The training involved experimenting with various hyperparameters, including batch sizes, dimensions of the latent space, and the number of training epochs, along with beta1 values, beta2 values, and learning rates. The optimisation algorithms tested were Adam, SGD, and RMSprop. They were selected after determining the most effective hyperparameter configuration while monitoring the Fréchet Inception Distance (FID) and Inception Score (IS) metrics (detailed in §4.3). It was noted that excessively increasing the latent space size beyond 100 did not significantly improve these metrics, suggesting that further size increases did not enhance image quality. Ultimately, training was conducted with a validated batch size of 64 images. This decision was made because while testing larger batch sizes (i.e. 128 samples), we observed marginal improvements in FID (1.03 on average) and IS (0.07 on average). Still, a significant increase in training time offsets these due to the higher computational load (42 % longer training time). The best hyperparameter settings used for the training and the explored ranges are reported in Table 2.

3.3.4. Post-processing stage

Following the generation of the wood colour texture, it must be correctly mapped onto the low-poly geometry derived via the Shrink-wrap technique. To ensure the virtual model exhibits a high fidelity to realism, the resolution of the textures must be judiciously fine-tuned following the physical dimensions of the model. This calibration is critical to prevent the appearance of pixelation or granularity when the texture is rendered on the model. The chosen texture resolution should correspond to the desired scale and level of detail on the model surface. For close-up visualisations or models with elaborate details, high-resolution textures are recommended. A pragmatic strategy involves selecting a texture resolution commensurate with the simulated material's intrinsic granularity, in this instance, the wood grain. This ensures that the texture faithfully reproduces the authentic visual effect, accurately reflecting the inherent material characteristics of various wood types. Notably, the pixel density must be carefully considered to optimise the 3D model visualisation with an XR Head-Mounted Display (HMD), such as the Varjo XR-3, which boasts a maximum resolution of 70 pixels per degree (PPD). At the minimum distance where the XR-3 headset preserves maximum focus, specifically at 30 cm, this density corresponds to 339.3 pixels per inch (PPI).

During the inference phase, the images generated by the trained CWGAN-GP model maintain the same size as the training images (470x270). However, this resolution is insufficient for seamless integration of the generated textures into virtual models, as the aesthetic quality would be compromised—corresponding to 72.9 PPI, equivalent to 15.0 PPD at a distance of 30 cm. Consequently, a Super-Resolution (SR) GAN model, specifically the Enhanced Super-Resolution Generative Adversarial Networks (ESRGAN) architecture [34], was employed in a post-processing stage to enhance the resolution of the generated images (i.e., 1880x1080) [35]. With SR, a resolution of 291.6 PPI is achieved, corresponding to just over 60 PPD, thus meeting the minimum requirement for achieving photorealistic resolution in VR, as per [36].

Given that the texture produced by the CWGAN-GP represents a finite region relative to the model's total surface area, a tiling process becomes essential even after applying the SR algorithm to maintain the correct scale of the wood surface pattern. This process involves the spatial repetition of a texture section across the entire model's surface. Achieving seamless tiling involves symmetrically mirroring a segment of the image to construct a motif that can be replicated without discernible discontinuities. This technique ensures the lifelike portrayal of the texture over extensive areas, preserving the material's uniformity and integrity. Incorporating variations such as stochastic offsets, rotation, or scaling of the tiled segments can further enhance surface heterogeneity and authenticity. This described procedure is fully automated within the framework.

3.4. Graphics engine integration

After generating the geometry and corresponding texture for the material, all files can be imported into the graphics engine used for XR development. Depending on the XR visualisation software, the integration process advances with material management and model positioning within the larger product assembly.

This step initiates importing the low-poly model into a graphic engine like Unity or Unreal Engine. The software automatically creates a new material that incorporates the previously generated maps. Applying the normal maps to the low-poly model enables the simulation of complex surface details without needing a high-density polygonal geometry, thus maintaining the model's efficiency. Concurrently, using the colour map ensures that the textures and colours of the component are displayed faithfully, in combination with other default textures such as normal and roughness textures. This integration in the material accurately represents the surface characteristics, contributing to a realistic and cohesive visual representation of the model.

At this stage, the iPVP is ready for use in the XR environment. Assuming the use of the developed framework to recreate the individual components of a product that make up the assembly, by repeating this framework procedure for each product element, it is possible to assemble the entire set, managing the configuration of each component independently. This procedure enables the creation of an XR product configurator, where the materials of individual components can be modified. In doing so, the user can view the product in various variants, experimenting with different combinations and configurations in an interactive and immersive manner.

