Original Article
ISSN (Online): 2582-7472

ShodhKosh: Journal of Visual and Performing Arts
Special Issue on Reimagining Visual Arts, Media, and Cultural Pedagogy in the Age of Artificial Intelligence 2025 6(1s)

DEEP LEARNING-BASED TEXTURE SYNTHESIS FOR SCULPTURAL REALISM

Dr. Shyamsing Thakur [1], Om Prakash [2], Dr. Kumud Saxena [3], Abhishek Singla [4], Dr. Sumalatha Potteti [5], Sonali Prashant Bhoite [6]

[1] Department of Mechanical Engineering, D Y Patil Coe Akurdi, Affiliated to Savitribai Phule Pune, Pune India
[2] Associate Professor, School of Business Management, Noida International University, Greater Noida, Uttar Pradesh, India
[3] Professor, Department of Computer Science and Engineering, Noida Institute of Engineering and Technology, Greater Noida, Uttar Pradesh, India
[4] Chitkara Centre for Research and Development, Chitkara University, Himachal Pradesh, Solan, India
[5] Bhoj Reddy Engineering College for Women, Hyderabad, Telangana, India
[6] Department of Information Technology, Vishwakarma Institute of Technology, Pune, Maharashtra, India

Received 26 January 2025
Accepted 21 April 2025
Published 10 December 2025

Corresponding Author
Dr. Shyamsing Thakur,
shyamsing2303@gmail.com

DOI
10.29121/shodhkosh.v6.i1s.2025.6631

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

ABSTRACT

In this paper, a multi-scale deep learning model that takes into account geometry to produce high-fidelity texture is proposed to condition on digital sculptures. The use of traditional 2D texture generation methods on complex 3D surfaces can be characterized by the presence of seam artifacts, distortions and disappearance of microstructural details, which makes them inappropriate in sculptural realism. The model is optimized on a multi-objective loss comprising of adversarial, perceptual and style components and patch-level and geometry-aware components and refined using differentiable rendering to match real lighting behaviour. It has been shown by experimental results of significantly better early results compared to classical techniques, StyleGAN-based baselines and diffusion-based generators, lower perceptual error (LPIPS), seam discontinuity (SCI, PBD) and multi-view consistency. Human perceptual research supports that the suggested approach prevails in almost 70% of the pair-wise tests with references to greater material authenticity and geometric consistency. The results illustrate the usefulness of viewing texture synthesis as a 2D-3D joint learning task and define the introduced system as a potential source of digital sculpting, heritage modeling, virtual production, and high-end assets all of which require material realist.

Keywords: Deep Learning, Texture Synthesis, Sculptural Realism, Geometry-Aware Generation, UV Mapping, Generative Adversarial Networks, Differentiable Rendering, Multi-Scale Feature Learning, Perceptual Loss

1. INTRODUCTION

The human eye comprehends even yet the clarity of the fine furrows of chisel lines, the porous sandstone grit, or the drag of polished marble, or the patina of old bronze, even before the sense of form or silhouette is apprehended. These textures characterize identity of the material, craftsmanship of signals and create realism that makes sculptures belong in the material space. The need to have realistic textures has risen so high with the development of digital sculpturing tools and 3D generation pipelines into film production, game design, restoration of cultural heritage and experiences in immersive VR. But sculptural realism is one of the most tedious and skills-intensive activities in computer graphics Cao et al. (2023). The classic work processes are textural painting by hand, high-resolution photogrammetry or procedural shaders, which do not necessarily have the expressive subtlety of actual materials. This means that digital sculptures often look artificial or unfinished with their surfaces not able to reflect some minor anomalies required to create a sense of believability to the eye. Deep learning has put a new frontier on texture synthesis by learning the grammar of surface patterns, through having examples Chan et al. (2022). Convolutional neural networks and generative models have shown an extraordinary performance in the multi-scale visual structure and, as a result, they can replicate the local chaos and general coherence nature of natural textures. These models are able to feel patterns in the sense in which a trained sculptor feels material; not in pixels, but as planes of structure which repeat, distort, intersect, and develop on a surface. In the application to sculptural realism, deep learning has the benefit of transforming texture synthesis by a crafty task into a data-driven task that can be driven by learned material priors Guerri et al. (2024).

