OPEN ACCESS

Original Article

DOI: 10.53389/RJAHS.2026050109

Segmentation of Charcot's Neuropathy (Diabetic Foot) Using Medical Images and Artificial Intelligence

Fawaz Nadeem[1], Zartasha Mustansar[1]*, Sumair Naseem Qureshi[2], Arslan Shaukat[3], Haider Ali[1, 4]

[1]School of Interdisciplinary Engineering & Sciences, National University of Sciences & Technology, Islamabad, Pakistan

[2]Shifa Medical Center, Islamabad, Pakistan

[3]Department of Computer Science, College of Electrical and Mechanical Engineering, Rawalpindi, Pakistan

[4]Department of Mechanical and Aerospace Engineering, Air University, Islamabad, Pakistan

Author`s Contribution

ZM Conception and design, ZM, AS, HA Collection and assembly of data, ZM, FN, SNQ Analysis and interpretation of the data, FN Statistical expertise, ZM Final approval and guarantor of the article

Article Info.

Received: April 15, 2026
Acceptance: May 25, 2026

Conflict of Interest: None

Funding Sources: None

Address of Correspondence

Zartasha Mustansar
School of Interdisciplinary Engineering & Sciences, National University of Sciences & Technology,
Email: zmustansar@sines.nust.edu.pk

A B S T R A C T

Background: Charcot neuroarthropathy (CN) is a diabetes-associated neuropathic condition that can cause progressive bone and joint destruction. The scarcity of annotated radiographs limits the development and evaluation of deep learning models for radiographic image segmentation.

Objectives: This study developed and evaluated a generative adversarial network (GAN)-augmented deep learning framework for segmentation of the affected bone/joint region in Charcot foot radiographs.

Methodology: Twenty plain radiographs from four patients with CN stages 0-2 were included. Fourteen original radiographs were used to train SinGAN-Seg, which generated 560 synthetic radiographs with corresponding segmentation masks. Synthetic-image similarity was assessed using single-image Fréchet Inception Distance (SIFID) and Inception Score (IS). Six U-Net-based segmentation architectures were trained using the synthetic dataset. Of the 560 synthetic image–mask pairs, 504 (90%) were used for training and 56 (10%) were reserved for testing. Segmentation performance was evaluated using the intersection over union (IoU).

Results: The SIFID values ranged from 77.56 to 125.81 for images generated from reference image A and from 139.92 to 196.91 for images generated from reference image B. Across the evaluated architectures, mean IoU values ranged from 0.642 to 0.757 on the held-out synthetic test images. R2U-Net obtained the highest observed mean IoU (0.757); however, no formal statistical comparison was performed, and this result should be interpreted cautiously.

Conclusion: GAN-generated image–mask pairs enabled preliminary training and evaluation of U-Net-based models for segmenting the annotated affected bone/joint region in Charcot foot radiographs. Larger patient-level, externally validated datasets are needed before conclusions can be made regarding diagnostic accuracy, clinical reliability, or generalizability.

Key words: Charcot neuroarthropathy; Charcot foot; plain radiography; synthetic data; generative adversarial networks; image segmentation; U-Net; intersection over union

Introduction

Charcot neuroarthropathy (CN), commonly termed Charcot foot, is a diabetes-associated neuropathic disorder characterized by progressive bone fragmentation, joint subluxation, collapse, and deformity.[1] Peripheral neuropathy and repetitive, unperceived microtrauma contribute to progressive structural destruction of the foot. If recognition and treatment are delayed, CN may lead to marked deformity, ulceration, instability, and increased risk of amputation.[2,3]

CN should be differentiated from other diabetic-foot complications. Diabetic foot ulceration is primarily a soft-tissue lesion, whereas osteomyelitis is an infectious process affecting bone.[4,5] Clinical and radiographic features can overlap, particularly in acute CN; however, management differs substantially. CN requires prompt immobilization and offloading, whereas osteomyelitis may require antimicrobial therapy and, in selected cases, surgical treatment.[1,6,7]

Plain radiography is widely available and commonly used for the initial assessment and follow-up of CN.[7] However, early radiographic abnormalities may be subtle, and interpretation can be complicated by disease-stage variation and overlap with other osseous abnormalities.[8] MRI and CT may provide complementary information in selected clinical cases, particularly when differentiation from osteomyelitis is necessary.[2,7] This study focuses exclusively on segmentation of plain foot radiographs; it does not evaluate multimodal imaging or diagnostic differentiation between CN and osteomyelitis.