Table 2
Hyperparameters for training CWGAN-GP generative model.

Hyps	Optimizer	Batch size	Latent space	Epochs	Beta1	Beta2	Learning rate
Range	Adam, SGD, RMSprop	{32, 64, 128}	{50, 80, 100, 200, 500}	{200, 500, 1000, 2000, 2500, 3500, 4000}	{0.4, 0.5}	{0.9, 0.999}	$\{10^{-3}, 10^{-4}, 10^{-5}\}$
Best	Adam	64	100	3500	0.5	0.9	10^{-4}

4. Validation, results and discussions

4.1. Case study

The iPVP generation framework was employed in a case study in collaboration with Benelli Armi Spa, a manufacturer of firearms for hunting and sporting use interested in exploring the new frontiers of Industry 4.0 technologies. The company established a tracking area dedicated to immersive XR to visualise their products, aimed at rapid prototyping and collaborative design. An MR application has been implemented through the Unreal Engine environment (by Epic Games, North Carolina – US). The app permits the user to hold a mono-material blank 3D-printed prototype and view through the Varjo XR-3 (by Varjo, Finland) HMD, a photorealistic virtual prototype projected over the physical one. The headset gives the possibility to experience the app from aesthetic and haptic standpoints. One of their products, the Lupo rifle [37], was the subject of the study, specifically focusing on the stock and forend components, which are often made of wood or highly customised materials (e.g., camouflage).

The case study involved simulating the entire CAD-to-XR workflow for both components, assessing the quality of the final virtual prototypes and comparing the satisfaction index between GAN-generated textures and the original ones used for training the GAN. A total of 30 individuals tested the XR experience.

The case study validated the CAD-to-XR methodology from a qualitative standpoint. An indicator was established to evaluate the perceived quality of iPVP in the XR experience, considering the geometry and material.

4.2. Geometric modelling workflow

A prior study [34] validated the approach for rapid geometry acquisition using Shrinkwrap-based modelling. This method, implemented in a case study, involved simulating a geometric update of a component, with slight alterations in external features in terms of shape and size. The study illustrated that once the modelling cage for the element was established, the methodology facilitated the quick creation of updated virtual prototype versions, each fully textured using the Shrinkwrap technique, reducing the time for obtaining the textured model with updated geometry by more than half (around −50 % compared to an entirely manual process).

The method for obtaining the updated geometry has proven effective for the components selected in this study. The automatic Shrinkwrap-based modelling method has made it possible to rapidly obtain models of stock and forend with some qualitative imperfections. The quality of the process is significantly affected by some critical issues:

- Quality of the Shrinkwrap effect at the site of elements subjected to Boolean operations.
- The presence of thin and deep features (e.g., finely engraved characters), which are not adequately captured in the normal process.

However, as explained in §4.4, these defects were not noticed by the users, who rated the geometric quality of the models as adequate.

Fig. 5 shows microscopic defects occasionally arising from using the combined Shrinkwrap and Boolean modifiers. Fig. 6 also illustrates the suboptimal aesthetic quality of fine engravings on the forend component. In both instances, these details, which can be pretty prominent when viewing the model on screen and zooming into the specific area in the 3D modelling software, become less noticeable in XR due to the viewer's inability to focus closely.

4.3. Generative texture creation

Diverse qualitative and quantitative metrics have been implemented to test the CWGAN-GP architecture's generative effectiveness and obtain a reference value for comparison. The latter is useful for benchmarking and providing a basis for comparing future improvements or even new architectures. Specifically, the approach involved comparing two different random sets of real images to obtain a reference value for comparison with the results of generated images. Subsequently, the experiment was repeated by comparing a set of generated images with a set of real images. This approach allows the authors to assess how much the generated images deviate from real ones regarding quality and diversity. For each set of generated images, the following metrics were tested:

- **Image quality metrics**

a. Fréchet Inception Distance (FID): this is a performance metric for evaluating the similarity (in terms of quality) between two sets of images. The literature has demonstrated that FID correlates well with human image quality evaluation and can detect model collapse within the class [38]. A lower FID value corresponds to higher quality.

b. Inception Score (IS): this is a previous version of the FID metric, which, differently, evaluates only the distribution of generated images. A higher IS value corresponds to higher quality.

