Available online at www.sciencedirect.com

ScienceDirect

Procedia CIRP 139 (2026) 268–273

www.elsevier.com/locate/procedia

13th CIRP Global Web Conference (CIRPe 2025)

6D Pose Estimation of Stator Windings from CT-Measurements

Erik Bär[a], Edwin Blum[a]*, Florian Stamer[a], Gisela Lanza[a]

[a] *wbk Institut für Produktionstechnik, Kaiserstraße 12, 76131 Karlsruhe, Germany*

* Corresponding author. *E-mail address:* edwin.blum@kit.edu

Abstract

In the manufacturing of electric motors, the quality of wire windings is a critical factor determining motor performance and reliability. Defective windings typically result in motor rejection during production, necessitating effective quality control methods. Non-destructive evaluation via computed tomography (CT) offers a promising approach for assessing internal components, although data analysis presents significant challenges. In this study, we introduce a novel method to extract the complete 6D pose of stator windings from CT measurements. The method applies a topology-preserving thinning algorithm to reduce segmented wires to their medial axis, from which positional and curvature data are derived. Local twist about the wire's axis is quantified by analysing image moments in orthogonal cross-sections along the winding. Validation with synthetic CT data and physical motor replicas demonstrates high precision in estimating wire positions and rotations, including intricate bending and twisting behaviours. Future work will focus on optimizing the method to enhance its robustness against the noise inherent in real-world CT scans.

Keywords: 3D image processing; Algorithm; X-Ray; Inspection; Quality Assurance

1. Introduction

Electric motors are fundamental components in a wide range of applications, including household appliances, industrial machinery, and particularly in electric mobility applications. In the European Union, they accounted for an estimated 56% of the EU's total electricity consumption in 2020 [1], making their efficiency and reliability crucial for achieving climate targets and promoting sustainability. A key factor influencing the efficiency and performance of electric motors is the quality of their internal wire windings. Modern quality control methods of electric motors focus primarily on end-of-line functional testing with external load, aiming to assess mechanical and electrical properties and defects [2]. Flat wire windings, favoured for their higher packing density in high-performance motors, introduce additional inspection challenges. For instance, in insertion-based winding processes, conductors must be precisely aligned and lie flat within the stator slots to maximize packing density; any bending or twisting can reduce the number of windings that fit and affect thermal or electrical performance. For better assessment of such structures, external quality control alone is insufficient for a comprehensive assessment. X-ray Computed Tomography offers a promising approach for non-destructive testing of internal structures [3]. CT works by taking multiple X-ray images from different angles around an object and using advanced algorithms to reconstruct a three-dimensional representation of its internal features. This technique allows for high-resolution volumetric data to be obtained without damaging the component being examined. However, the subsequent extraction of relevant geometric parameters, such as the exact position and twisting of individual wires, remains a difficult task.

A range of prior studies have addressed the challenge of reconstructing wirelike structures from volumetric CT data, focusing on the extraction of geometry, orientation estimation or path tracing. One such method, proposed by Emerson et al.

10.1016/j.procir.2025.08.190

[4], applies a supervised method for segmenting and tracing individual fibres in polymer composites using dictionary-based feature recognition as well as centreline tracking. This approach enables segmentation of straight fibres in composite materials but is constrained to low curvature paths with a known general orientation and relies on annotated training data. Motor stator windings, however, vary widely in cross-section, slot shape, curvature, and twist; assembling a sufficiently comprehensive training database would be time- and labour-intensive, limiting practical applicability. In the context of fibre-optic coils, Pillon et al. [5] present a CT-based technique that applies global thresholding and radial slicing of the volume to extract 2D centroids of fibres. These centroids are linked across slices to reconstruct individual fibre paths. However, the method is tailored to cylindrical winding geometries and lacks generality for arbitrary layouts. Tiseanu et al. [6] propose a reconstruction method for superconducting cables based on analysing slices along a predefined straight axis. Within each slice, intensity profiles are fitted using Gaussian functions and thresholding to detect conductor centres, which are then connected using polynomial curves. This method performs well for straight cable sections but is not suited for arbitrary 3D bending. A more domain-specific approach to stator winding inspections is an MATLAB-based algorithm developed by Hoole et al. [7]. They use a hybrid approach of global thresholding for segmentation in less complex areas of the winding while using watershed-segmentation for volumes with curved wire paths. As watershed-segmentation leads to over-segmentation, Hoole et al. [7] implemented a semi-automatic approach to assign the extracted volumes to each other to reconstruct the complete winding. While this hybrid approach allows for complex wiring structures of different geometries, it relies on a trained expert to define where in the CT-scanned volume which approach should be used, as well as correcting the semi-automatic assignment of volumes, making the overall workflow labour-intensive and time-consuming.

