

TransPointNet++: A Hybrid Model for 3D Medical Orthotic Surface Reconstruction from Point Clouds*

Sumaia ABDULWAHAB AL-WAHASHI¹ and Aicha BEN MAKHLOUF^{1,2}

¹Department of Computer Science, College of Computer, Qassim University
Buraydah, 52571, Qassim, Saudi Arabia

²Université de Sousse, Ecole Nationale d'Ingénieurs de Sousse,
LATIS- Laboratory of Advanced Technology and Intelligent Systems, 4023, Sousse, Tunisia

Correspondence should be addressed to: Sumaia ABDULWAHAB AL-WAHASHI, 461217096@qu.edu.sa

*Presented at the 47th IBIMA International Conference, 29-30 June 2026, Madrid, Spain

Abstract

3D modeling of the anatomical structures based on 3D scan information is essential in the production of customized medical instruments. Raw point cloud data is however, in most cases quite noisy, incomplete and is sampled quite irregularly making it difficult to reconstruct with much reliability. This paper presents a comparative analysis of advanced deep learning approaches for 3D reconstruction, including a pointbased model, PointNet++, and a Transformer-based implicit surface reconstruction model, as well as a proposed hybrid framework that integrates both paradigms. The performance of these approaches is evaluated using standard geometric metrics, including Chamfer Distance, Hausdorff Distance, and Normal Consistency, in addition to classification-based measures such as Precision, Recall, and F-score. Experimental results demonstrate that while PointNet++ effectively preserves local geometric details and achieves high precision, the Transformer-based model provides improved global context modeling and surface smoothness, and the proposed hybrid model achieves the best overall performance. Specifically, it provides a balanced trade-off between geometric accuracy, surface continuity, and structural consistency, outperforming individual models across multiple evaluation criteria. These findings highlight the advantages of integrating complementary 3D representation paradigms and provide practical insights for selecting appropriate reconstruction techniques in the automated and scalable design of personalized anatomical models, reducing the need for manual intervention in medical applications.

Index Terms: 3D reconstruction, 3D scan, point cloud, PointNet++, Transformer, hybrid model, geometric learning, surface modeling

Introduction

Recent developments in deep learning have had a profound effect on the medical world, especially on the medical imaging and 3D data analysis. With the application of artificial intelligence to clinical processes, interpretation and readability of complex anatomical data has been more precise and efficient [1]. Of these applications, the 3D reconstruction is instrumental in the reconstruction of entire anatomical structures out of incomplete, noisy and sparsely sampled data [2].

The use of orthotic devices to help sedate and improve the work of the musculoskeletal system applies to different biomechanical disorders [3]. Conventionally, they have been designed using manual methods like plaster

Cite this Article as: Sumaia ABDULWAHAB AL-WAHASHI and Aicha BEN MAKHLOUF Vol. 2026 (12) "TransPointNet++: A Hybrid Model for 3D Medical Orthotic Surface Reconstruction from Point Clouds " Communications of International Proceedings, Vol. 2026 (12), Article ID 4723726, <https://doi.org/10.5171/2026.4723726>

computer-aided design (CAD) that are time consuming, laborious, and demand specialized skills [4]. With the development of 3D scanning technologies, it is now possible to obtain patient-specific anatomical data in the point-cloud format. Point clouds are a flexible and efficient 3D geometry representation, but are unstructured by nature and can be susceptible to noise, broken regions, and an uneven sampling density [5], [6]. These complications render it hard to rebuild the proper and fine high quality 3D reconstructed surfaces.