- **Image diversity metrics**

a. Structural Similarity (SSIM): this index provides a similarity measure by comparing two images based on luminance, contrast, and structural similarity features. A lower value indicates higher diversity among the generated images.

b. Learned Perceptual Image Patch Similarity (LPIPS): essentially, it calculates the similarity between activations of two image patches

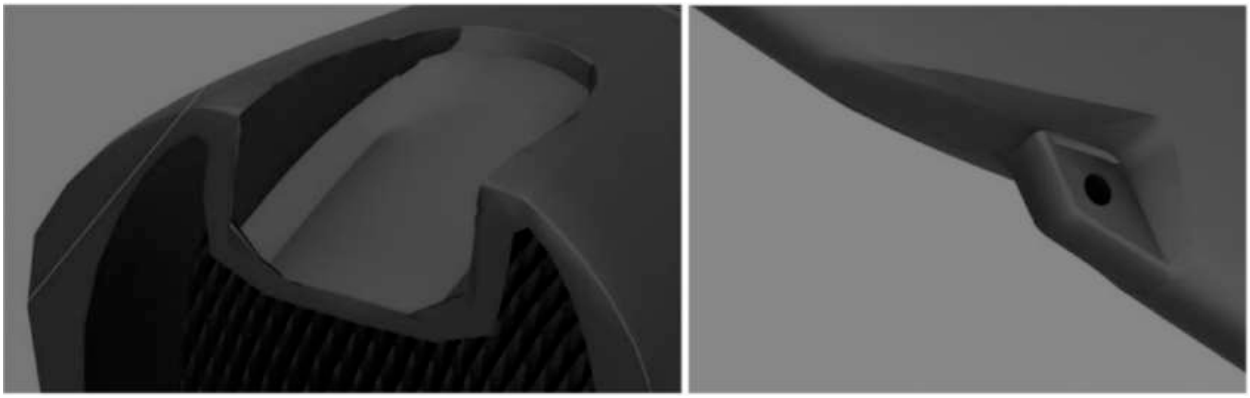

Fig. 5. Microscopic shading defects in Boolean and Shrinkwrap modifier combination effect.

Fig. 6. Aesthetic rendering of an engraved inscription reproduced through the detail normal map.

for some predefined networks. A low LPIPS score indicates that image patches are perceptually similar.

The results in Table 3 demonstrate that the metrics are very similar, highlighting that the images are generated in a "realistic" manner. Results in the LPIPS metric underline that the perceived quality is highly comparable to that of real images. In addition, when observing the column of the SSIM metric, obtaining such small values ensures the diversity of the generated images concerning the original ones. However, since SSIM is based on pixel differences, this can also be due to the potentially lower quality of the generated images. This observation is confirmed by the FID and IS metrics results, where the images demonstrate a slightly lower quality than real ones. Also, looking at the visual goodness of the generated images (Fig. 7), it is pretty noticeable that pixels are more discontinuous. As mentioned earlier, this naturally leads to greater diversity in metrics.

Continuing with the qualitative analysis of the generated images, it has been observed that the best performance comes from high-quality classes, as images exhibit more details (i.e., a richer and more varied pattern of wood grain), resulting in a better pixelated and less noisy generation. Indeed, GANs commonly perform better with this kind of image. This result is because detailed images provide more information for the model to learn from, which helps generate more realistic and diverse outputs. On the other hand, when dealing with less detailed data, GANs might struggle to learn sufficient variability, leading to outputs that are too similar. This similarity can reduce the generalisation efficiency of the model, as it becomes less effective in handling a wide range of different inputs. Essentially, the richness and variety in the training data play a critical role in these architecture's effectiveness and generalisation capabilities in image generation.

4.3.1. Ablation study

We conducted an ablation study comparing the performance of the CWGAN-GP model with several other GAN variants, including the traditional GAN, Conditional GAN (CGAN), and Wasserstein GAN (WGAN). The results from this ablation study demonstrated that the CWGAN-GP model outperformed the other models in generating high-quality textures with superior realism and diversity (see Table 4). Specifically:

- **Traditional GAN**: the traditional GAN model showed instability during training, resulting in higher FID and lower IS values, indicating lower image quality. The SSIM values were also higher, reflecting less diversity among the generated images.
- **Conditional GAN (CGAN)**: CGAN improved over the traditional GAN by incorporating conditional information, but it underperformed compared to CWGAN-GP. The FID and IS metrics were better than those of the conventional GAN. Yet, they fell short of the performance achieved by CWGAN-GP, especially in handling complex texture patterns.
- **Wasserstein GAN (WGAN)**: although WGAN was more stable than the traditional GAN, it still did not match the texture quality of CWGAN-GP. The FID values were closer to those of CWGAN-GP. Still, the IS and SSIM metrics indicated that WGAN struggled to capture intricate details, resulting in less convincing textures.