Deep-learning segmentation methods can delineate anatomical structures and pathological regions at the pixel level in medical images. U-Net is a widely used encoder–decoder architecture designed for biomedical image segmentation and can be trained with relatively limited annotated data through appropriate data-augmentation strategies.[9] However, supervised segmentation still requires sufficiently large, representative, and accurately annotated datasets. This is challenging for CN because the condition is relatively uncommon, image sharing is constrained by privacy considerations, and expert mask annotation is resource intensive.[10,11]

Synthetic-image generation may help augment small annotated datasets. SinGAN-Seg is a generative framework that produces synthetic medical image–mask pairs from limited source data and was developed to support segmentation-model training when access to large real-image datasets is restricted.[12] However, generated images may remain dependent on their original source radiographs; therefore, synthetic-data performance should not be interpreted as independent clinical validation or evidence of diagnostic generalizability.

Accordingly, this study aimed to develop and conduct a preliminary internal evaluation of a GAN-augmented deep-learning framework for segmentation of an annotated affected bone/joint region in CN radiographs. SinGAN-Seg was used to generate synthetic radiograph–mask pairs from a limited clinical dataset, and six U-Net-based architectures were compared using intersection over union (IoU) on a held-out synthetic test subset. This study evaluates segmentation performance only; it does not assess diagnostic accuracy, disease progression, prognosis, or clinical outcomes.

Materials and Methods

Data acquisition, ethics, and annotation

This retrospective study included 20 plain foot radiographs obtained from four patients with CN treated at a local hospital in Pakistan. The included radiographs represented CN stages 0-2. Ethical approval was obtained from the relevant institutional review body (approval no. 06-2023-ASAB-01/07). All radiographs were anonymized before preprocessing and analysis.

The radiographs were pre-processed using ImageJ and GIMP. Preprocessing included cropping, rescaling to 128 × 128 pixels, and conversion to the required image format. Fourteen original radiographs were selected as source images for SinGAN-Seg-based synthetic-data generation.

Binary segmentation masks were initially prepared by the researcher. Pixels belonging to the annotated affected bone/joint region were labelled 255 (white), while background pixels were labelled 0 (black). The masks were reviewed and validated by an orthopaedic surgeon (S.N.Q.), who is a co-author of this study. Each source radiograph and its corresponding mask were combined into a four-channel input comprising three RGB channels and one binary-mask channel. The distribution of original radiographs by patient sex and radiographic projection is summarized in Table 1.

Table 1: Distribution of original radiographs from patients with CN. The dataset included 20 radiographs from four patients with CN stages 0-2

Patient	Gender	Profile view X-ray	Posteroanterior X-ray	Total
A	Male	2	2	4
B	Female	4	6	10
C	Female	2	2	4
D	Male	1	1	2
Total		9	11	20

Synthetic-data generation

Synthetic radiograph–mask pairs were generated using SinGAN-Seg, a SinGAN-based framework designed to generate synthetic images together with corresponding segmentation masks from limited source data. Fourteen four-channel radiograph-mask inputs were used to train the SinGAN-Seg models. The training process used an image pyramid with scales 1-9, enabling image synthesis at multiple resolutions. Random noise was introduced by the generative process to produce variations while retaining the overall visual structure of the source radiograph and its associated mask.

A neural style-transfer procedure was subsequently applied to the generated images to modify their visual appearance. The resulting synthetic radiographs and their corresponding masks were used to create the segmentation dataset. In total, 560 synthetic radiograph–mask pairs were generated. An overview of the SinGAN-Seg-based synthetic-data generation workflow is shown in Figure 1.