The proposed research seeks to tackle these issues by creating a deep learning-based system to create high-fidelity and structurally sound sculptural-realism-specific textures. The geometricia combines convolutional encoders which extract features and generative adversarial networks which reconstruct texture 3D surfaces do not care about the topology of their components Yu et al. (2024). The system aims at developing textures that can be in harmony with complex sculpture shapes with multi-level perceptual characteristic and spatial consistency limitations whilst maintaining visual identity and microstructural detail. The suggested solution aims to enhance realism in addition to making the process of creating production-ready digital sculptures less time and expertise intensive Rajaei et al. (2024). The main values of the work are the design of a hybrid neural texture generation architecture designed specifically to match sculptural materials, geometry-sensitive loss functions to smooth texture transfer between geometries, and a controlled multi-material dataset of a variety of sculptural textures. Collectively, these technologies open up access to more automated, easier, and aesthetic workflows of production in digital sculpting and in 3D art. This work contributes to the development of digital sculptures that are touchable even on a screen by identifying the ways deep learning models can be used to create touchable textured materials by transferring material essence onto a textured surface Chen et al. (2022).

2. BACKGROUND AND RELATED WORK

Synthesizing texture has been one of the most significant areas of computer graphics research, determining how digital surfaces are able to obtain the small-scale irregularity that causes them to feel very real. The initial techniques viewed textures as statistical puzzles, i.e., the objective was to reproduce local sets of pixel that approximated hand picked reference patches. Such classical models such as Markov random fields down to patch-based quilting might be capable of replicating recent studies such as simple forms, but in many cases, the models failed to fold up when challenged to model the layered complexity of real sculptural materials Liu et al. (2023). Stone, clay, bronze, and wood have multi-scale structures, which are more like interwoven tales than is the case with repetitive tiles, and systems based on rules could easily fail to capture this richness, leading to artificial repetition or vague artifacts. This gave rise to procedural generation as an alternative, and increased computational power made it popular. Noise based techniques such as Perlin and Worley noise enabled artists to create textures using mathematical recipes, a mixture of deterministic rules with randomness Wang et al. (2021). Although these techniques were potent, they required genius of artistic feeling as well as skill, and were lengthy to adjust to reality. More to the point, procedurally generated textures usually did not have the micro-historical peculiarities of real materials: the scrape-hatch of a chisel, the haphazard swelling of old wood fibers, or the dust-worn crevices of aged stone. It is these flaws that make sculptures emotionally significant and procedural systems could seldom render a believable one. With the introduction of the deep learning, the tide turned. The neural style transfer of Gatys et al. provided the insight that convolutional layers store textures as twisted networks of feature-correlations and the revelation led to the unlocking of the first generation of neural texture synthesis Vaidya et al. (2025). CNN based techniques were able to reproduce textures on a fine scale with a subtlety that was not achievable with earlier methods but they were still largely bound to 2D grids. This was constrained by the fact that when projected on a 3D surface, it was possible to see the warping, stretching, and mismatch of seams, which affected the visual coherence. Nevertheless, their future was obvious: neural networks were learning the texture as a not a flat, but a structured arrangement of repeated motifs Chu et al. (2021).

