This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Puncture Guidance from Non-Contrast X-Ray for Cardiac Pacemaker Implantation

Evangelos Spatharis[1(✉)], Giulio Falasconi[2], Diego Penela[3], Alessia Chiara Latini[3], Andrea Saglietto[4,5], David Soto-Iglesias[2], Antonio Berruezo[2], Oscar Camara[1], and Julia Camps[1]

[1] Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain
evangelos.spatharis01@estudiant.upf.edu
[2] Teknon Medical Centre, Barcelona, Spain
[3] Clinical Institute Humanitas IRCCS, Rozzano, Milan, Italy
[4] Division of Cardiology, Cardiovascular and Thoracic Department, Città della Salute e della Scienza Hospital, Turin, Italy
[5] Department of Medical Sciences, University of Turin, Turin, Italy

Abstract. Obtaining reliable axillary vein access is a critical first step of a successful Cardiac Implantable Electronic Device (CIED) implantation. To support puncture-site selection, clinicians rely on Venography, which visualizes the vein with an iodinated contrast agent. However, Venography X-ray imaging is invasive, nephrotoxic, and not always feasible. In its absence, access is performed blindly using external anatomical landmarks, which carries the risk of pneumothorax, arterial puncture, and brachial plexus injury. In this study, we investigate how much of the puncture-relevant information that Venography provides can be recovered from a non-contrast X-ray alone. A dataset of 95 patient cases was assembled, each with paired non-contrast and contrast sequences, as well as clinician-validated puncture-region boxes. A U-Net was trained to reconstruct the axillary vein from the mean of the non-contrast sequence. In a frozen 30-patient test set, the puncture point derived from the predicted vein fell inside the Venography-informed puncture region in 93.3% of patients, including all 17 patients in which the blind landmark puncture could have missed it, and reduced the mean distance to 3.79 mm, against 6.88 mm for a population atlas prior and 9.44 mm for the blind landmark. These retrospective results suggest that non-contrast X-ray images contain puncture-relevant, patient-specific information to approximate Venography-informed puncture guidance and motivate contrast-free guidance as a direction for prospective investigation.

Keywords: Puncture-site localization · Contrast-free guidance · CIED implantation · Image-guided intervention.

1 Introduction

Cardiac Rhythm Disorders (CRDs), or arrhythmias, are disturbances in the heart's electrical activity that constitute a major global health burden, sig-

nificantly compromising cardiac output and quality of life [1, 2]. When pharmacological therapy alone is insufficient, patients rely on Cardiac Implantable Electronic Devices (CIEDs) such as pacemakers and implantable cardioverter-defibrillators [3]. Millions of CIEDs are implanted every year, and one of the most critical steps of the procedure is obtaining reliable venous access. Following the 2021 ESC guidelines on cardiac pacing [4], clinicians can reach the heart through three main routes, and the axillary vein is often the preferred target due to a higher reported success rate (97.6%) than the cephalic vein (78.2%) [5] and the close proximity of the subclavian vein to the lung and pleura [6].

Access is established with the modern Seldinger technique [7]. However, both the standard [8] and the modern variant are performed blindly: the needle is introduced using surface landmarks, typically just inferior to the clavicle at the junction of its medial and middle thirds, without direct imaging. The complications of an imprecise blind puncture include, but are not limited to, pneumothorax, inadvertent arterial puncture, and brachial plexus injury [9–11]. These risks are amplified by the high anatomical variability between patients, making reliance on static landmarks far from ideal.

To mitigate these dangers, clinicians rely on Venography, an invasive X-ray imaging procedure that injects an iodinated contrast dye to visualize the vein. The optimal puncture site lies under the clavicle and within the medial border of the first rib. Despite its value, Venography is nephrotoxic, can provoke severe allergic reactions, is time-consuming, and requires an additional intravenous line, so patients with renal insufficiency or contrast allergy are often excluded [12]. When it cannot be used, clinicians revert to external landmarks and the risks of blind puncture return. Machine learning has been applied extensively to medical image segmentation [13] with the U-Net [14] reconstructing anatomies reliably across modalities. Generative models have also been used to synthesize contrast-enhanced angiography from non-contrast inputs in order to avoid Iodinated dye [15, 16]. However, no prior work has addressed the localization of the axillary vein from a non-contrast X-ray for puncture guidance, which is the specific gap this study targets.