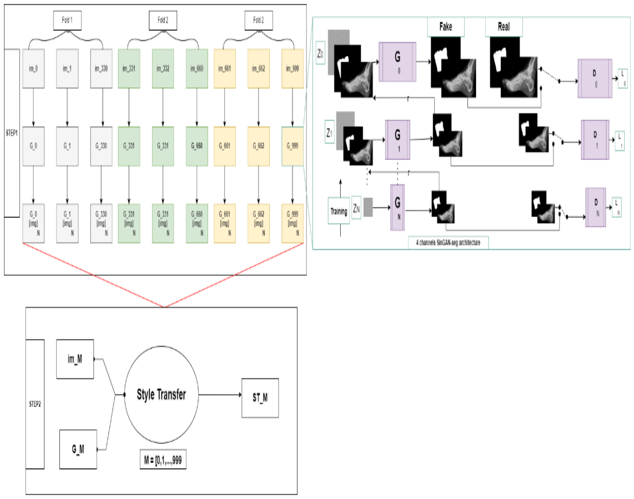

Figure 1: Overview of the SinGAN-Seg-based workflow used to generate synthetic radiograph–mask pairs from source radiographs and corresponding binary masks[12]

Synthetic-image quality assessment

Synthetic-image quality was assessed using single-image Fréchet Inception Distance (SIFID), a single-image adaptation of the Fréchet Inception Distance (FID) framework, and Inception Score (IS). SIFID measures the distance between feature representations of a generated image and its reference image; lower values indicate closer similarity in the evaluated feature space.[13] IS was used as an additional descriptive measure of generated-image feature characteristics.[14] These metrics were used to describe synthetic-image properties and do not establish clinical realism, pathological validity, diagnostic value, or equivalence to original radiographs. The SIFID implementation used in this study follows the FID formulation shown in Equation 1.

$$FID = \left\|\mu_r - \mu_g\right\|_2^2 + Tr\left(\Sigma_r + \Sigma_g - 2\sqrt{\Sigma_r \Sigma_g}\right) \quad (1)$$

The SIFID formulation is based on the Fréchet Inception Distance (FID) metric.[13] For a reference image distribution and a generated image distribution, the feature means are denoted by μ_r and μ_g, respectively, and the corresponding covariance matrices are denoted by Σ_r and Σ_g. Tr denotes the matrix trace. IS is defined as follows[14]:

$$\mathrm{IS}(G) = \exp\left(\mathbb{E}_{\mathbf{x}\sim p_g} D_{KL}(p(y \mid \mathbf{x}) \parallel p(y))\right) \quad (2)$$

where p(y|x) is the predicted class distribution for generated image x, p(y)is the marginal class distribution, D_{KL} is the Kullback–Leibler divergence, and E_x denotes expectation over generated images.

Dataset partitioning and evaluation

The 560 synthetic radiograph–mask pairs were divided at the image level into a training subset of 504 pairs (90%) and a held-out test subset of 56 pairs (10%). The test subset consisted exclusively of synthetic images; original clinical radiographs were not used as the final test set.

Because data partitioning was performed at the image level rather than at the patient or source-radiograph level, synthetic images generated from the same source radiograph could have been represented in both the training and test subsets. The reported results should therefore be interpreted as internal performance on synthetic data and not as independent patient-level or external validation.

Segmentation-model training

Segmentation experiments were performed in Google Colaboratory using a 16 GB NVIDIA Tesla T4 GPU. Six U-Net-based architectures from the Keras U-Net library were evaluated: U-Net with a ResNet152V2 encoder, U-Net++[15] with a VGG19 encoder, attention U-Net[16] with a VGG19 encoder, U-Net with an EfficientNetB7 encoder, recurrent residual U-Net (R2U-Net)[17], and attention U-Net trained from scratch.

All synthetic radiograph–mask pairs were resized to 128 × 128 pixels. Each architecture was trained for a maximum of 50 epochs with early stopping. Batch sizes for Models 1–6 were 4, 14, 5, 10, 5, and 3, respectively. Training stopped after 24, 31, 7, 31, 20, and 19 epochs, respectively.

Models 1-3 used ImageNet-initialized encoders. Models 4–6 were trained from scratch. Leaky ReLU activation was used for Models 1-4 and 6, whereas R2U-Net used ReLU activation. Exact records of the optimizer, learning-rate schedule, loss-function configuration, validation fraction, and early stopping patience were not retained. This limits exact reproducibility and is acknowledged as a study limitation.