Deep learning techniques have been widely developed to process 3D data directly from point clouds. Point-based approaches such as PointNet++ learn hierarchical geometric features from unordered point sets and are highly effective in capturing fine-grained local structures. In contrast, Transformer-based implicit representation models encode global contextual relationships and represent 3D geometry as continuous functions, enabling the reconstruction of smooth and coherent surfaces. Despite their strengths, these approaches differ in their ability to preserve structural detail versus surface smoothness. This paper presents a comparative analysis of PointNet++, a Transformer-based reconstruction model, and a proposed hybrid framework that integrates both approaches for 3D reconstruction from imperfect point cloud data. The evaluation is conducted based on geometric accuracy, structural detail preservation, and surface quality. The results show that while PointNet++ excels in local feature accuracy and the Transformer-based model provides superior global consistency and smoothness, the proposed hybrid model achieves the most balanced overall performance across all evaluation criteria. These findings highlight the complementary nature of local and global representations and support the selection of hybrid strategies for accurate and scalable medical 3D reconstruction.

State Of the Art

Medical Imaging and Segmentation-Based Approaches

Deep learning has been extensively applied in medical imaging, particularly for segmentation and diagnostic analysis. Eleonora et al [8] proposed a convolutional neural network (CNN)-based workflow for automatic bone segmentation from cone-beam computed tomography (CBCT) images, followed by 3D surface reconstruction for orthopedic applications. While the method demonstrates improved accuracy compared to manual segmentation, it is limited to image-based processing and does not extend to full geometric reconstruction or orthotic design. Similarly, Miao et al [9] introduced a 3D ensemble framework (3D-ONFHNet) utilizing volumetric MRI data for the diagnosis and staging of osteonecrosis. The model achieves high diagnostic performance; however, it focuses on classification rather than reconstructing anatomical geometry. In a related direction, Reyhaneh et al [10] proposed a multi-view deep learning approach to estimate 3D anatomical landmarks from 2D radiographic images. Although the method improves measurement accuracy and reduces manual annotation effort, it does not generate complete 3D models suitable for downstream design tasks. These studies demonstrate the effectiveness of deep learning in medical image analysis; however, they primarily address segmentation and diagnosis, leaving a gap in end-to-end reconstruction and device generation.

3D Reconstruction from Partial and Noisy Data

A large body of research has focused on reconstructing complete 3D anatomical structures from incomplete or noisy inputs. Matthias et al [11] conducted a comprehensive evaluation of multiple state-of-the-art models, including GAN-based approaches, transformer-based architectures such as PoinTr and SnowflakeNet, and VAE-based models for reconstructing partial foot scans. The study highlights the strengths of transformer-based models in capturing fine details; however, it relies on controlled datasets, limiting generalization to real-world scenarios. Adnan et al [12] proposed a deep learning framework that directly operates on point cloud data to reconstruct smooth and complete anatomical surfaces from sparse inputs. While the method achieves high geometric accuracy, it does not address orthotic generation. Similarly, Nolan et al [13] combined statistical shape modeling with deep learning to reconstruct anthropomorphic limb geometry from partial observations, leveraging learned anatomical priors. Although effective, this approach may struggle with highly diverse anatomical variations. More advanced techniques have been introduced in recent years. Oliver et al [14] proposed an implicit neural representation model that learns continuous deformation fields for articulated foot reconstruction. This approach enables high-resolution outputs and flexible modeling but requires complex training and significant computational resources. Prasoon et al [19] further explored deep learning methods for reconstruction, augmentation, and registration, emphasizing the ability of neural networks to learn anatomical priors and improve robustness to noise and missing data. Additionally, Nolan et al [15] and Kyle et al [20] addressed reconstruction under limited observations, where Kyle proposed a deep learning model capable of reconstructing foot geometry from a single RGB image using parametric shape representations. While this method

demonstrates strong performance in unconstrained environments, it relies heavily on learned priors and may not capture patient-specific fine details accurately. Jie-Wen et al [14] contributed a high-resolution multi-view point cloud dataset of the foot-ankle complex during gait, enabling the development of robust reconstruction models under dynamic and occluded conditions. However, the work focuses on data generation rather than algorithmic design.