4.4. CAD-to-XR workflow

The framework for implementing the CAD-to-XR workflow has been validated through an MR application and the involvement of testers. The validation aims to verify, through an HMD, the quality of the iPVPs generated via the Shrinkwrap method using textures generated via GAN algorithms (Fig. 8).

The MR application enables users to choose the texture material for two parts of the virtual prototype (forend and stock). The choice made by the tester involves loading a virtual model ready for the viewer. For each configurable part, the tester can choose between two materials (woods belonging to the same category), one real and one generated via

Table 3
Quantitative results regarding the quality of images created by the CWGAN-GP model for all wood classes in terms of FID, IS, SSIM, and LPIPS by comparing two sets of original images (ref) and sets of original and generated images.

Class	FID (ref) ↓	FID ↓	IS (ref) ↑	IS ↑	SSIM (ref) ↓	SSIM ↓	LPIPS (ref) ↑	LPIPS ↑
1	121.647	139.467	2.177	1.689	0.228	0.143	0.539	0.535
2	109.327	125.737	2.112	1.795	0.212	0.121	0.552	0.601
3	100.923	132.654	2.304	1.759	0.220	0.118	0.680	0.692
4	109.540	115.856	1.932	1.796	0.188	0.122	0.654	0.692

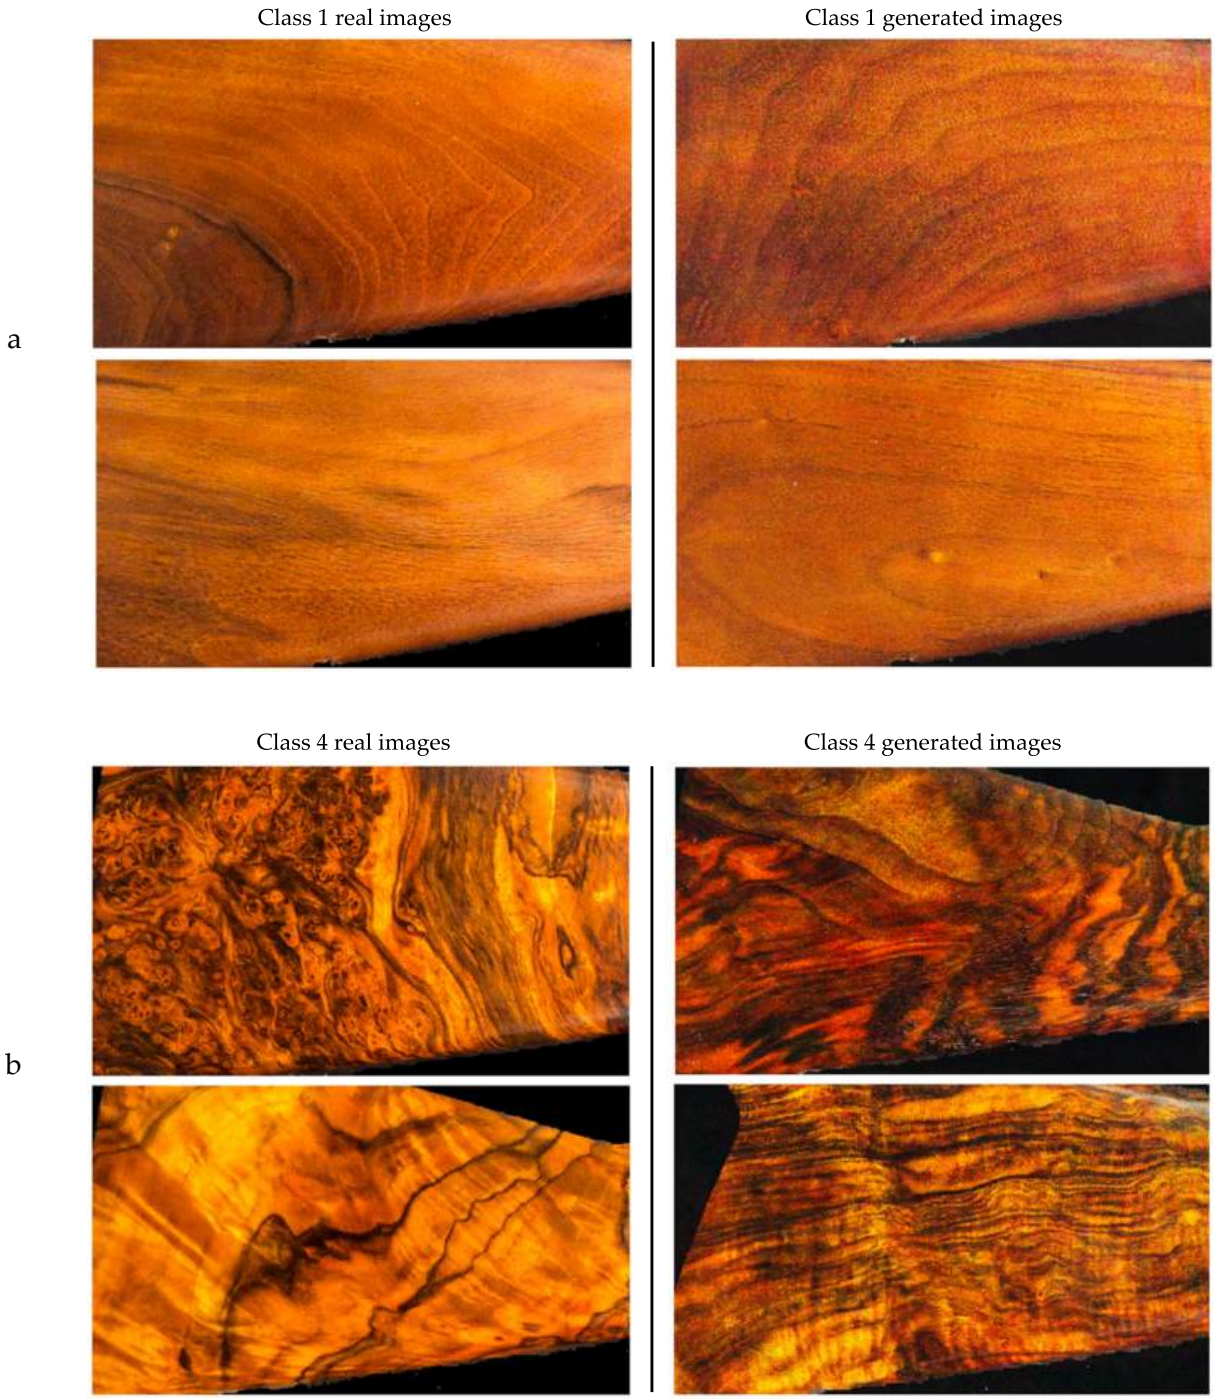

Fig. 7. Textures generated for two wood quality classes (1 and 4, the lowest and the highest), with a reference of real images for visual comparison. a) Examples of real dataset images (left) and generated images (right) for class 1; b) examples of real dataset images (left) and generated images (right) for class 4.

Table 4
Ablation study results (averaged on all classes). CWGAN-GP achieves the best overall performance, with the lowest FID and SSIM values and the highest IS and LPIPS scores. This results indicates that it produces more realistic and diverse textures than the other models tested.

Model	FID ↓	IS ↑	SSIM ↓	LPIPS ↑
CWGAN-GP	**128.428**	**1.759**	**0.126**	**0.630**
Traditional GAN	175.457	1.303	0.250	0.507
CGAN	158.322	1.466	0.195	0.556
WGAN	135.671	1.601	0.150	0.612