Performance metric

Segmentation performance was assessed using intersection over union (IoU), also termed the Jaccard index. IoU was calculated between the predicted segmentation mask Pand the reference mask Gas:

$$IoU = \frac{|P \cap G|}{|P \cup G|} \quad (3)$$

DOI: 10.53389/RJAHS.2026050109

IoU was calculated for each synthetic test image, and mean IoU across the 56-image held-out synthetic test subset was reported for each model.

Results

Twenty original radiographs from four patients with CN stages 0-2 were included. Fourteen original radiographs were selected as source images for SinGAN-Seg-based synthetic-data generation. A total of 560 synthetic radiograph-mask pairs were generated while preserving the principal visual and anatomical characteristics of the source images. These synthetic image–mask pairs were used for the subsequent segmentation experiments. The reported SinGAN-Seg training time was approximately 1 hour and 30 minutes per run in the Google Colaboratory environment.

Synthetic-image quality assessment

Synthetic-image quality was evaluated qualitatively by visual inspection and quantitatively using Single-Image Fréchet Inception Distance (SIFID) and Inception Score (IS). For images generated from source image A, SIFID values ranged from 77.56 to 125.81. For images generated from source image B, values ranged from 139.92 to 196.91. Lower SIFID values indicate closer feature-distribution similarity between a generated image and its respective source image; however, because this assessment was based on a limited number of source images, these values should be interpreted as descriptive rather than as evidence of clinical realism or diagnostic validity. Representative source radiographs, masks, and corresponding SinGAN-Seg-generated image–mask pairs are shown in Figure 2.

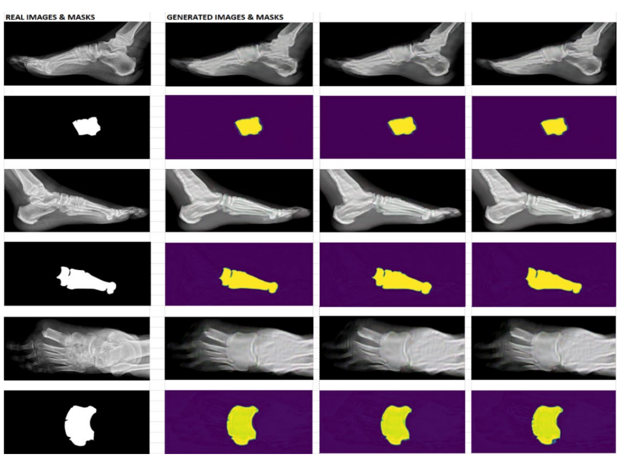

Figure 2: Representative examples of source radiographs and corresponding binary masks used for SinGAN-Seg training, together with generated synthetic radiograph–mask pairs.

Illustrative changes in generated radiographs after the style-transfer step are presented in Figure 3.

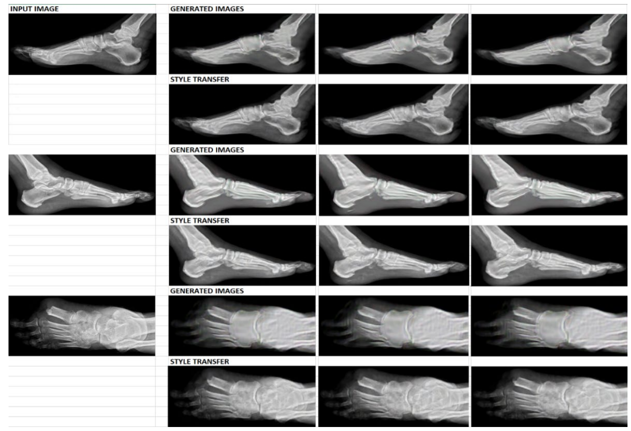

Figure 3: Illustrative examples of synthetic radiographs before and after style transfer.[18] The figure shows visual changes associated with the applied style-transfer setting and does not constitute a clinical-validity assessment.

The SIFID values for ten generated images from each of two source radiographs are presented in Table 2. The Inception Score results are shown in Figure 4 and provide an additional descriptive assessment of the diversity and feature characteristics of generated images. These generative-image metrics do not establish whether synthetic radiographs retain clinically valid pathological features.

Table 2: Single-image Fréchet Inception Distance (SIFID) values for synthetic images generated from two source radiographs.