Digital Orthotic Design and Additive Manufacturing

The integration of 3D scanning and additive manufacturing has led to significant advancements in orthotic design. Leandro et al [16] presented a comprehensive overview of 3D printing in orthopedics, describing the full pipeline from data acquisition to fabrication. Although the approach enables customization, it remains time-consuming and dependent on manual processes. Venkateswaran et al [17] proposed a multiview 3D reconstruction system using depth sensors and point cloud registration techniques for orthotic brace design. The system achieves high geometric accuracy but still requires manual CAD intervention. Similarly, Betsy et al [18] introduced a rapid workflow for user-tailored orthosis design, integrating biocompatibility testing and finite element analysis. Despite improving accessibility, the process is not fully automated. Vaia et al [19] proposed an integrative computational design framework combining parametric modeling and finite element analysis for personalized orthopedic devices. While this approach enhances customization and performance optimization, it still relies on clinician interaction and does not incorporate deep learning for automated design generation. Zhigang et al [20] reviewed innovations in 3D-printed orthotics, highlighting improvements in customization, material efficiency, and biomechanical performance. However, the study emphasizes the lack of standardized workflows and long-term clinical validation. Chang-Jin et al [21] demonstrated the clinical effectiveness of personalized 3D-printed prostheses for reconstructing bone defects, achieving accurate anatomical restoration and improved functional outcomes. Despite these advances, such approaches are primarily surgical and do not focus on external orthotic devices.

Deep Learning for Orthotic Design and Foot Analysis

Recent studies have explored the integration of deep learning with orthotic design and foot analysis. Pui-ling et al [22] proposed a deep learning model for classifying 3D foot types in diabetic patients using geometric features extracted from 3D scans. The model achieved high accuracy and supports clinical decision-making; however, it does not generate orthotic models. Joshua et al [23] conducted a clinical evaluation of AI-assisted 3D-printed orthotics using plantar pressure data and gait analysis. The results demonstrated improved biomechanical performance and user comfort, highlighting the benefits of data-driven orthotic design. However, the system relies on additional biomechanical data and does not address reconstruction from raw scans. Deqiang et al [24] introduced an unsupervised learning framework for modeling reference bony shapes using a surface deformation network for orthognathic surgical planning. The approach demonstrates the ability of deep learning models to learn anatomical variability and generate plausible 3D structures without requiring extensive labeled data. In a related study, Deqiang et al [25] also proposed a point cloud-based deep learning method for estimating reference anatomical shapes, further emphasizing the effectiveness of data-driven modeling in capturing complex geometric variations. Despite these contributions, both approaches are primarily focused on surgical planning and do not extend to orthotic device generation. Additionally, Vaia et al [26] and Betsy et al [27] emphasize semi-automated workflows, where design optimization is guided by computational tools but still requires expert input. This highlights a key limitation in current systems: the lack of fully automated pipelines that integrate reconstruction and orthotic design. This research proposes a hybrid framework that integrates the PointNet++ model for hierarchical feature extraction with a Transformer-based implicit surface reconstruction module. By combining the strengths of local geometric feature learning and global contextual modeling, the proposed approach aims to achieve accurate, smooth, and geometrically consistent reconstruction of anatomical structures from incomplete and noisy input data.

Methodology

Overview

The proposed methodology contrasts three different deep learning architectures: an autoencoder based on Transformer models, a PointNet++ hierarchical architecture, and a unique Hybrid model that incorporates both styles through bidirectional cross-attention fusion. Each of these models receives a sparse input point cloud to reconstruct a denser surface representation that accurately captures the geometry of the original surface. In fact, the proposed Hybrid model aims to improve upon individual PointNet++ and Transformer models through local geometric accuracy and global contextual modeling, respectively. The Hybrid model consistently outperforms each of the six-evaluation metrics, due

to the combined effect of dual-decoder ensemble architecture, coverage-aware loss, and feature-weighted seed selection.




Fig. 1. Process Overview. Source: Authors' own work.