GAN. Each tester was offered different pairs of woods to cover the four available categories. The materials can be selected using a number (1 and 2) and are shown randomly. The natural wood corresponds to a virtual prototype generated with a traditional approach. Conversely, the GAN wood corresponds to a virtual model obtained via the Shrinkwrap method with a GAN-generated texture. The two choices allow testers to compare, within an MR viewer, iPVPs made traditionally with those obtained through the method proposed in the paper.

Each tester was asked, for each component, to select the two available materials and evaluate the quality of the iPVP. The evaluation took place by handling the 3D-printed blank prototype as it is commonly used. At the end of each session, the tester was instructed to indicate the

a b

Fig. 8. MR application developed for a sporting rifle. The application presents several rifle stock (a) and forend (b) materials. For these components, only real wood and GAN wood are selectable through numbers on the screen.

best material (i.e., iPVP) in terms of photorealistic representation. No maximum time was imposed for the choice.

The results were analysed to evaluate user preferences in choosing the models generated according to the traditional approach and the one proposed in the paper. The more significant the difference in the choice, the more distant the photorealism of the iPVP generated with the proposed method will be.

The validation involved 30 participants, namely employees and selected customers of Benelli Armi S.p.A. (Italy). The average age is 39.2 years. Testers are familiar with sporting rifles but do not have a background concerning MR devices. They are also familiar with virtual prototypes. Still, no one has used XR devices before this validation.

Table 5 presents the percentage of preferences for two different types of materials visualised by the tester for each model. In the pairwise comparison between the iPVPs (real vs GAN woods) for the stock and forend, 50 % is the target value. This value means that users cannot differentiate between the two iPVPs. Thus, the iPVPs generated using the proposed framework are equivalent to those created employing the traditional approach.