Generated image	SIFID: Source image A	SIFID: Source image B
1	112.62	196.91
2	97.54	153.09
3	77.56	187.15
4	113.76	161.30
5	101.42	171.66
6	112.44	156.55
7	125.81	160.23
8	100.94	149.08
9	99.08	152.43
10	102.20	139.92
Range	**77.56 - 125.81**	**139.92 - 196.91**

Lower SIFID values indicate closer similarity in the evaluated Inception-feature space. These results are descriptive and do not establish clinical realism or diagnostic validity.

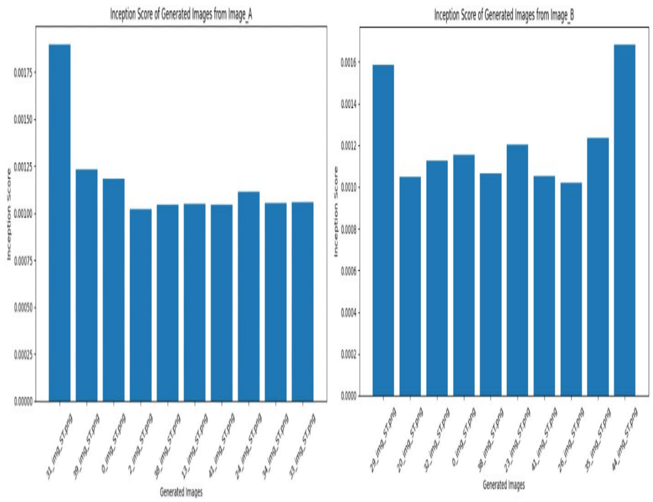

Figure 4: Inception Score values for synthetic images generated from source images A and B. The metric provides a descriptive assessment of generated-image feature characteristics and does not validate clinical realism.

Segmentation-model training

The 560 synthetic image–mask pairs were divided at the image level into 504 training pairs and 56 held-out synthetic test pairs. The models were trained for up to 50 epochs using early stopping. The completed training duration varied among models, ranging from 7 epochs for attention U-Net with an ImageNet-initialized VGG19 backbone to 31 epochs for U-Net++ with VGG19 and U-Net with EfficientNetB7.

Training and validation loss and Dice-coefficient trajectories for the six models are presented in Figure 5. The observed fluctuations, particularly in validation curves, may be attributed to the small synthetic dataset, differences in batch size and model complexity, the image-level partitioning procedure, and variation among generated samples. Therefore, these curves should be interpreted as indicators of optimization behaviour rather than definitive evidence of model generalizability or absence of overfitting.

Comparative segmentation performance

The six U-Net-based models were evaluated on the same 56-image held-out synthetic test set using IoU. Mean dataset IoU ranged from 0.6425 to 0.7572. R2U-Net had the highest observed mean IoU (0.7572), followed by attention U-Net with an ImageNet-initialized VGG19 backbone (0.7275). However, no standard deviations, confidence intervals, per-image distributions, or formal statistical comparison were available. Therefore, these numerical differences should be considered descriptive, and no model can be concluded to be definitively superior. Descriptive IoU performance for the six architectures on the same 56-image held-out synthetic test subset is summarized in Table 3.