Data Description

This research analyses the data from the MANO (hand Model with Articulations and Non-rigid defOrmations) dataset [?], which is a free parametric hand dataset that provides realistic 3D hand meshes in the form of point clouds along with their surface normal. This advanced dataset has a fully differentiable statistical model of the human hand that adapts to both shape and pose, and is trained on 1,500 3D hand scans from 31 individuals whose hand shape vary greatly. Each sample from MANO consists of an XYZ point cloud file that is a vector of the form $[x, y, z, nx, ny, nz]$. Each vector is a coordinate of a point within three-dimensional space and the accompanying surface normal at that point. The relevance of the ground-truth normal presented in this work is for the evaluation metrics of normal consistency.

Manual Finger Segmentation

To direct the task of reconstruction towards the finer anatomical levels, the individual finger parts were separated from the entire point cloud of the MANO hand. The procedure took into account the finger parts and selective inspection of the geometry to obtain separate point clouds for the fingers, which displayed the surface topology that is characteristic to the reconstruction task. The resulting step of this segmentation was a dataset containing 500-point clouds of fingers while retaining the six-dimensional constructions of $[x, y, z, nx, ny, nz]$. Sample sizes differ across these datasets because of the a given finger form variation and the granularity of the MANO meshes.

Dataset Split

The 500 samples were divided into three non-overlapping pre-defined splits to ensure a rigorous evaluation: The adopted

TABLE I
DATASET SPLIT STATISTICS

Split	N Samples	Percentage	Purpose
training	350	70%	Model optimization
Validation	50	10%	Hyperparameter tuning and early stopping
Test	100	20%	Final performance evaluation

70:10:20 split was selected to ensure a robust and fair evaluation of all investigated models. A large training subset (70%) was allocated to enable effective learning of 3D geometric features, while 10% of the data was reserved for validation to support hyperparameter tuning and early stopping. The remaining 20% was kept as an unseen test set for final performance assessment. This partitioning strategy is commonly employed in deep learning research and provides a suitable balance between model development and unbiased evaluation.

Data Pre-Processing

The samples substantially differ from each other in their point cloud files especially with respect to the number of points, their spatial sizes, and the unity of their origin. To maintain the consistency of samples, a pre-processing step is adhered to all samples across all models. The pre-processing steps are the following:

- **Loading and Validity Filtering:** Every XYZ file is checked, and only rows with six numerical entries separated by spaces are kept. NaN and infinite rows are also removed for numerical stability with the models.
- **Unit-Sphere Normalization** The point clouds are scaled and positioned. Each point cloud is then scaled to a bounding unit sphere. The centroid and scale factor is stored for later to change the denormalized point cloud to its original form.
- **Random Point Sampling** The input sizes of the point clouds are inconsistent. Clouds are encoded to a size of 1,024 by uniform random sampling. For point clouds greater than 1,024, sampling is done with replacement. The decoder needs 2,048-point clouds to remain a constant denser reference surface during reconstruction.

Model Architecture

- 1) **Transformer model:** The purpose of this model is to use deep learning techniques based on the transformer model to capture the global relationships of unordered 3D point clouds to perform surface reconstruction. The model takes a point cloud consisting of a collection of points in 3D space with a set of features and converts those points into high-dimensional embeddings using a shared Multi-Layer Perceptron (MLP). The embeddings are treated as tokens and passed to a Transformer Encoder. A Stack of Self-Attention layers is used. The stacking of self-attention layers captures the long-range interdependencies and global-contextual information, which is necessary to understand complex geometric relationships. The proposed model incorporates a seed-anchored decoding model as opposed to the traditional encoder-decoder Transformer model. A collection of representative points, referred to as seeds, is obtained from the point cloud using the Farthest Point Sampling (FPS) method to obtain uniform coverage to represent the entire object. These representative point features are then enhanced with a self-attention layer of light architecture in a self-attention layer to aid in the preservation of local coherence. Each enhanced seed feature is then processed by a prediction head which is an MLP, and in this case, produces normal vectors. A small number of points close to the surface and normal to the surface are generated. The orientation and standoff of the reconstructed points are closer to the surface of the point cloud, as is the final point cloud sample which is obtained by combining all of the points generated by each of the modified seeds.