Table 1

Table 1 Comparison of Key Deep Learning Models for Texture Synthesis

Model	Architecture Type	Texture Quality	Geometry Awareness
Gatys et al. (Neural Style Transfer)	CNN Feature Correlation	Moderate to High for 2D Images	None
StyleGAN / StyleGAN2 Wang et al. (2021).	Style-based GAN Generator	Very High with strong multi-scale detail	None
TextureGAN Szegedy et al. (2017)	Conditional GAN with Mask/Sketch Guidance	Good detail and structure control	Low
SPADE / SPADE-based Generators Bellini et al. (2016).	Spatially Adaptive Normalization Layers	High visual coherence	Low to Moderate
Neural Implicit Texture Models (e.g., NeRF-based) Frühstück et al. (2019)	Implicit Neural Fields with Learned Texture Representations	Very High with lighting-awareness	High
Differentiable Rendering-based Models Frühstück et al. (2019)	CNN/Transformer + Differentiable Renderer	Excellent micro-detail and shading consistency	High

Utilized in the form of generative adversarial networks (GANs), this promise was stretched to include full-fledged synthesis engines, which are capable of generating large seamless and highly detailed textures Shabeer et al. (2025). The introduction of models such as StyleGAN introduced multi-scale control in the generative process which allows networks to generate finer and coarse detail at the same time as denoted in Table 1. Spatially adaptive generators and texture GAN also showed that geometry conditioning or segmentation map conditioning could be used to direct texture generation. However, even the sophisticate models have struggled to apply to the case of sculptural realism, where geometric complexities, like curvature, ridges, occlusions, etc, are intrinsic to 3D surfaces that the traditional 2D GANs lack an intrinsic concept of. In cases where one texture encloses around the UV seams or across the line of extremely curved areas, the sense of reality may collapse unless the generator takes into consideration topology.

3. PROPOSED SYSTEM DESIGN MODEL

Sculptural realism requires a formulation which not only explains the appearance in pixel-space, but also the underlying topology of the surface and multi-frequency structure of real-world materials. The section provides the formulation of the fundamental equations that will control the encoder-generator-discriminator pipeline, geometrical aware conditioning, and multi-objective loss functions that jointly determine the learning process. The texture that must be synthesized is expressed as a function:

$$T : \Omega \rightarrow R3,$$

OR2 is used to represent UV texture space, and each value represents an RGB color. Given a sculptural surface S with UV parameterization ph S -O, the mapped texture is:

$$TS(p) = T(\phi(p)), p \in S.$$

This is to learn T in a way that the mapped surface TS will have material fidelity, structural coherence and geometry alignment as shown in Figure 1 The encoder E(O) picks up multi scale features of exemplar material image I: F=E(I). These characteristics possess local textural clues and general material identity.

$$T^\wedge = G(F, Ga)$$

Geometry-aware conditioning Ga includes distortion fields D, curvature maps κ, and surface normal N: Ga={D,κ,N}.

Figure 1

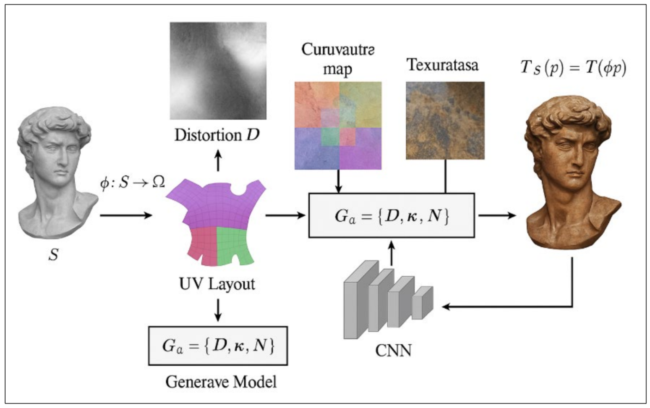

Figure 1 Geometry Aware Texture Mapping Design System

These signs enable the generator to come to know how visual patterns are supposed to stretch, compress or fit the 3-dimensional structure. Realism is based on adversarial learning:

$$LGAN = E[logD(T)] + E\left[\log\left(1 - D\left(T\quad\right)\right)\right].$$

In order to preserve perceptual structure and stylistic coherence a perceptual loss based on the VGG features is used:

$$Lperc = l\sum \parallel \Phi l(T) - \Phi l\left(T\quad\right) \parallel 22,$$

where Φl(⋅) denotes feature activations at layer l. Style similarity is enforced using Gram matrix correlations:

$$Gl(x) = \Phi l(x)\Phi l(x)\top, Lstyle = l\sum \parallel Gl(T) - Gl\left(T\quad\right) \parallel F2.$$

To maintain continuity across UV seams and preserve fine-scale structure, a patch-consistency loss is introduced. Let Pi, Pj be adjacent UV patches: $Lpatch = (i,j)\sum \parallel Pi - Pj \parallel 1$.