In this study, we developed a deep-learning workflow to estimate the optimal puncture region from non-contrast X-ray imaging alone. A segmentation model localizes the axillary vein from the non-contrast X-ray, and overlaying this prediction onto the non-contrast image lets a clinician select a candidate puncture region for CIED implantation. We evaluate the predicted puncture region against the Venography-informed reference on a held-out test set and compare it with the blind landmark puncture region and a population atlas built from non-contrast anatomy alone, isolating the patient-specific information the model recovers beyond average anatomy. Our retrospective evaluation indicates that the puncture region derived from our segmentation tool predictions recovers patient-specific puncture information beyond the atlas, agreeing with the Venography-informed puncture region in cases where blind puncture could miss.

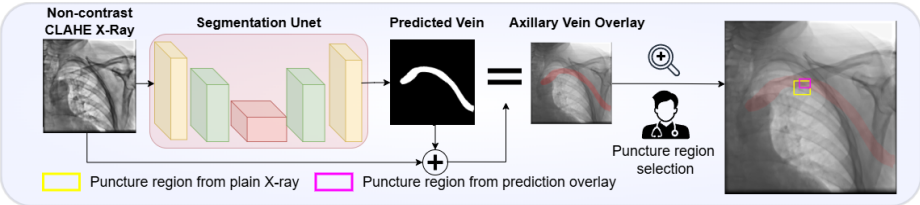


Fig. 1. Overview of the proposed workflow. A non-contrast X-ray image is passed through the segmentation tool (U-Net) to identify the axillary vein; overlaying the prediction on the plain image lets a clinician select a candidate puncture region. The yellow box is the puncture region from the plain X-ray; the magenta box is the puncture region from the prediction overlay.

2 Methods

Our workflow for puncture-site guidance from non-contrast X-ray follows three phases: data preparation, segmentation, and clinical evaluation. We assemble a dataset of 95 patients with paired non-contrast and contrast sequences and clinician-validated puncture regions. Each sequence is registered to remove motion artifacts, and by temporal aggregation we produce a single non-contrast X-ray input image and an axillary vein mask. A U-Net segmentation tool is then trained on the non-contrast images using these masks as targets. Finally, we derive a clinically validated candidate puncture region from the predicted vein and evaluate its agreement with the Venography-informed reference, against both the blind puncture region and the population atlas. Fig. 1 illustrates the segmentation and puncture-region selection steps.

2.1 Venography Dataset

For this study, we assembled a dataset of 95 patient cases scheduled for CIED implantation at two hospitals (Teknon Medical Centre, Barcelona, Spain, and IRCCS Humanitas Research Hospital, Milan, Italy). The study followed the required ethical procedures and all patients signed informed consent. We froze 30 patients (≈31%) as a held-out test set and used the remaining 65 for training and validation. Each case comprises two X-ray temporal sequences of images, with contrast (Venography) and without. Sequences have a different number of frames across patients and between the contrast and non-contrast acquisitions, ranging from 20 to 100 frames per sequence. In addition, for each case, there are two rectangular boxes validated by the clinician that mark plausible puncture regions. The first box is based only on the first rib and the clavicle, where puncture could be possible while avoiding complications. The latter is drawn according to a Venography-informed decision, in the same area but restricted to where the axillary vein is actually present, effectively making it smaller and more specific. We refer to the former as the Blind Region of Puncture (Blind-ROP) and to the latter as the Venography Region of Puncture (Venography-ROP). As

each ROP represents a region of plausible puncture locations, we take the center of the box as a candidate safe puncture point.