Fig. 2. Transformer Model Architecture. Source: Authors' own work.

PointNet++ model

The PointNet++ model learns hierarchical geometric features from point clouds by analyzing the local and global structures in the data. Unlike the Transformer-based model, which emphasizes global attention, PointNet++ uses a multiscale hierarchical approach to model local neighborhoods. The encoder comprises a series of Set Abstraction (SA) layers, where the SA layers perform three functions in particular: point sampling, neighborhood grouping, and

local feature extraction. In the SA layers, a small batch of points is sampled using a technique called Farthest Point Sampling (FPS). Next, the neighborhood group points are grouped, and the grouped points are processed with shared Multilayer Perceptron's (MLPs), which are then followed by a maxpooling operation to extract local features. The model is able to characterize and store a hierarchy of features on the point clouds to capture the increasingly coarse and fine geometric characteristics by stacking SA layers, each of which is configured with a smaller point set. Point-wise features are restored by the decoder utilizing Feature Propagation (FP) layers, which do inverse distance weighted interpolation to upsample features from the lower representation layers back to the original point set. Skip connections are used to combine interpolated features with earlier layer features, preserving spatial detail and improving reconstruction quality. Similar to the Transformer-based model, the final reconstruction is achieved through a seed-anchored residual decoding module. A set of seed points is selected using FPS from the original input, and their corresponding features are refined through MLP layers. Each seed then generates multiple output points via predicted residual offsets and normals. This strategy ensures that the reconstructed points are spatially consistent with the input geometry. Overall, the proposed PointNet++ architecture effectively captures local geometric structures through hierarchical learning while maintaining global coherence through feature propagation, making it well-suited for high-quality 3D surface reconstruction.




Fig. 3. PointNet++ Model Architecture. Source: Authors' own work.

Hybrid model

This hybrid model integrates the strengths both PointNet++ and the Transformer architectures to build an effective and accurate system for the 3D point cloud reconstruction from sparse noisy input. The system has both a local learning component and a feature global learning component to allow geometric learning and contextual modeling to coexist. The architecture has two parallel branches. The first branch incorporates PointNet++ layers to derive local architectural feature hierarchies. PointNet++ Set Abstraction layers perform the task of local feature extraction through multiple shared MLPs. Set Abstraction layers also perform local feature maxpooling and layered neighborhood groupings through the ball query, and point sample using the farthest point sample. These multiple layers perform a geometric feature compression from a fine to a coarse structure. Feature Propagation layers then upsample and interpolate the fine local features to the original point cloud and spatial resolution, respectively. The second branch uses a Transformer Encoder to model the global geometric relationships of the point cloud. The input features are first mapped to a high-dimensional embedding space. The highly dimensional embedding space features are then refined through multiple self-attention layers, allowing the model to learn the global relationships of the point cloud. A fusion module with bidirectional cross-attention is then introduced to integrate the transformer representation with the PointNet++ representation. In this module, global features attend to local features and local features attend to global features, allowing them to exchange mutual information. This style of crossattention pairs coarse representation and local geometric detail with fine pixel detail and extended context. This representation for the model provides the 3D geometric information needed. For reconstruction, the model utilizes a dual-seed-anchored decoding technique. A set

of representative seed points is selected from the input using feature-guided FPS to ensure good spatial coverage. Two parallel decoders are then applied: A fine decoder, which generates points with small residual offsets, focusing on high-precision surface details. A coarse decoder, which produces points with larger offsets, improving coverage and capturing broader geometric structures. Each decoder predicts multiple points per seed in the form of residual displacements relative to the seed positions, along with their corresponding normal vectors. This residual formulation ensures that generated points remain close to the underlying geometry while allowing flexibility in reconstruction.

Evaluation Metrics

Six geometric metrics are defined to evaluate the reconstruction quality. All metrics are computed on the test set and the full, original point clouds (all available points) as ground truth. This is used to avoid inflating results from a random, re-sample set that is too similar to the trained distribution. A distance threshold of $\tau = 0.05$ (relative to the unit sphere) is used for precision and recall.