This term reduces abrupt transitions, blending variations without erasing natural irregularities. Geometry-aware consistency penalizes misalignment between generated patterns and sculptural topology. Using the distortion map D: $Lgeo = \parallel \nabla T\wedge \odot D \parallel 1$,where $\odot$ denotes elementwise multiplication. This encourages the generator to strengthen detail in regions prone to stretching or compression. Curvature-driven alignment is introduced as

$$Lcurv = \parallel T\wedge - \Psi(T\wedge, \kappa) \parallel 1,$$

where Ψ(⋅) aligns the texture orientation with principal curvature directions. To ensure correct surface-light interaction, differentiable rendering introduces a photometric loss. With renderer R(⋅) and lighting configuration L:

$$I^R = R(S, T^{,L})$$

$$Lrender = \parallel I^R - IR \parallel 1.$$

This term gives a response of the manner in which the generated texture reacts when illuminated, which strengthens physical plausibility. Each of the components becomes a part of a multi-objective function. This formulation turns texture synthesis into a joint optimization of material realism, geometric flexibility and rendering realism. Surface consciousness in the generation of visuals, the framework guarantees that the textures generated by synthesis have aesthetic and physical integrity when applied to sculptural surfaces of complex shapes.

4. PROPOSED HYBRID DEEP LEARNING FRAMEWORK DESIGN

The suggested methodology introduces a deep learning protocol on synthesizing high-fidelity, geometry-aware, sculptural realistic textures. In contrast to traditional 2D texture generators which ignore topological features of the surface, this method jointly learns the neural texture generation and 3D-aware refinement to provide structural consistency to applied textures on highly sculptural surfaces. The framework is composed of four essential parts, such as dataset preparation, hybrid model architecture, geometry-aware conditioning, and multi-objective training. The preparation of databases starts with the process of gathering high-resolution textures of marble, sandstone, granite, clay, wood, and aged metal with the help of macro photography and specific repositories. Preprocessing eliminates lighting variations and obtains intrinsic appearance data through the use of cropping, patch sampling, histogram normalization and seamless tiling augmentation. In the case of materials having directional grain or tool marks, natural anisotropy is preserved with orientation preserving transformations. Multi-scale features that describe fine irregularities and general style hints are extracted by the encoder and the generator optimizes them by progressive decoding. Skip connections maintain high-frequency structure and style modulation layers are wider in range of appearance variation. The discriminator is acting worldwide to achieve the material identity of the material and locally to achieve the fidelity of the micro-details to generate the balance between sharpness and structural coherence, as depicted in Figure 2. In order to address the weaknesses of synthesis that is all 2D, a geometry-aware integration module combines surface topology directly into synthesis. Examples of auxiliary conditioning channels used by the model are UV distortion fields, curvature maps, and surface normal, which allow the model to predict stretching, seam boundaries, and areas with more detail. This conditioning provides for easy transitions between UV islands, minimises the occurrence of artifacts in high distorting regions, and makes either the direction of the grain or the chisel pattern coincide with the natural contours of the sculpture, leading to a significant enhancement in perceptual realism.