While previous CT-based wire reconstruction methods have contributed valuable insights, they typically rely on restrictive geometric assumptions, supervised learning schemes, or user-guided segmentation steps. These methods often focus on either the spatial trajectory or the orientation of wire segments, but not on both in a unified framework. The method introduced in this work addresses these limitations by enabling unsupervised reconstruction of individual conductors under arbitrary 3D bending and twisting. It combines skeletonization, graph-based topology extraction, and moment-based analysis to jointly recover the full 6D pose, comprising translational position, curvature, and torsional twist, of each segment. Although our pipeline requires manual tuning of filtering and thresholding parameters, this calibration need only be performed once for a given set of scan parameters and winding design, dramatically reducing overall labour compared to other existing approaches.

2. Methodology

2.1. Overview of Processing Pipeline

To extract the spatial arrangement and local twist of flat-wire stator windings from volumetric CT data, a custom Python-based processing pipeline was developed. The methodology comprises a modular sequence of image processing and analysis steps tailored to the unique geometrical and material properties of the windings under investigation. The procedure begins with the import and preprocessing of high-resolution CT-scans, followed by a global threshold-based segmentation that isolates the flat-wire winding from the background. A subsequent skeletonization step reduces the segmented volume to a one-voxel-wide medial axis representation, which is then converted into a graph structure to facilitate further processing by allowing efficient analysis through graph-functions and faster cleaning and pruning. Along the enhanced skeleton paths, local cross-sections are extracted orthogonally, enabling a slice-wise shape analysis to determine the positional trajectory and angular orientation of each winding strand.

2.2. CT Data Import

The CT data is provided in multiple files, including a binary volume file and accompanying metadata. The data is structured as a voxel grid, which serves as a three-dimensional representation composed of discrete volumetric elements (voxels). Each voxel is assigned a grey value that primarily corresponds to the density of the material, reflecting the attenuation of X-rays as they pass through different substances within the scanned volume. In the binary file, the voxel grid is stored as a one-dimensional array, which can be reshaped into a three-dimensional array with the associated metadata. Each element in this array represents a voxel's grey value at its respective spatial coordinates. This approach enables fast, memory-efficient access to the data and ensures compatibility with scientific Python libraries of which several are employed in subsequent steps.

2.3. Segmentation

Prior to segmentation, we implemented a Gaussian filter, which is applied to the whole volume, reducing noise and improving the homogeneity of the grayscale values within the wires. This step is particularly effective in attenuating minor surface irregularities or artifacts that may compromise the accuracy of subsequent binarization. A global intensity threshold is then applied to segment the wire material from the background. Given the high contrast and overall quality of the CT scans used in this study, this simple approach yields sufficiently accurate results. The parameters for Gaussian smoothing and thresholding were selected through manual tuning, guided by visual inspection and segmentation performance across representative datasets. Although the pipeline is unsupervised in its execution, these parameters may require adaptation when applied to scans with substantially different voxel intensity profiles or noise characteristics. This need for re-tuning arises only when the scan parameters or the analysed structure change significantly.

2.4. Skeletonization and Graph Conversion

Following binary segmentation, the volumetric winding structures are transformed into their medial axis representations through 3D skeletonization. For this task we utilize a thinning algorithm from Lee et al. [8], which iteratively removes outer voxels while preserving the topology, ultimately yielding 1-voxel-wide skeletal structures representing the medial axis of the original volume. While efficient and well-established, this method is not without limitations. Prior studies have demonstrated its susceptibility to imperfect segmentation and rough surface quality [9]. In this work, similar limitations were observed, resulting in unwanted structures such as spurious branches, clumps, and loops, as illustrated in Figure 2.1.