Results from Table 5 show that the iPVPs obtained following the traditional approach received more preferences from the testers than the proposed framework. However, the deviation in preference from the target of 50 % is 7 % for the rifle and 3 % for the forend. This outcome means that the quality of the iPVPs generated through the two methods is comparable overall. Thus, the benefits achievable with the proposed framework are not negatively influenced by the quality of the results.

Table 5
Results for the CAD-to-XR workflow validation.

	Real wood Traditional approach	GAN wood Proposed framework
Stock	57 % (17 testers)	43 % (13 testers)
Forend	53 % (16 testers)	47 % (14 testers)

5. Conclusions

This study proposed a framework for rapidly generating interactive photorealistic virtual prototypes used in XR applications (CAD-to-XR workflow). First, a Generative Adversarial Network, specifically a CWGAN-GP model, creates high-quality textures. Second, such textures are UV mapped onto a simplified 3D model through Shrinkwrap-based modelling. The latter, being larger than the model, can automatically UV mapping textures on 3D CAD models that are different from the original version. The resulting photorealistic virtual prototype is then transferred to an MR HMD for the user experience.

The framework was used to design a sporting rifle. The case study demonstrated that the image metrics for generated and real textures are very similar, highlighting that the images are generated in a "realistic" manner. The virtual prototyping process based on the Shrinkwrap approach is faster than the traditional one (around −50 %), where virtual prototypes are entirely manually generated (geometry and texture) after every revision of the original CAD model. Moreover, the photorealistic virtual prototype obtained through the proposed framework and visualised in the HMD looks like a prototype generated with a traditional approach (real texture and manual virtual prototyping process). Testers were not able to significantly appreciate differences.

The framework is still a methodological concept that deserves a more robust implementation into commercial software tools. For example, the virtual prototyping process still implies manually running Blender's commands (e.g., import of the geometries and textures, duplication, and application of Blender's modifiers). By employing Blender's application programming interface (API) it will be possible to automate the process. Additionally, it is expected to use generative algorithms to create normal and roughness material texture maps in the future. This solution will increase the quality of the iPVPs without losing the lightness of the models.

Concerning texture generation, the primary limitation includes the potential for overfitting due to the relatively small dataset and the manual nature of some processes within the framework, such as texture application. Additionally, while the CWGAN-GP model performed well, it has limitations in handling very detailed textures, as indicated by the slightly higher SSIM and FID values. Future research opportunities include expanding the dataset, automating more of the workflow, and exploring other generative models to enhance texture quality further.

CRediT authorship contribution statement

Riccardo Rosati: Writing – review & editing, Writing – original draft, Methodology, Data curation, Conceptualization. **Paolo Senesi:** Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Conceptualization. **Barbara Lonzi:** Resources, Project administration, Investigation, Funding acquisition. **Adriano Mancini:** Writing – review & editing, Supervision, Project administration, Funding acquisition. **Marco Mandolini:** Writing – review & editing, Writing – original draft, Validation, Supervision, Methodology, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The authors do not have permission to share data.

Acknowledgement

This work was supported by the 2014/2020 ERDF Regional Operational Programme of the Marche Region in the context of the research project 4USER – User and product development: from the virtual Experience to the Regeneration of the model (CUP: B36G20001210007).

References

[1] C. Zhang, Z. Wang, G. Zhou, F. Chang, D. Ma, Y. Jing, W. Cheng, K. Ding, D. Zhao, Towards new-generation human-centric smart manufacturing in Industry 5.0: a systematic review, Adv. Eng. Inf. 57 (2023) 102121, https://doi.org/10.1016/j.aei.2023.102121.

[2] D. Mourtzis, J. Angelopoulos, N. Panopoulos, A literature review of the challenges and opportunities of the transition from industry 4.0 to society 5.0, Energies 15 (2022) 6276, https://doi.org/10.3390/en15176276.

[3] L. Adriana Cárdenas-Robledo, Ó. Hernández-Uribe, C. Reta, J. Antonio Cantoral-Ceballos, Extended reality applications in industry 4.0. – a systematic literature review, Telematics Inform. 73 (2022) 101863, https://doi.org/10.1016/j.tele.2022.101863.

[4] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial networks, Commun. ACM 63 (2020) 139–144, https://doi.org/10.1145/3422622.

[5] M. Mirza, S. Osindero, Conditional generative adversarial nets, ArXiv. 10.48550/arXiv abs/1411.1 (2014) 1411.1784.