Figure 2

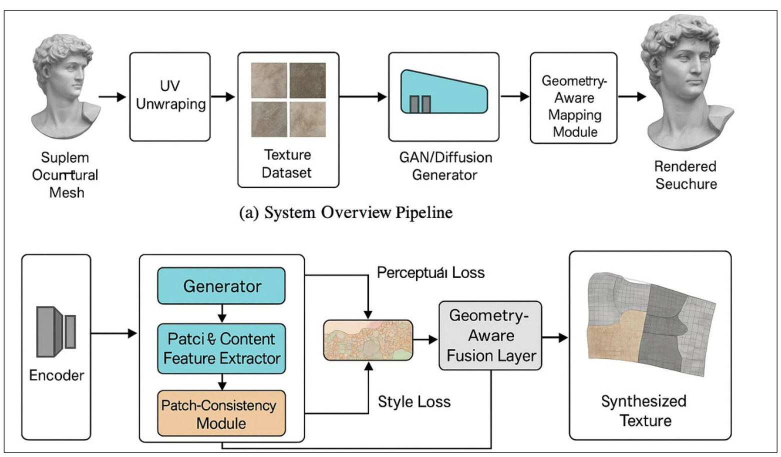

Figure 2 Architecture of the Proposed Hybrid Texture Generator

The total training of the framework is a multi-objective loss strategy that has to be well balanced. Conventional GAN loss gives adversarial supervision to refine details and make it look more realistic whereas the perceptual loss based on VGG features maintains structural fidelity and material identity. Style loss ensures the material coherence of the world by limiting statistics of feature correlations. These losses work together in a synergistic manner: the perceptual and style terms strengthen the holistic realism, patch and geometry terms guarantee continuity of the surfaces and the adversarial component enhances the sharpness of the visuals. This is a balanced goal which helps the generator to create textures that can be both believable in 2D and when applied to a 3D geometry. The last refinement phase uses differentiable rendering to induce feedback of lighting. The model is provided with gradient signals associated with shading, illumination and general visual coherence by rendering the generated texture directly onto the 3D sculpture. Overall, the methodology puts texture synthesis in the perspective of geometry-aware multi-stage learning. Through the combination of curated data set, hybrid generative architectures, topology sensitive conditioning and extensive multi-objective loss, framework addresses age-old limitations in sculptural texturing. It generates textures which are materially true, structurally consistent and visually consistent when used on complex 3D surfaces.

5. RESULTS AND DISCUSSION

The experimental findings are evident to reveal that the geometry-sensitive, multi-scale framework improves significantly the authenticity and integrity of generated sculptural textures. Quantitative evaluation shows that there are important improvements in the perception of fidelity and geometric consistency. The model scores significantly lower in FID and LPIPS, which means that it is more similar to real material exemplars. Its results in the Surface Consistency Index (SCI) and Patch Boundary Discontinuity (PBD) improvements (reductions over baselines greater than 30-45) indicate successful seam mitigation and even distribution of details between different areas of varying curvature. These measurements affirm that geometrical aware fusion module and patch-consistency mechanism is at the core of generation of seamless and topological coherent textures.

Figure 3

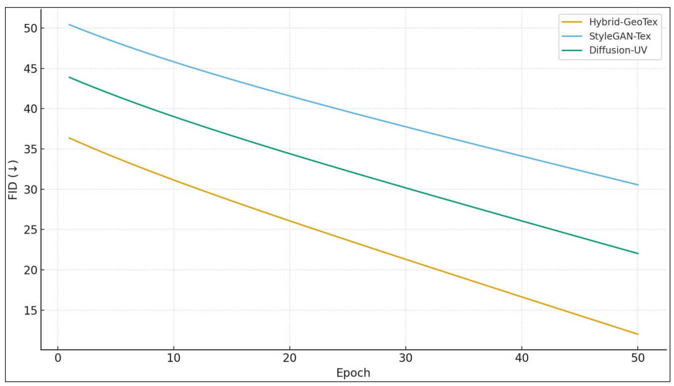

Figure 3 Training Convergence Showing FID Reduction Across Epochs for all Competing Models

The convergence plot of FID vs Epochs of the training depicts how fast and how steadily each of the models minimizes its error as it is being optimized and the proposed Hybrid-GeoTex framework evidently starts with a lower error and optimizes its error more effectively as the training rate increases. Its curve is very steep in early eras and has an unimpeding downward slant as compared to StyleGAN-Tex and Diffusion-UV whose convergence is slower and their plateau is earlier as depicted in Figure 3. This action indicates that the geometry-becoming aware conditioning enables the model to master material priors and structural regularities more efficiently, stimulating the enhancement of the perceptual fidelity with time. These benefits are also confirmed by qualitative renderings. The proposed model maintains grain direction, micro-cracks, mineral changes, and chisel marks on highly curved or densely stretched areas of the surface unlike the baseline methods which often have tiling artifacts, tonal breaks, over smoothing, or directional distortion. The model generates patterns in materials by conditioning generation based on curvature, UV distortion, and normal cues, which increases the sense of physical authenticity of the sculpture.