- **Chamfer Distance:** Sum of bidirectional mean squared nearest-neighbor distances between predicted and GT point sets.
- **Precision:** Fraction of predicted points within threshold τ of any GT point




Fig. 4. Hybrid Model Architecture. Source: Authors' own work.

- **Recall:** Fraction of GT points within threshold τ of any predicted point.
- **F-Score:** Harmonic mean of Precision and Recall
- **Normal Consistency:** Mean $|\cos\theta|$ between matched predicted and GT normal at nearest neighbors.
- **Hausdorff Distance:** Symmetric maximum directed Hausdorff distance: $\max(h(P,Q), h(Q,P))$

Hausdorff Distance is the maximum distance from a point in one set to the nearest point in the other set. Focuses on the "worst-case" mismatch. If a single point is far away, the Hausdorff distance is high. The Normal Consistency metric checks whether the reconstructed normal of the surface are consistent with the orientation of the original surface and is essential for processes like physical simulation and rendering. Chamfer distance is the sum of squared distances between each point in one set to its nearest neighbor in the other set. It is robust to outliers and works well for comparing 3D point clouds. Our comparison between the models will consider the less Chamfer distance.

TABLE II: QUANTITATIVE EVALUATION ON TEST SETS. TR = TRANSFORMER, PN2 = POINTNET++, HYB = HYBRID

Metric	TR-Test	PN2-Test	HYB-Test
Precision	0.9936	0.9954	0.9958
Recall	0.9986	0.9983	0.9993
F-Score	0.9960	0.9969	0.9975
Normal Consist.	0.9893	0.9802	0.9922
Chamfer Dist.	0.0006836	0.0006965	0.0006488
Hausdorff Dist.	0.0681	0.0687	0.0638

Development Results

Implementation Details

All models are implemented in Python 3.10 and PyTorch

2.2. Model training and evaluation are done on an NVIDIA T4 GPU made available via Google Colaboratory. KD-tree nearest-neighbor queries and Hausdorff distance are computed using the SciPy library. Other library dependencies are the data utilities, Poisson, Ball-Pivoting surface reconstruction of scikit-learn, Open3D and Matplotlib.

The models are trained on the same 350-sample training split, validated on 50 samples, and evaluated on the heldout 100-sample test set. The evaluation protocol uses the full original point clouds (all available points) as ground truth, rather than a re-sampled subset, to prevent artificially inflated metrics. The Transformer and PointNet++ models are each trained for 100 epochs, while the Hybrid model is trained for up to 150 epochs with early stopping (patience = 20 epochs). All models use the AdamW optimizer with an initial learning rate of 1×10^{-4} and cosine annealing to 1×10^{-6} .

TABLE III: EXPERIMENTAL CONFIGURATION FOR ALL THREE MODELS

Property	Transformer	PointNet++	Hybrid
Batch size	8	8	8
Encoder depth	6 layers	SA1-SA3	8 layers +SA/FP
Outpoints	2,048	2,048	2,048

Chamfer Distance measures the overall geometric fidelity of the reconstruction. The Hybrid model achieves a test Chamfer Distance of 0.0006488, then the Transformer achieves 0.0006836 and outperforms PointNet++ 0.0006965. The Hybrid achieves the best result, the lowest absolute value, confirming that the dual-decoder ensemble with fine-scale offsets (scale = 0.02) produces predictions that are, on average, closest to the ground truth surface. The PointNet++ model exhibits the highest Chamfer Distance, suggesting that its SA/FP pipeline alone does not provide sufficient coverage of the dense target surface despite its strong local feature learning.




Fig. 5. Hybrid Model Result. Source: Authors' own work.

Conclusion

This paper presented a comprehensive evaluation of three deep learning models for 3D point cloud surface reconstruction from the custom MANO-derived finger dataset. The experimental results demonstrate that the proposed Hybrid model consistently outperforms both the Transformer and PointNet++




Fig. 6. Transformer Model Result. Source: Authors' own work.