[6] M. Arjovsky, S. Chintala, L. Bottou, Wasserstein Generative Adversarial Network, In: Proceedings of the 34th International Conference on Machine Learnin, 2017: pp. 214–223.

[7] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, A. Courville, Improved training of wasserstein GANs, ArXiv. 1704.00028 (2017) https://github.com/igul222/improved_wgan_training. (accessed December 13, 2023).

[8] W. Xian, P. Sangkloy, V. Agrawal, A. Raj, J. Lu, C. Fang, F. Yu, J. Hays, TextureGAN: Controlling Deep Image Synthesis with Texture Patches, in: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, IEEE, 2018, pp. 8456–8465, https://doi.org/10.1109/CVPR.2018.00882.

[9] N. Chen, Z. Xu, Z. Liu, Y. Chen, Y. Miao, Q. Li, Y. Hou, L. Wang, Data Augmentation and intelligent recognition in pavement texture using a deep learning, IEEE Trans. Intell. Transp. Syst. 23 (2022) 25427–25436, https://doi.org/10.1109/TITS.2022.3140586.

[10] D.J.V. Lopes, G.F. Monti, G.W. Burgreen, J.C. Moulin, G. dos Santos Bobadilha, E. D. Entsminger, R.F. Oliveira, Creating high-resolution microscopic cross-section images of hardwood species using generative adversarial networks, Front. Plant Sci. 12 (2021), https://doi.org/10.3389/fpls.2021.760139.

[11] P. Milgram, F. Kishino, A taxonomy of mixed realityvisual displays, IEICE Trans. Inf. Syst. E77-D (1994) 1–15.

[12] S. Branislav, C. Dragan, Mixed reality and three-dimensional computer graphics, IntechOpen London (2020), https://doi.org/10.5772/intechopen.77405.

[13] Y. Tang, H. Gu, CAD Model's Simplification and Conversion for Virtual Reality, in: 2010 Third International Conference on Information and Computing, IEEE, 2010, pp. 265–268, https://doi.org/10.1109/ICIC.2010.338.

[14] J. Harlan, B. Schleich, S. Wartzack, Linking a game-engine with CAD-software to create a flexible platform for researching extended reality interfaces for the industrial design process, in: Proceedings of the 31st Symposium Design for X (DFX2020), The Design Society, 2020: pp. 169–178. 10.35199/dfx2020.18.

[15] T. Dunming, Z. Gang, Y. Lu, GPU Based Compression and Rendering of Massive Aircraft CAD Models, in: 2012 International Conference on Virtual Reality and Visualization, IEEE, 2012, pp. 50–55, https://doi.org/10.1109/ICVRV.2012.8.

[16] M. Lorenz, M. Spranger, T. Riedel, F. Pürzel, V. Wittstock, P. Klimant, CAD to VR – a methodology for the automated conversion of kinematic CAD models to virtual reality, Procedia CIRP. 41 (2016) 358–363, https://doi.org/10.1016/j.procir.2015.12.115.

[17] E. Prada, A. Kolár, Possibilities of convert cad models for real time rendering software, Technical Sci. Technol. 3 (2020) 220–228, https://doi.org/10.25140/2411-5363-2020-3(21)-220-228.

[18] CAD Interop, Create a ShrinkWrap from a complex model, (2024). https://www.cadinterop.com/en/your-needs/cad-weight-reduction/create-a-shrinkwrap-from-a-complex-model.html (accessed January 11, 2024).

[19] Blender, Shrinkwrap Modifier, (2024). https://docs.blender.org/manual/en/latest/modeling/modifiers/deform/shrinkwrap.html (accessed January 11, 2024).

[20] J. Berhouet, M. Slimane, M. Facomprez, M. Jiang, L. Favard, Views on a new surgical assistance method for implanting the glenoid component during total shoulder arthroplasty. Part 2: from three-dimensional reconstruction to augmented reality: Feasibility study, Orthop. Traumatol. Surg. Res. 105 (2019) 211–218, https://doi.org/10.1016/j.otsr.2018.08.021.

[21] J.-F. Chauvette, I. Lee Hia, R.D. Farahani, R. Plante, N. Piccirelli, D. Therriault, Non-planar multinozzle additive manufacturing of thermoset composite microscaffold networks, Compos. B Eng. 256 (2023) 1359–8368, https://doi.org/10.1016/j.compositesb.2023.110627.

[22] C.L. Valle, G.T. Carranza, R.C. Rumpf, Conformal frequency selective surfaces for arbitrary curvature, IEEE Trans. Antennas Propag. 71 (2023) 612–620, https://doi.org/10.1109/TAP.2022.3216960.