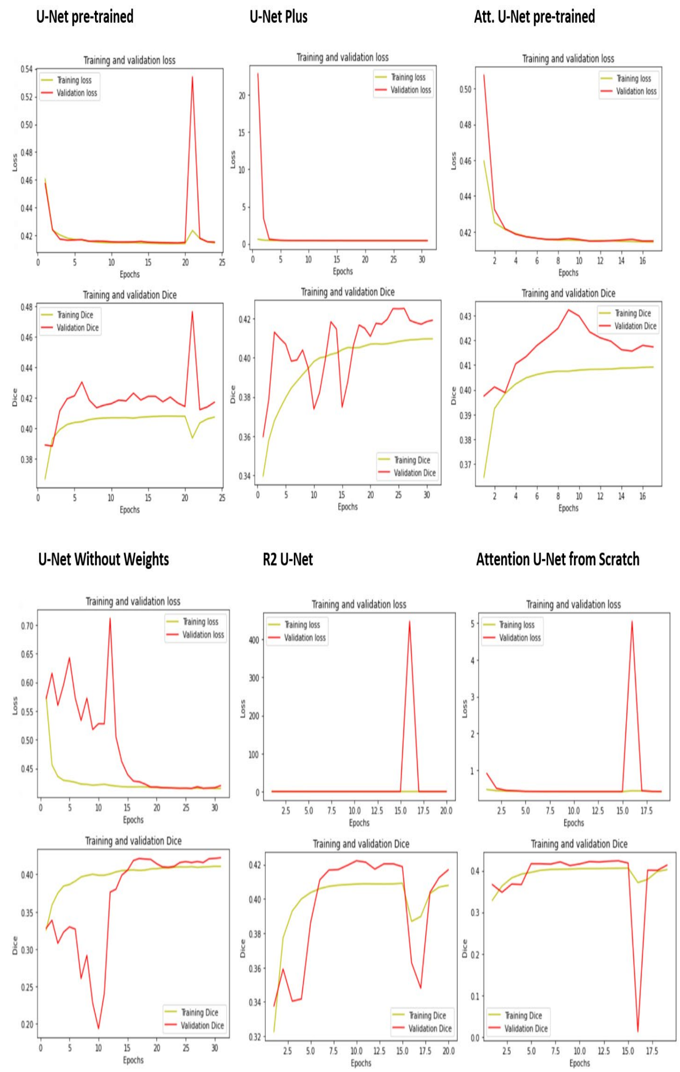

Figure 5: Training and validation loss and Dice-coefficient curves for the six U-Net-based segmentation models across completed epochs. Curve variation reflects optimization behaviour during training and should not be interpreted as evidence of independent generalizability.

The representative prediction examples in Figure 6 demonstrate that all six architectures generated masks with visual overlap with the corresponding reference masks for the illustrated synthetic test image. Nevertheless, these examples are illustrative only. Because the test set was composed of synthetic images and image-level splitting allowed possible source-image dependence between the training and test subsets, the reported IoU values represent internal synthetic-data performance rather than independent patient-level validation.

Table 3: Descriptive segmentation performance of U-Net-based architectures on the held-out synthetic test subset (n = 56)

Model	Encoder / initialization	Batch size	Completed epochs	Example single-image IoU	Mean test IoU
U-Net	ResNet152V2 / ImageNet	4	24	0.8034	0.7179
U-Net++	VGG19 / ImageNet	14	31	0.7473	0.6425
Attention U-Net	VGG19 / ImageNet	5	7	0.8120	0.7275
U-Net	EfficientNetB7 / trained from scratch	10	31	0.8851	0.7177
R2U-Net	No encoder / trained from scratch	5	20	0.8410	0.7572
Attention U-Net	No encoder / trained from scratch	3	19	0.8460	0.7044

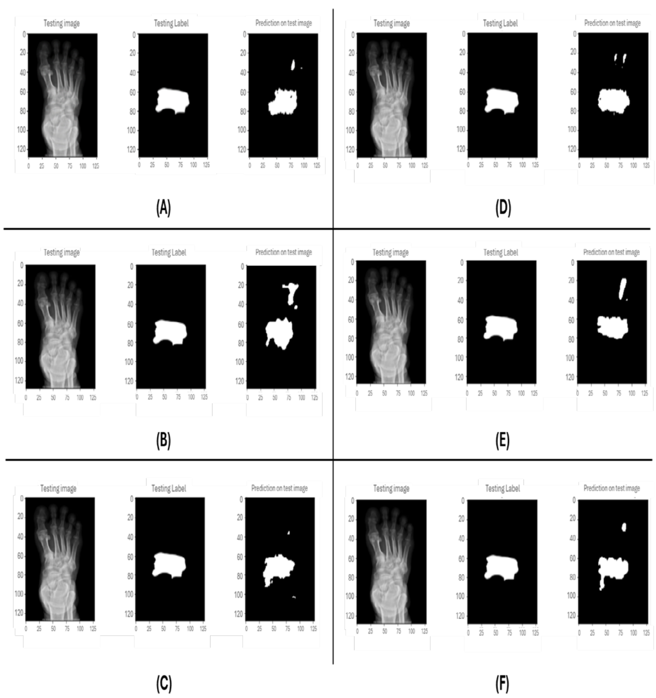

Figure 6: Representative segmentation predictions generated by six U-Net-based architectures for the same synthetic test radiograph: (A) U-Net with ResNet152V2 encoder; (B) U-Net++ with VGG19 encoder; (C) attention U-Net with VGG19 encoder; (D) U-Net with EfficientNetB7 encoder; (E) R2U-Net; and (F) attention U-Net trained from scratch.