2.2 X-ray Image Sequence Preprocessing

Each sequence consists of continuous frames from a single X-ray recording. The contrast sequence was recorded while the dye flowed through the patient's axillary vein, so that different parts of the vein become visible as the sequence progresses. To annotate the vein and to compress each sequence into a single composite image, we must correct two motion artifacts: localized intra-sequence respiratory motion and inter-sequence positional shift. We minimize both with rigid Euclidean transformations optimized by Enhanced Correlation Coefficient (ECC) maximization to a tolerance of $\epsilon = 10^{-6}$ [17]: intra-sequence frames are first registered to a five-frame mean baseline to reduce motion artifacts, after which a single global warp between the temporal medians of the two registered stacks is applied to every non-contrast frame. This achieves sub-pixel translation, ECC > 0.98, and a Structural Similarity Index [18] of 0.90 ± 0.05 at the inter-sequence stage.

Let $\{I^{(t)}_{\text{non-contrast}}\}_{t=1}^{N}$ and $\{I^{(t)}_{\text{contrast}}\}_{t=1}^{M}$ denote the registered non-contrast and contrast sequences. The input of the segmentation tool is CLAHE-enhanced [19] temporal mean of the non-contrast sequence, $\bar{I}_{\text{non-contrast}} = \text{mean}_t(I^{(t)}_{\text{non-contrast}})$. To localize the vein we compute the maximum absolute frame-to-frame temporal derivative of each sequence:

$$\begin{aligned} D_{\text{contrast}} &= \max_t \left| I^{(t+1)}_{\text{contrast}} - I^{(t)}_{\text{contrast}} \right|, \\ D_{\text{non-contrast}} &= \max_t \left| I^{(t+1)}_{\text{non-contrast}} - I^{(t)}_{\text{non-contrast}} \right|. \end{aligned} \tag{1}$$

Then, we subtract the non-contrast map from the contrast map to minimize the motion shared by both, producing an enhancement map as follows:

$$E = \max(0,\ D_{\text{contrast}} - D_{\text{non-contrast}}), \tag{2}$$

which highlights the contrast-enhanced vascular region. E was binarized via hysteresis thresholding, with low and high bounds set to the mean of E plus 0.5 and 0.6 standard deviations respectively, followed by morphological opening/closing. Keeping the largest connected component yielded a binary axillary-vein mask we used as ground truth for the segmentation. Residual noise was manually corrected by us and all masks were verified under the supervision of the clinician. The pipeline is illustrated in Fig. 2.

2.3 Axillary Vein Segmentation

Given the small size of our training set (65 patients), we use a plain U-Net with a ResNet-34 [20] encoder pretrained on ImageNet [21], which exploits transfer learning [22]. The mean vein occupies only 5.84% ($\pm$1.31) of the image, giving a