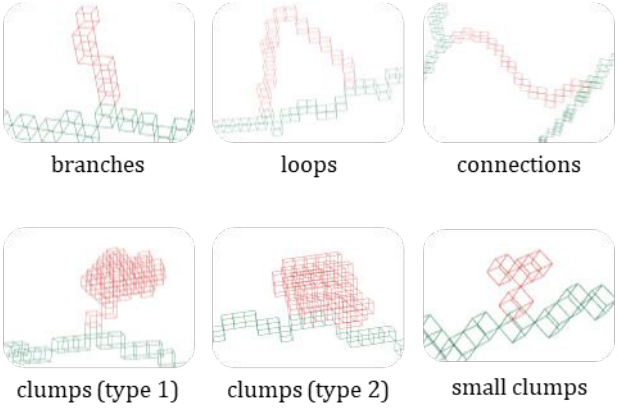

Figure 2.1: Artifacts in generated Skeleton Structure

To enable further analysis and refinement, the voxel-based skeleton is converted into a graph representation. In this graph model, each skeleton voxel is represented as a node, identified and labelled by its 3D coordinates in the Voxel volume. Edges are created between nodes that are 26-neighborhood connected in the voxel space, preserving the topological continuity of the skeleton. This transformation into an unweighted graph enables a more computationally efficient and flexible manipulation of the skeleton structure by leveraging graph-theoretic tools.

2.5. Pruning of the skeleton structure

The resulting graph is used to perform a sequence of customized pruning operations, which have been designed and tuned specifically to remove artifacts that arise in this context while preserving the primary conductor path. This is critical for the subsequent extraction of cross-sectional planes and the accurate determination of local strand orientation and twist. Specifically, the artifact types as shown in Figure 2.1 are targeted and eliminated.

Hair-like branches are non-physical offshoots connected via a junction node with at least three edges and terminating in an endpoint with one edge. These structures are detected by traversing paths starting from all endpoints. If a path reaches a junction or another endpoint before exceeding a predefined maximum length, it is marked as a hair and removed. This prevents the pruning of legitimate wire endings, while removing branches as well as small, disconnected subgraphs.

Loops result from segmentation artifacts where continuous holes appear inside the segmented volume. They manifest in the skeleton as closed cycles bounded by two or more junctions. Loops are detected by computing all cycle bases in the graph and filtered based on a maximum length threshold. For loops bounded by exactly two junctions, the longer of the two paths between the junctions is removed. When three junctions are present, the longest of the paths between each pair is removed to eliminate the unwanted loop, while ensuring that the remaining path maintains connectivity. In cases where junctions are closely spaced and potentially represent a single crossing, their spatial arrangement is considered to merge them virtually before applying the pruning.

Connecting structures between parallel wires, often appearing when the segmentation fails to separate closely adjacent strands, can also form large cycles in the graph structure. These are detected similarly to small loops but are distinguished by their longer path lengths and spatial configuration. In cases with two junctions, the shorter connection between the two junctions is removed, while in case of three junctions, geometric analysis is used to identify which connection to remove. If two junctions form approximately right angles with adjacent branches, the section between them is identified as the connecting bridge and removed.

Large clumps of skeleton nodes form when enclosed voids appear in the segmented volume. They are identified by nodes with high degrees and are classified into two types: clumps connected to the skeleton via a single narrow strand, and those directly embedded into the skeleton with two connection points. In both cases the connection to the overall structure is analysed by searching the surrounding nodes of the high-degree nodes. In case of type 1 clumps, only one connection point to the skeleton structure is found, and all nodes upstream can be removed. In the second case, two connection points are found and only the shortest path between the two is preserved while discarding all other nodes within the clump.

Residual small clumps, often appearing as triangular configurations of three nodes with three edges each, are typically remnants from earlier pruning operations. These are removed by identifying their limited connectivity to the surrounding graph and applying local clean-up steps similar to the removal of larger clumps.

The outcome of these combined pruning operations is illustrated in Figure 2.2, where the removed voxels are highlighted in red, and the skeleton retained for further processing is depicted in green.