Discussion

This preliminary study examined whether SinGAN-Seg-generated radiograph–mask pairs could support training of U-Net-based models for segmentation of an annotated affected bone/joint region in CN radiographs. Across the six evaluated architectures, mean IoU values on the held-out synthetic test subset ranged from 0.6425 to 0.7572. R2U-Net showed the highest observed mean IoU; however, because per-image variability measures, confidence intervals, and formal statistical comparisons were unavailable, the observed differences should be interpreted descriptively.

The study addresses the scarcity of annotated radiographs in CN. Most image-based artificial-intelligence work in the diabetic-foot domain has focused on diabetic foot ulcers, osteomyelitis, thermography, or classification tasks rather than pixel-level segmentation of a Charcot-related region on plain radiographs.[4,6,8] Direct comparison with prior studies remains limited because datasets, image modalities, annotation targets, data partitions, and reported performance metrics differ.

The segmented affected bone/joint region could potentially provide a visual overlay to support closer image review and, if validated on independent clinical radiographs, may assist quantitative monitoring of the annotated region over time. However, the current model does not diagnose CN, distinguish CN from osteomyelitis, determine disease stage, predict disease progression, or replace clinical assessment.

Several limitations should be considered. The original dataset comprised only 20 radiographs from four patients, and 14 source radiographs were used for synthetic-data generation. The held-out test subset included synthetic rather than independent clinical radiographs. In addition, data were partitioned at the image level, so generated images derived from the same source radiograph could have appeared in both the training and test subsets. This dependence may inflate internal performance estimates and prevents assessment of patient-level generalizability. Masks were initially created by one researcher and reviewed by one orthopaedic surgeon; inter-rater agreement was not assessed. Exact records of several training hyperparameters were also unavailable, which limits exact reproducibility. Finally, the study did not include external validation, comparison with human readers, or independent cases of osteomyelitis or diabetic foot ulceration.

Future studies should use larger multicentre datasets with patient-level separation among training, validation, and test sets. Complete experiment logging, independent expert annotations, inter-rater reliability analysis, external validation using original clinical radiographs, and comparisons with clinician performance are required before such a model can be considered for clinical decision support.

Conclusion

This study provides preliminary evidence that SinGAN-Seg-generated radiograph–mask pairs can be used to train and internally evaluate U-Net-based models for segmentation of an annotated affected bone/joint region in CN radiographs. Across six architectures, mean IoU values on a held-out synthetic test subset ranged from 0.6425 to 0.7572, with R2U-Net showing the highest observed value. These findings are limited by the small number of patients, dependence on synthetic data, image-level data partitioning, incomplete hyperparameter traceability, and lack of independent clinical or external validation. Therefore, the results should not be interpreted as evidence of diagnostic accuracy, disease-progression prediction, prognosis, clinical reliability, or generalizability. Larger patient-level and externally validated studies are required before this approach can be considered for clinical decision support.

Generative AI usage Declaration

Generative AI-based language tools were used only for correcting grammatical and typographical errors in the manuscript. All intellectual content, analyses, and conclusions are entirely the work of the authors.

Data Availability Statement

Data is available from the corresponding author on reasonable request.