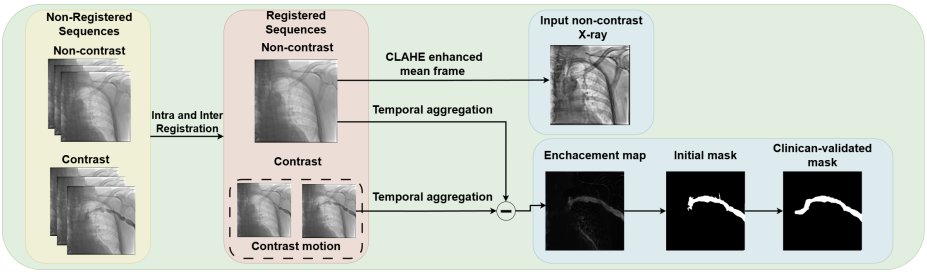


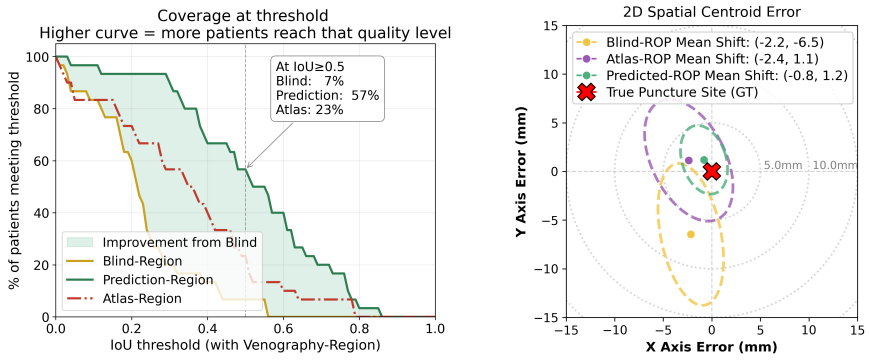
Fig. 2. Illustration of the preprocessing workflow: intra and inter sequence registration, enhancement map generation and the final input and output of the segmentation tool.

heavy class imbalance, which we handle at the loss level with an equally weighted Dice [23] and Binary Cross-Entropy (BCE) loss, $\mathcal{L} = 0.5\,\mathcal{L}_{\text{Dice}} + 0.5\,\mathcal{L}_{\text{BCE}}$. We apply only a $\pm 5\%$ intensity jitter and deliberately omit spatial augmentation, since the patients are consistently positioned during acquisition and the vein masks therefore occupy similar, predictable locations that the network can exploit. The network is optimized with AdamW [24] under a cosine-annealing schedule and trained with 5-fold cross-validation, retaining the best-validation-Dice checkpoint per fold. All test-set predictions are produced by the ensemble of the five retained checkpoints. Their probability maps are averaged and the mean map is binarized at 0.5.

2.4 Evaluation Experiments

Atlas baseline. As a non-learned reference we built a probabilistic vein atlas [25] from the 65 training patients. Each patient is mapped into a canonical 512×384 frame by a similarity transform that aligns the two manually annotated by us clavicle landmarks (medial and lateral) onto fixed positions. The clavicle is used because it both bounds the puncture region and is clearly visible on every plain X-ray. The per-pixel vein frequency peaks at $P = 0.877$ under the clavicle and decays outward, which confirms the conserved anatomy near the puncture site. A binary atlas overlay is produced by thresholding the atlas at $P = 0.3$ and mapping it back into each test patient's image space using only that patient's clavicle landmarks, i.e., using non-contrast information alone. The atlas thus reflects the upper bound achievable from a fixed, population-averaged anatomical prior, isolating the contribution of the learned model beyond average anatomical structure.

Puncture-site agreement. Four boxes are defined per patient under one identical annotation protocol, differing only in their information source: Blind-ROP, Venography-ROP, the Predicted Region of Puncture (Predicted-ROP), and the Atlas Region of Puncture (Atlas-ROP). Predicted-ROP and Atlas-ROP are drawn following the same framework as Venography-ROP, but instead of real



(a) (b)

Fig. 3. Puncture-site localization metrics. (a) Patient coverage by IoU threshold across methods: a higher curve means more patients reach that overlap level. (b) 2D spatial centroid error, with solid dots showing the mean X/Y shift and dashed ellipses one standard deviation of the covariance relative to the true puncture site.

contrast they use the prediction of the model or the atlas overlaid on the plain X-ray. Agreement of a candidate box B with the reference V (Venography-ROP) is their IoU. To express how much of the contrast-informed benefit a candidate recovers relative to the blind baseline, we define

$$\text{Recovery} = \frac{\text{IoU}(B, V) - \text{IoU}(\text{Blind}, V)}{1 - \text{IoU}(\text{Blind}, V)}, \tag{3}$$

where 0 matches the blind guess and 1 reproduces the Venography-informed box exactly. We also report centroid containment, whether the candidate's center lies inside V, and the clinically interpretable physical centroid error, the centroid-to-centroid distance in millimeters obtained from each patient's DICOM pixel spacing. All metrics are computed on the held-out test set ($n = 30$), which played no role in training, atlas construction, or model selection.