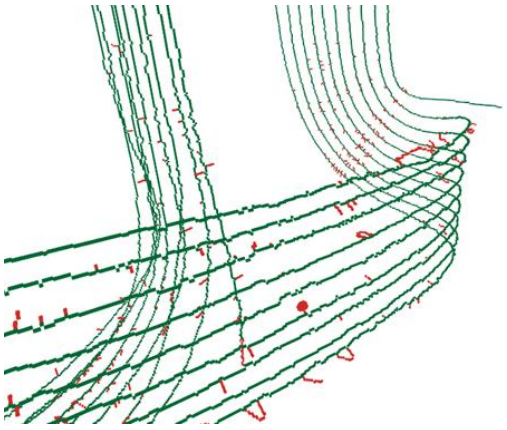

Figure 2.2: Pruning results, showcasing removed voxels in red and the refined skeleton structure in green

2.6. Cross Section Analysis and Twist Detection

Following the topology-preserving skeletonization and specialized pruning of the conductor paths, each continuous wire trajectory is labelled to distinguish between individual windings. Along each wire, discrete positions are sampled at regular intervals. At each sample point, a local tangent vector is calculated based on skeleton points located symmetrically ahead and behind. This tangent describes the local wire direction and is used to compute the 2-dimensional bending as well as a rotation matrix, which aligns the wire segment with the vertical axis. A small volume centred at the sample point is extracted from the original, Gaussian-filtered CT data and rotated to align the wire with the z-axis. From this rotated volume, the central slice, representing a cross-section orthogonal to the wire direction, is selected for analysis. The cross-section is segmented using a threshold, and its contours are identified using an border-following algorithm from Suzuki & Be [10]. For each contour, image moments are computed, and the in-plane twist angle is derived from the normalized second-order central moments. To ensure robustness, only contours with sufficient point count and centre point proximity to the skeleton are considered. The closest valid contour is selected, and its centroid and twist angle are recorded. The resulting dataset contains, for each sampling point, the wire label, 3D positions, tangent vector, local cross-section centre, and computed twist angle, allowing for complete geometric reconstruction of the wire path.

3. Results and Discussion

3.1. Quantitative Evaluation with Synthetic CT Data

To assess the algorithm's accuracy under controlled conditions, synthetic CT datasets were generated using CAD-modelled flat-wire geometries. These included straight, twisted, bent (over both the thin and thick side), and bent-plus-twisted configurations with a cross-section of 4 by 10 millimetres, imported with a resolution of 0.2 mm per voxel and free of noise and imaging artifacts. This setup enabled direct quantitative comparison between the extracted positional and rotational data and the known ground-truth values. The positional deviation refers to the difference between the algorithmically detected skeleton and the CAD-defined wire centreline, while the rotational deviation quantifies the difference between the calculated and actual twist angles at corresponding locations along the wire.

The results (Table 3.1) demonstrate sub-voxel standard positional deviation across all configurations, confirming that the extracted wire centrelines closely match the ground-truth geometries. The rotational deviation remains low for simple geometries but increases significantly in complex cases where curvature and twist are present simultaneously. A notable outlier is the maximum rotational deviation exceeding 33°, which was traced to local failures in the skeletonization process: the derived tangent vector at that point did not align well with the actual curvature of the wire, causing the cross-section analysis to deviate from orthogonality and resulting in inaccurate twist estimation. Furthermore, the computed volume overlap appears relatively low, especially for the shorter straight and twisted wire models. This is attributable to the algorithm's current implementation of tangent estimation, which requires points before and after the analysis position. As a result, twist estimation is undefined near the wire ends, leading to a lack of volume overlap in these regions. This effect disproportionately impacts the short synthetic strands. However, comparisons within groups of wires of equal length indicate that twisting alone does not significantly degrade volume reconstruction fidelity. This is evident from the comparable overlap percentages between the straight and twisted model, both measuring 100 millimetres in length, and likewise among the three bent configurations, each approximately 157 millimetres long. These comparisons are valid because the models within each group share the same physical length and are therefore equally affected by the algorithm's current inability to compute twist at the ends of the wire, resulting in a systematic underestimation of volume overlap near the boundaries.