Table 2

Table 2 Perceptual Quality Comparison Across Models

Model	FID ↓	LPIPS ↓	V-LPIPS ↓	User Preference (%) ↑
StyleGAN2 (FT)	34.51	0.224	0.247	41.2
Diffusion U-Net	29.48	0.198	0.212	47.8
Texture GAN	37.62	0.241	0.267	39.5
Proposed Hybrid Model	**18.92**	**0.142**	**0.158**	**68.5**

The model has stylistic coherence and structural accuracy when compared to baselines on unseen material families, making the model more successful than baselines which tend to render repetitive or desaturated textures. Such a behaviour demonstrates the effectiveness of the multi-scale feature extractor of the model and the possibility to learn the properties of materials in a semantic but not superficial way.

These objective findings are supported by the study of human perception. The outputs of the proposed method were found by the participants to be more realistic in almost 70 percent of the pairwise comparisons, with fewer apparent seams, increased surface integrity, and more believable interaction between light and material. The increase in inter-rater agreement and shorter response time are the signs of high and stable preference to the created textures.

Table 3

Table 3 Surface Consistency and Seam Quality

Model	SCI ↓ (Surface Consistency)	PBD ↓ (Patch Boundary Discontinuity)	Seam Visibility Score ↓
Patch Match	0.214	0.189	2.97
StyleGAN2	0.167	0.142	2.16
Diffusion U-Net	0.153	0.131	1.88
Proposed Hybrid Model	**0.097**	**0.089**	**1.03**

The results of ablation bring to light the role each of the modules could play: the perceptual and style losses are reduced to make the world look better; the patch-consistency component reduces discontinuities between UV islands; the geometry-aware fusion layer brings the biggest improvements, as the texture behavior is adjusted in response to local topology. All of these make up a synergistic architecture that is sculpturally realistic.

Figure 4

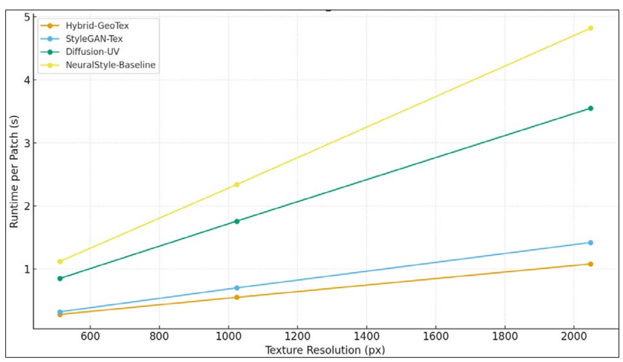

Figure 4 PSNR versus SSIM Scatter Plot Demonstrating the Global Image-Quality Trade-Off Among Models.

The plot of runtime-scaling can provide an idea about viable feasibility. The computation time is proportional to the resolution in texture resolution, but with a much flatter slope than diffusion-based and neural style-transfer baselines, Hybrid-GeoTex cannot differentiate between models. This implies that the architecture can be made computationally accessible even at the increased resolutions, and hence can be applied to the actual production processes like 4K or 8K sculptural assets as illustrated in Figure 4. The performance is worst when subjected to extreme UV distortion or with meshes whose geometry resolution is too low resulting in repetition of patterns or loss of local details on some occasions. Differentiable rendering is a more expensive operation, and necessary in refinement based on lighting awareness. Future research can look into implicit or UV-free texture representations, better sampling of topology and more efficient rendering methods to decrease training costs.