3 Results

3.1 Puncture-Site Localization

Anatomical agreement. As reported in Table 1, the Predicted-ROP center fell inside the Venography-ROP in 28 of 30 patients (93.3%), against 13 of 30 (43.3%) for Blind-ROP and 19 of 30 (63.3%) for the atlas. More importantly, among the 17 patients in which the blind center fell *outside* the Venography region, the predicted center fell inside in all 17. The predicted puncture point was correct in the cases where landmark selection could have missed the Venography-informed region.

Table 1. Puncture-region box agreement with the Venography-ROP on the 30 test patients. Best value per row in bold.

Metric	Blind-ROP	Atlas-ROP	Predicted-ROP
IoU with Venography-ROP	0.23	0.34	**0.52**
Recovery	—	0.14	**0.35**
Centroid inside Venography-ROP	43.3%	63.3%	**93.3%**
Centroid error (mm)	9.44	6.88	**3.79**

Physical distance. The predicted center lies on average 3.79 ± 2.37 mm from the Venography-informed center, against 6.88 ± 4.25 mm for the atlas and 9.44 ± 4.53 mm for the blind landmark (Fig. 3b). The reduction in spread is as notable as the reduction in mean: the predicted error ellipse is tightly clustered around the reference, whereas the blind error is both larger and far more variable, particularly along the vertical axis. In 25 of 30 patients the distance to the reference decreased, by 7.4 mm on average. In the remaining 5, where Blind-ROP was closer, it increased by only 3.2 mm.

Overlap and recovery. The Predicted-ROP reaches a mean IoU of 0.52 with the Venography-ROP, more than double the blind baseline (0.23) and well above the atlas (0.34), and it recovers 0.35 of the blind-to-Venography gap on average, more than twice the atlas (0.14). In addition, paired comparison of per-patient ROP IoU confirms the advantage is not driven by average anatomy: the model exceeds the atlas in 23 of 30 patients (77%), a statistically significant difference (paired t-test, $t = 3.50$, $p = 0.0015$). The coverage curves in Fig. 3a show that the predicted box lies above both baselines at every IoU threshold, so the advantage is consistent across the test set rather than carried by a few patients.

3.2 Segmentation Quality

As the puncture-site guidance is derived directly from the predicted segmentation, its accuracy conditions the reliability of the guidance. We measure reconstruction quality by the Dice overlap between the predicted and actual vein masks restricted only to the area above the first rib and an Electrophysiology (EP) region around the overlap of the first rib and the clavicle, namely, around the puncture region. For both regions, the predictions reach a Dice of 0.79. The distal part of the axillary vein is not clinically relevant to CIED implantation. The cases in Fig. 4 show that even if the distal part is not correctly identified, the predictions and the ground truth converge where it is clinically relevant. We further observed that the validation and test Dice scores differed by at most 0.07 in both regions, showing similar performance to unseen patients rather than overfitting on a dataset of this size.