References

1. Rosskopf AB, Loupatatzis C, Pfirrmann CWA, Böni T, Berli MC. The Charcot foot: a pictorial review. Insights Imaging. 2019 Dec 5;10(1):77. doi:10.1186/s13244-019-0768-9
2. Mascio A, Comisi C, Cinelli V, Pitocco D, Greco T, Maccauro G, et al. Radiological Assessment of Charcot Neuro-Osteoarthropathy in Diabetic Foot: A Narrative Review. Diagnostics. 2025 Mar 19;15(6):767. doi:10.3390/diagnostics15060767
3. Benard T, Lequint C, Jugnet AC, Bouly M, Winther T, Penfornis A, et al. Delayed Diagnosis of Bilateral Neuroarthropathy: Serious Impact on the Development of Charcot's Foot, a Case Report. Medicina (B Aires). 2022 Nov 30;58(12):1763. doi:10.3390/medicina58121763
4. Fahmy AM. Assessing the Efficacy of AI Models in Medical Imaging Analysis for the Early Identification of Diabetic Foot Ulcers and Gangrene: A Review. Premier Journal of Science. 2025 Jul 30. doi:10.70389/PJS.100093
5. Kim J. The pathophysiology of diabetic foot: a narrative review. Journal of Yeungnam Medical Science. 2023 Oct 31;40(4):328–34. doi:10.12701/jyms.2023.00731
6. Cakir M, Tulum G, Cuce F, Yilmaz KB, Aralasmak A, Isik Mİ, et al. Differential Diagnosis of Diabetic Foot Osteomyelitis and Charcot Neuropathic Osteoarthropathy with Deep Learning Methods. Journal of Imaging Informatics in Medicine. 2024 Mar 15;37(5):2454–65. doi:10.1007/s10278-024-01067-0
7. Martín Noguerol T, Luna Alcalá A, Beltrán LS, Gómez Cabrera M, Broncano Cabrero J, Vilanova JC. Advanced MR Imaging Techniques for Differentiation of Neuropathic Arthropathy and Osteomyelitis in the Diabetic Foot. RadioGraphics. 2017 Jul;37(4):1161–80. doi:10.1148/rg.2017160101
8. Álvaro-Afonso FJ, Tardáguila-García A, López-Moral M, Sanz-Corbalán I, García-Morales E, Lázaro-Martínez JL. Using Artificial Intelligence for Detecting Diabetic Foot Osteomyelitis: Validation of Deep Learning Model for Plain Radiograph Interpretation. Applied Sciences. 2025 Aug 1;15(15):8583. doi:10.3390/app15158583
9. Ronneberger O, Fischer P, Brox T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Navab, N., Hornegger, J., Wells, W., Frangi, A. (eds) Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015 [Internet]. Springer, Cham; 2015. Available from: https://doi.org/10.1007/978-3-319-24574-4_28 doi:https://doi.org/10.1007/978-3-319-24574-4_28
10. Renard F, Guedria S, Palma N De, Vuillerme N. Variability and reproducibility in deep learning for medical image segmentation. Sci Rep. 2020 Aug 13;10(1):13724. doi:10.1038/s41598-020-69920-0
11. Willemink MJ, Koszek WA, Hardell C, Wu J, Fleischmann D, Harvey H, et al. Preparing Medical Imaging Data for Machine Learning. Radiology. 2020 Apr;295(1):4–15. doi:10.1148/radiol.2020192224
12. Thambawita V, Salehi P, Sheshkal SA, Hicks SA, Hammer HL, Parasa S, et al. SinGAN-Seg: Synthetic training data generation for medical image segmentation. PLoS One. 2022 May 2;17(5):e0267976. doi:10.1371/journal.pone.0267976
13. Heusel M, Ramsauer H, Unterthiner T, Nessler B, Hochreiter S. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In: Advances in Neural Information Processing Systems. 2017.
14. Salimans T, Goodfellow I, Zaremba W, Cheung V, Radford A, Chen X. Improved Techniques for Training GANs [Internet]. 2016 Jun 10. Available from: http://arxiv.org/abs/1606.03498
15. Zhou Z, Siddiquee MMR, Tajbakhsh N, Liang J. UNet++: A Nested U-Net Architecture for Medical Image Segmentation [Internet]. 2018 Jul 18. Available from: http://arxiv.org/abs/1807.10165
16. Oktay O, Schlemper J, Folgoc L Le, Lee M, Heinrich M, Misawa K, et al. Attention U-Net: Learning Where to Look for the Pancreas [Internet]. 2018 May 20. Available from: http://arxiv.org/abs/1804.03999
17. Alom MZ, Hasan M, Yakopcic C, Taha TM, Asari VK. Recurrent Residual Convolutional Neural Network based on U-Net (R2U-Net) for Medical Image Segmentation [Internet]. 2018. Available from: https://arxiv.org/abs/1802.06955
18. Gatys LA, Ecker AS, Bethge M. Image Style Transfer Using Convolutional Neural Networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2016. p. 2414–23.