Fig. 7. PointNet++ Model Result. Source: Authors' own work.

baselines on the held-out test set of 100 samples. The Hybrid model achieves the best performance Chamfer Distance (0.0006488).

These results validate the architectural hypothesis motivating the Hybrid design: combining local geometric feature learning (PointNet++ branch) with global shape context modelling (Transformer branch) through bidirectional crossattention, feature-weighted seed selection, and dual-decoder ensemble learning produces a reconstruction model that is strictly superior to either component architecture in isolation. The results also confirm that the custom MANO-derived finger dataset poses a meaningful reconstruction challenge with discriminative difficulty across all six metrics. Figure 8 illustrates the steps followed in order to reconstruct the final orthotic model from a MANO hand scan.




Fig. 8. Steps of the orthotic model reconstruction. Source: Authors' own work.

References

- G. Eason, B. Noble, and I. N. Sneddon, "On certain integrals of Lipschitz-Hankel type involving products of Bessel functions," *Phil. Trans. Roy. Soc. London*, vol. A247, pp. 529–551, April 1955.
- J. Clerk Maxwell, *A Treatise on Electricity and Magnetism*, 3rd ed., vol. 2. Oxford: Clarendon, 1892, pp.68–73.
- I. S. Jacobs and C. P. Bean, "Fine particles, thin films and exchange anisotropy," in *Magnetism*, vol. III, G. T. Rado and H. Suhl, Eds. New York: Academic, 1963, pp. 271–350.
- K. Elissa, "Title of paper if known," unpublished.
- R. Nicole, "Title of paper with only first word capitalized," *J. Name Stand. Abbrev.*, in press.
- Y. Yorozu, M. Hirano, K. Oka, and Y. Tagawa, "Electron spectroscopy studies on magneto-optical media and plastic substrate interface," *IEEE Transl. J. Magn. Japan*, vol. 2, pp. 740–741, August 1987 [Digests 9th Annual Conf. Magnetics Japan, p. 301, 1982].
- M. Young, *The Technical Writer's Handbook*. Mill Valley, CA: University Science, 1989.
- Tiribilli, E. and Bocchi, L. (2024) 'Deep learning-based workflow for bone segmentation and 3D modeling in cone-beam CT orthopedic imaging', *Applied Sciences*, 14(17), p. 7557.
- Yang, M., Hsiang, F., Li, C., Chen, X., Zhang, C., Sun, G., Lou, Q., Zhu, W., Zhao, H., Liu, F. and Ding, X. (2025) 'Development and evaluation of a 3D ensemble framework for automatic diagnosis of early osteonecrosis of the femoral head based on MRI: A multicenter diagnostic study', *Frontiers in Surgery*, 12, p. 1555749.
- Rostamian, R., Shariat Panahi, M., Karimpour, M., Kashani, H.G. and Abi, A. (2025) 'A deep learning-based multi-view approach to automatic 3D landmarking and deformity assessment of lower limb', *Scientific Reports*, 15(1), p. 534.
- Jager, M., Eberhardt, J. and Cunningham, D.W. (2024) '3D reconstruction of partial foot scans using different state-of-the-art neural network approaches', *Footwear Science*, 16(2), pp. 105–114.