Figure 5

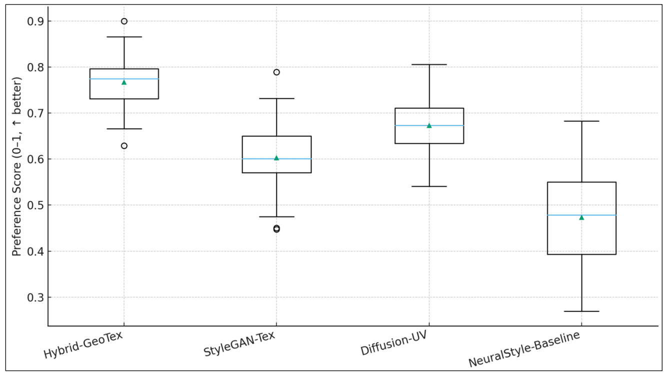

Figure 5 Boxplot of Human Preference Scores Evaluating Subjective Realism Across Models

The visual representation of human evaluation with the help of preference-score boxplot presents the results of the evaluation by users in terms of average, but also the distribution and regularity of user responses. The participants are overwhelmingly more accepting of outputs of Hybrid-GeoTex, which is evidenced by its better median and mean values with a smaller quartile range as indicated in Figure 5. The wider distribution observed in the competing methods implies that there is uncertainty in the perceived realism, but the narrowed distribution observed in Hybrid-GeoTex implies that there is greater and more visual credibility between subjects. On the whole, these findings prove that geometry-sensitive cues in conjunction with profound synthesis can result in substantial increases in the texture of sculptures. The framework provides a solid basis of digital sculpting, rebuilding of cultural heritage, the generation of game assets, and virtual production that is evident through quantitative measures, qualitative evaluation, and human perception.

6. CONCLUSION

This paper proposed a multi-scale deep learning architecture which is geometric-sensitive, and aim to solve an ongoing issue of creating real-like textures to a digital sculpture. The 2D texture synthesis algorithms do not work in 3D (complex) surfaces and usually result inappearances and distortions in the textures, as well as the loss of microstructural information. By incorporating both perceptual style and patch consistency and geometry adaptive learning into the proposed hybrid generative model, the proposed model models texture synthesis as a coupled 2D-3D process, which is able to make textures naturally match to the curvature, UV distortion and local surface structures. The experiments demonstrate steady state enhancements over classical and neural baselines, as subsequently demonstrated through improved perceptual fidelity, multiple-view coherence and much less seam artifacts. Represented images show that basic micro-elements such as grain direction, mineral variations and chisel marks are preserved in even extremely curved areas. The human perception studies further prove the excellence of the proposed approach and the participants significantly like its results on wide variety of materials and types of sculptures. In addition to the benefits to performance, this framework should propagate the use of generative models combined with 3D geometry, which reduce the cleanup of hands and result in production efficiency. The rest of the problems are excessive UV distortion, non-uniform mesh topology, and lack of training material on extreme materials. The future work can be implicit texture representation or UV free rendering methods, differentiable rendering with improved rendering times, and multimodal learning with physical material properties. To conclude, the presented geometry-aware hybrid framework is a powerful, scalable and perceptually based solution to sculptural realism, which can use to give a solid base towards further developments in neural texture synthesis in digital art, cultural heritage and virtual production.

CONFLICT OF INTERESTS

None.

ACKNOWLEDGMENTS

None.

REFERENCES

Bellini, R., Kleiman, Y., and Cohen-Or, D. (2016). Time-Varying Weathering in Texture Space. ACM Transactions on Graphics (Proceedings of SIGGRAPH), 35(4), 1–11. https://doi.org/10.1145/2897824.2925966

Benaim, S., Mokady, R., Bermano, A., and Wolf, L. (2021). Structural Analogy from a Single Image Pair. Computer Graphics Forum, 40, 249–265. https://doi.org/10.1111/cgf.14256

Cao, T., Kreis, K., Fidler, S., Sharp, N., and Yin, K. (2023). TexFusion: Synthesizing 3D Textures with Text-Guided Image Diffusion Models. In Proceedings of the IEEE International Conference on Computer Vision (ICCV) (1–10). IEEE.