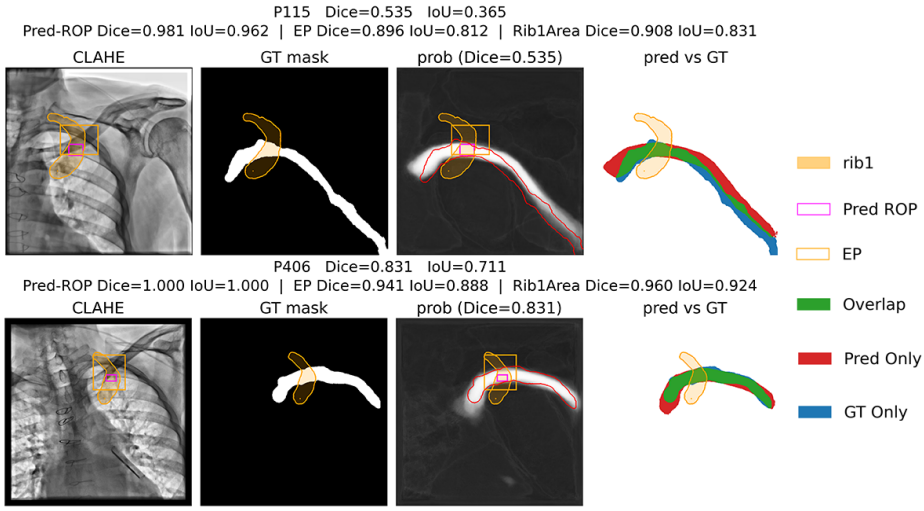


Fig. 4. Representative cases. From left to right: input image, axillary vein mask, non-contrast X-ray image with prediction overlay and vein mask outline, pixel-wise comparison of prediction and vein mask. The orange box is the blind EP-relevant region, the magenta box is the Predicted-ROP, and the orange-filled region is the first rib.

4 Discussion and Conclusion

In this study, our results indicate that a non-contrast X-ray imaging sequence contains puncture-relevant patient-specific information to guide the clinician towards the Venography-informed puncture site. On the clinical evaluation metrics, the predicted vein yields accurate guidance: the point derived from it fell inside the Venography-informed region in 93% of the tested patients against 43% for the blind landmark, at less than half the mean distance and with a tighter, more predictable spread. The most important cases are the ones the blind puncture could have missed. In all 17 patients whose blind landmark puncture site fell outside the Venography-informed region, the predicted site landed inside it, providing guidance precisely where blind puncture was most likely to fail. Notably, this guidance does not require complete vein reconstruction. Recovering a thin vessel in full is intrinsically difficult, but distal error does not affect the selected point, so a partial segmentation suffices. The atlas baseline lets us attribute this improvement to patient-specific image features rather than to anatomy alone. Built from non-contrast landmarks only, the atlas acts as a more informed blind puncture, and the conserved position of the vein makes it a strong prior rather than a weak baseline. Yet, the prediction outperformed the atlas in per-patient overlap and also showed lower mean physical error (3.79 vs 6.88 mm), recovering patient-specific variation that a fixed population prior cannot capture.

A few boundaries frame what this study can claim. Firstly, the evaluation is retrospective. We compare puncture-site annotations on already-collected images rather than guiding a live puncture, so the results are bound to our 30-patient test set. Venography-informed boxes, although verified by more than two clinicians, may vary between experts. The whole-image Dice plateau near 0.60 indicates that the current model and dataset cannot fully reconstruct the distal vein. Finally, the results come from a single trained segmentation network, and even though the X-ray imaging sequences originate from different capturing set-ups, extended adaptability still needs to be confirmed, especially in prospective setting. However, these same boundaries point to a clear path forward. The most immediate is scale: a larger dataset would tighten the localization estimates. Given the paired sequences already in hand, a second direction is to synthesize contrast frames from non-contrast inputs with generative models, which would open the same evaluation to far larger cohorts needing only plain X-rays.