- Anwer, A., Malik, J., Khurshid, K., Elhayek, A., Battrawy, R., Stricker, D., Naveed, S.S. and Safder, M.U. (2025) 'Scene flow based deep network for hand reconstruction using depth images', *Scientific Reports*, 15(1), p. 32726.
- Li, J.W., Ye, Z.H., Zhou, Q., Song, J., He, Y., Fei, B. and Chen, W.M. (2025) 'A Multi-View High-Resolution Foot-Ankle Complex Point Cloud Dataset During Gait for Occlusion-Robust 3D Completion', *arXiv preprint*, arXiv:2507.11037.
- Boyne, O., Charles, J. and Cipolla, R. (2022) 'FIND: An unsupervised implicit 3D model of articulated human feet', *arXiv preprint*, arXiv:2210.12241.
- Lunscher, N. and Zelek, J. (2017) 'Deep learning anthropomorphic 3D point clouds from a single depth map camera viewpoint', in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pp. 689–696.
- Vinodkumar, P.K., Karabulut, D., Avots, E., Ozcinar, C. and Anbarjafari, G. (2024) 'Deep learning for 3D reconstruction, augmentation, and registration: A review paper', *Entropy*, 26(3), p. 235.
- Fogarty, K., Yang, J., Patodi, C.K., Foster, J., Bhanti, A., Chacko, S., Oztireli, C. and Bonde, U. (2025) 'Best Foot Forward: Robust Foot Reconstruction in-the-wild', in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4373–4382.
- Venkateswaran, N., Hans, W.J. and Padmapriya, N. (2021) '3D design of orthotic casts and braces in medical applications', *Advances in Materials and Processing Technologies*, 7(1), pp. 136–149.
- Chaparro-Rico, B.D., Martinello, K., Fucile, S. and Cafolla, D. (2021) 'User-tailored orthosis design for 3D printing with PLACTIVE: A quick methodology', *Crystals*, 11(5), p. 561.
- Tsiokou, V., Papatheodorou, A., Ntenekou, D., Zouboulis, P. and Karatza, A. (2023) 'An Integrative Computational Design Workflow and Validation Methodology for 3D-Printed Personalized Orthopedic Devices: Case Study of a Wrist-Hand Orthosis (WHO)', *Processes*, 11(7), p. 2204.
- Ejnisman, L., Gobbato, B., de Franc, a Camargo, A.F. and Zancul, E. (2021) 'Three-dimensional printing in orthopedics: From the basics to surgical applications', *Current Reviews in Musculoskeletal Medicine*, 14(1), pp. 1–8.
- Qu, Z., Yue, J., Song, N. and Li, S. (2024) 'Innovations in threedimensional-printed individualized bone prosthesis materials: Revolutionizing orthopedic surgery: A review', *International Journal of Surgery*, 110(10), pp. 6748–6762.
- Yon, C.J., Choi, B.C., Lee, J.M. and Lee, S.W. (2025) 'Personalized 3Dprinted prostheses for bone defect reconstruction after tumor resection in the foot and ankle', *Journal of Functional Biomaterials*, 16(2), p. 62.
- Kubach, J., Pasurka, M., Lueg, J. and Betsch, M. (2025) 'Clinical evaluation of novel custom 3D-printed meshed-silicone orthotics utilizing standing foot scans and dynamic gait data', *Biomedical Engineering and Computational Biology*, 16, p. 11795972251371476.
- Li, P.L., Xiao, Q.F., Yick, K.L., Liu, Q.L. and Zhang, L.Y. (2025) 'A novel deep learning approach to classify 3D foot types of diabetic patients', *Scientific Reports*, 15(1), p. 13819.
- Xiao, D., Deng, H., Lian, C., Kuang, T., Liu, Q., Ma, L., Lang, Y., Chen, X., Kim, D., Gateno, J. and Shen, S.G. (2021) 'Unsupervised learning of reference bony shapes for orthognathic surgical planning with a surface deformation network', *Medical Physics*, 48(12), pp. 7735–7746.
- Xiao, D., Lian, C., Deng, H., Kuang, T., Liu, Q., Ma, L., Kim, D., Lang, Y., Chen, X., Gateno, J. and Shen, S.G. (2021) 'Estimating reference bony shape models for orthognathic surgical planning using 3D point-cloud deep learning', *IEEE Journal of Biomedical and Health Informatics*, 25(8), pp. 2958–2966.
- Romero, J., Tzionas, D. and Black, M.J. (2022) 'Embodied hands: Modeling and capturing hands and bodies together', *arXiv preprint*, arXiv:2201.02610.