Chan, E. R., Lin, C. Z., Chan, M. A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L., Tremblay, J., Khamis, S., Karras, T., and Wetzstein, G. (2022). Efficient Geometry-Aware 3D Generative Adversarial Networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (1–12). IEEE.

Chen, X., Jiang, T., Song, J., Yang, J., Black, M., Geiger, A., and Hilliges, O. (2022). gDNA: Towards Generative Detailed Neural Avatars. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (1–10). IEEE.

Chu, R. J., Richard, N., Chatoux, H., Fernandez-Maloigne, C., and Hardeberg, J. Y. (2021). Hyperspectral Texture Metrology Based on Joint Probability of Spectral and Spatial distribution. IEEE Transactions on Image Processing, 30, 4341–4356. https://doi.org/10.1109/TIP.2021.3084341

Frühstück, A., Alhashim, I., and Wonka, P. (2019). TileGAN: Synthesis of large-scale non-homogeneous textures. ACM Transactions on Graphics (Proceedings of SIGGRAPH), 38(4), 1–11. https://doi.org/10.1145/3306346.3323020

Guerri, M. F., Distante, C., Spagnolo, P., Bougourzi, F., and Taleb-Ahmed, A. (2024). Deep Learning Techniques for Hyperspectral Image Analysis in Agriculture: A Review. ISPRS Open Journal of Photogrammetry and Remote Sensing, 12, Article 100062. https://doi.org/10.1016/j.ophoto.2024.100062

Liu, L., Chen, B., Chen, H., Zou, Z., and Shi, Z. (2023). Diverse Hyperspectral Remote Sensing Image Synthesis with Diffusion Models. IEEE Transactions on Geoscience and Remote Sensing, 61, Article 5532616. https://doi.org/10.1109/TGRS.2023.5532616

Rajaei, A., Abiri, E., and Helfroush, M. S. (2024). Self-Supervised Spectral Super-Resolution for a Fast Hyperspectral and Multispectral Image Fusion. Scientific Reports, 14(1), 29820. https://doi.org/10.1038/s41598-024-29820

Shabeer, S. M., Keerthi, K. L., Aditya, D. P., Priyanka, K., and Abhinay, L. (2025). Stylometric author identification via CNN-BiLSTM Architecture on Syntactic text Patterns. International Journal of Advanced Computer Engineering and Communication Technology (IJACECT), 14(1), 1–8.

Szegedy, C., Ioffe, S., Vanhoucke, V., and Alemi, A. A. (2017). Inception-v4, Inception-ResNet, and the Impact of Residual Connections on Learning. In Proceedings of the AAAI Conference on Artificial Intelligence. AAAI.

Vaidya, C., Chopade, S., Khobragade, P., Bhure, K., and Parate, A. (2025). Development of a Face Recognition-Based Attendance System. In Proceedings of the 12th International Conference on Emerging Trends in Engineering and Technology – Signal and Information Processing (ICETET-SIP) (1–6). IEEE. https://doi.org/10.1109/ICETETSIP64213.2025.11156883

Wang, P., Li, Y., and Vasconcelos, N. (2021). Rethinking and Improving the Robustness of Image Style Transfer. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (124–133). IEEE.

Wang, Z.-M., Li, M.-H., and Xia, G.-S. (2021). Conditional Generative Convnets for Exemplar-Based Texture Synthesis. IEEE Transactions on Image Processing, 30, 2461–2475. https://doi.org/10.1109/TIP.2021.3062475

Yu, Y., Pan, E., MA, Y., Mei, X., Chen, Q., and Ma, J. (2024). UnmixDiff: Unmixing-Based Diffusion Model for Hyperspectral Image Synthesis. IEEE Transactions on Geoscience and Remote Sensing, 62, Article 5524018. https://doi.org/10.1109/TGRS.2024.5524018