In summary, in a retrospective evaluation, a non-contrast X-ray-based predicted-vein workflow produced puncture-site guidance closer to the Venography-informed reference region than blind landmark-based localization, and closer than a population anatomical prior placed from the same non-contrast information. This supports the feasibility of recovering puncture-relevant information from non-contrast X-ray and identifies contrast-free puncture-site guidance as a direction to be confirmed prospectively.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Antzelevitch, C., Burashnikov, A.: Overview of basic mechanisms of cardiac arrhythmia. Cardiac Electrophysiology Clinics **3**(1), 23–45 (2011)
2. Lippi, G., Sanchis-Gomar, F., Cervellin, G.: Global epidemiology of atrial fibrillation: an increasing epidemic and public health challenge. International Journal of Stroke **16**(2), 217–221 (2021)
3. Massaro, G., et al.: A historical perspective of cardiac implantable electronic device infection: how a menace can drive technological and clinical improvement. Hearts **2**(2), 202–212 (2021)
4. Glikson, M., et al.: 2021 ESC guidelines on cardiac pacing and cardiac resynchronization therapy. European Heart Journal **42**(35), 3427–3520 (2021)
5. Chan, N.-Y., Kwong, N.-P., Cheong, A.-P.: Venous access and long-term pacemaker lead failure: comparing contrast-guided axillary vein puncture with subclavian puncture and cephalic cutdown. EP Europace **19**(7), 1193–1197 (2017)
6. Liu, P., et al.: Optimized axillary vein technique versus subclavian vein technique in cardiovascular implantable electronic device implantation: a randomized controlled study. Chinese Medical Journal **129**(22), 2647–2651 (2016)
7. Higgs, Z., Macafee, D., Braithwaite, B., Maxwell-Armstrong, C.: The Seldinger technique: 50 years on. The Lancet **366**(9494), 1407–1409 (2005)
8. Seldinger, S.I.: Catheter replacement of the needle in percutaneous arteriography: a new technique. Acta Radiologica **39**(5), 368–376 (1953)

9. Kusminsky, R.E.: Complications of central venous catheterization. Journal of the American College of Surgeons **204**(4), 681–696 (2007)
10. McGee, D.C., Gould, M.K.: Preventing complications of central venous catheterization. New England Journal of Medicine **348**(12), 1123–1133 (2003)
11. Mol, T.N., Gupta, A., Narain, U.: Brachial plexus compression due to subclavian artery pseudoaneurysm from internal jugular vein catheterization. Indian Journal of Nephrology **27**(2), 148–150 (2017)
12. Caltekin, İ., Akyol, P.Y., Topal, F.E., Karagöz, A., Ünlüer, E.E.: Contrast-induced nephropathy and allergic reaction in patients who were given intravenous contrast material in emergency department. Journal of Emergencies, Trauma, and Shock **10**(1), 48–49 (2017)
13. Litjens, G., et al.: A survey on deep learning in medical image analysis. Medical Image Analysis **42**, 60–88 (2017)
14. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI), LNCS, vol. 9351, pp. 234–241. Springer (2015)
15. Xu, Q., Rügamer, D., Bender, A., Maros, M.E.: Deep generative models in digital subtraction angiography (DSA) and X-ray angiography: a systematic review. Artificial Intelligence Review **59**(2), 77 (2026)
16. Huang, D.-X., et al.: CAS-GAN for contrast-free angiography synthesis. In: IEEE Symposium on Computational Intelligence in Image, Signal Processing and Synthetic Media (CISM), pp. 1–6. IEEE (2025)
17. Evangelidis, G.D., Psarakis, E.Z.: Parametric image alignment using enhanced correlation coefficient maximization. IEEE Transactions on Pattern Analysis and Machine Intelligence **30**(10), 1858–1865 (2008)
18. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE Transactions on Image Processing **13**(4), 600–612 (2004)
19. Pizer, S.M., Johnston, R.E., Ericksen, J.P., Yankaskas, B.C., Muller, K.E.: Contrast-limited adaptive histogram equalization: speed and effectiveness. In: First Conference on Visualization in Biomedical Computing, pp. 337–345. IEEE (1990)
20. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778 (2016)
21. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 248–255 (2009)
22. Pan, S.J., Yang, Q.: A survey on transfer learning. IEEE Transactions on Knowledge and Data Engineering **22**(10), 1345–1359 (2010)
23. Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Jorge Cardoso, M.: Generalised Dice overlap as a deep learning loss function for highly unbalanced segmentations. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support (DLMIA), LNCS, vol. 10553, pp. 240–248. Springer (2017)
24. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations (ICLR) (2019)
25. Evans, A.C., Kamber, M., Collins, D.L., MacDonald, D.: An MRI-based probabilistic atlas of neuroanatomy. In: Magnetic Resonance Scanning and Epilepsy, NATO ASI Series, vol. 264, pp. 263–274. Plenum Press, New York (1994)