[23] C.A. Bridger, P.D. Reich, A.M. Caraça Santos, M.J.J. Douglass, A dosimetric comparison of CT- and photogrammetry- generated 3D printed HDR brachytherapy surface applicators, Phys. Eng. Sci. Med. 45 (2022) 125–134, https://doi.org/10.1007/s13246-021-01092-1.

[24] J. Zhu, Y. Xu, L. Wang, A stationary SVBRDF material modeling method based on discrete microsurface, Comput. Graphics Forum 38 (2019) 745–754, https://doi.org/10.1111/cgf.13876.

[25] H. Tan, J. Zhu, Y. Xu, X. Meng, L. Wang, L. Yan, Real-time microstructure rendering with MIP-Mapped Normal Map samples, Comput. Graphics Forum 41 (2022) 495–506, https://doi.org/10.1111/cgf.14448.

[26] K. Qian, Y. Li, K. Su, J. Zhang, A measure-driven method for normal mapping and normal map design of 3D models, Multimed. Tools Appl. 77 (2018) 31969–31989, https://doi.org/10.1007/s11042-018-6207-y.

[27] Blender, Modifiers, (2024). https://docs.blender.org/manual/en/latest/modeling/modifiers/index.html (accessed January 11, 2024).

[28] J. Knodt, Z. Pan, K. Wu, X. Gao, Joint UV optimization and texture baking, ACM Trans. Graph. 43 (2024) 1–20, https://doi.org/10.1145/3617683.

[29] M. Zheng, T. Li, R. Zhu, Y. Tang, M. Tang, L. Lin, Z. Ma, Conditional Wasserstein generative adversarial network-gradient penalty-based approach to alleviating imbalanced data classification, Inf. Sci. 512 (2020) 1009–1023, https://doi.org/10.1016/j.ins.2019.10.014.

[30] R. Rosati, L. Romeo, G. Cecchini, F. Tonetto, L. Perugini, L. Ruggeri, P. Viti, E. Frontoni, Bias from the Wild Industry 4.0: Are We Really Classifying the Quality or Shotgun Series?, In: Pattern Recognition. ICPR International Workshops and Challenges. ICPR 2021. Lecture Notes in Computer Science, 2021: pp. 637–649. 10.1007/978-3-030-68799-1_46.

[31] V.M. Vargas, P.A. Gutiérrez, R. Rosati, L. Romeo, E. Frontoni, C. Hervás-Martínez, Exponential loss regularisation for encouraging ordinal constraint to shotgun stocks quality assessment, Appl. Soft Comput. 138 (2023) 110191, https://doi.org/10.1016/j.asoc.2023.110191.

[32] R. Rosati, L. Romeo, V.M. Vargas, P.A. Gutiérrez, C. Hervás-Martínez, E. Frontoni, A novel deep ordinal classification approach for aesthetic quality control classification, Neural Comput. & Applic. 34 (2022) 11625–11639, https://doi.org/10.1007/s00521-022-07050-6.

[33] V.M. Vargas, P.A. Gutiérrez, R. Rosati, L. Romeo, E. Frontoni, C. Hervás-Martínez, Deep learning based hierarchical classifier for weapon stock aesthetic quality control assessment, Comput. Ind. 144 (2023) 103786, https://doi.org/10.1016/j.compind.2022.103786.

[34] X. Wang, K. Yu, S. Wu, J. Gu, Y. Liu, C. Dong, Y. Qiao, C.C. Loy, ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks, In: Computer Vision – ECCV 2018 Workshops. ECCV 2018. Lecture Notes in Computer Science, 2019: pp. 63–79. 10.1007/978-3-030-11021-5_5.

[35] Z. Wang, J. Chen, S.C.H. Hoi, Deep learning for image super-resolution: a survey, accessed December 13, 2023, IEEE Trans. Pattern Anal. Mach. Intell. 43 (2019) 3365–3387, http://arxiv.org/abs/1902.06068.

[36] J. Lu, T. Lian, J. Jia, Display and imaging system sharpness modeling and requirement in high-resolution VR and AR, Electronic Imaging. 35 (2023), https://doi.org/10.2352/EI.2023.35.12.ERVR-213, 213–1-213–6.

[37] Benelli Armi S.p.A., Benelli Lupo, (2024). https://benelli.it/en/lupo (accessed January 11, 2024).

[38] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, S. Hochreiter, GANs trained by a two time-scale update rule converge to a local Nash equilibrium, in: Advances in Neural Information Processing Systems, Long Beach, CA, USA, 2017: pp. 6627–6638. 10.18034/ajase.v8i1.9.