Table 3.1: Evaluation Metrics for Synthetic CT Wire Models with Varying Geometries

	Straight	Twisted	Curved (thin side)	Curved (thick side)	Curved and twisted
Standard positional deviation [Voxel]	0,000	0,641	0,215	0,225	0,446
Max. positional deviation [Voxel]	0,000	2,236	0,823	0,902	2,264
Standard rotational deviation [degrees]	0,000	0.096	0,000	1,231	6,343
Max. rotational deviation [degrees]	0,000	0.539	0,000	4,614	33,413
Computed volume overlap [%]	88,0	85,9	92,1	91,2	89,7

3.2. Qualitative evaluation on Real CT Data

Real-world applicability was examined by applying the algorithm to CT data from a high energy-dense motor replica. The use of a replica with an additively manufactured polymer housing was chosen to ensure higher scan quality and avoid artifacts, such as beam hardening, excessive noise and artifacts. The evaluation specifically examined the sensitivity of the skeletonization and segmentation pipeline to various pre-processing parameters, particularly the Gaussian smoothing filter and global threshold values. We manually tuned these parameters to establish default settings, which served as a baseline for our analysis. The results demonstrated that low filter values preserved detail but also amplified noise-related artifacts, such as branches and clumps. Conversely, high filter values effectively suppressed noise but often led to the merging of adjacent wire segments or the loss of finer structures. While increasing the threshold values could mitigate merging between wire strands, excessively high thresholds introduced new issues, such as holes in the segmented wire, resulting in clumps and discontinuities in the skeleton.

To illustrate the sensitivity of the cleanup process, we varied the Gaussian filter by ±0.2 around the baseline σ=0.8 and adjusted the threshold by ±2000 around T=45000. Under this

variation, the number of artifacts detected and removed changed significantly, ranging from 1670 to 3648, compared to 2097 at the default settings. The differences in artifact counts can be attributed to the interplay between noise suppression and detail preservation, highlighting how these parameter choices can significantly impact the outcome. Once the parameters are fine-tuned for a specific set of scan conditions, they could be consistently applied to similar scans, streamlining future analyses.

Using the default settings established through systematic testing, we removed 5,109 of the 105,813 voxels associated with the 2,097 detected artifacts, resulting in no leftover unwanted structures. While these totals will vary with scan settings and the object under study, they demonstrate that our pruning pipeline can be effectively tuned to eliminate all structures that would otherwise hinder further analysis. Using this setup, the algorithm successfully extracted the wire positions and rotations from the segmented volume with high reliability across most of the structure. The resulting analysis, depicted in Figure 3.1, clearly visualizes the wire's three-dimensional skeleton along with vectors indicating the direction and magnitude of twist computed from the cross-sectional data. This comprehensive pose extraction illustrates the algorithm's capability to capture detailed geometric and rotational characteristics directly from real CT scans. The entire processing pipeline, from loading the 651 × 1084 × 1693 voxel volume to producing these outputs, was completed in approximately six minutes on a mid-range computer (Ryzen 5 2600, 16 GB RAM).

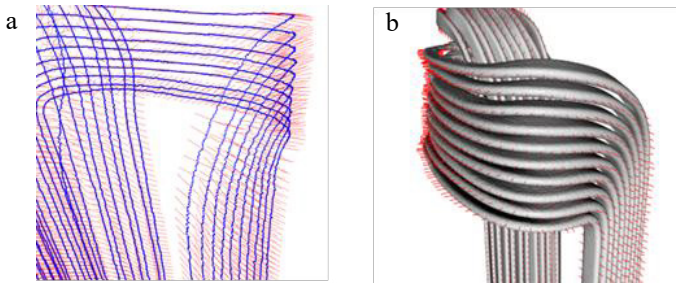

Figure 3.1: (a) Extracted and pruned skeleton (in blue) with local rotation vectors (in red); (b) Mesh of the extracted flat-wire skeleton from real CT data overlaid with rotational vectors

4. Conclusion

This work presents a novel algorithm for the automated analysis of flat-wire windings in electric motors using CT data. The algorithm is capable of extracting the complete 6D Pose of flat-wire stator windings, providing detailed geometric information which is not accessible with conventional CT analysis tools or standard motor testing methods. The algorithm's strength lies in its flexibility of the analysed winding structure and its high degree of automation. It enables detailed inspection of the winding structure, which is particularly relevant for quality control in high-energy-density electric motors using flat-wire coils. In these applications, precise detection of positioning and twist is crucial for ensuring efficient winding and reliable long-term operation.

Quantitative and qualitative evaluations demonstrate that the algorithm performs robustly across a range of wire geometries. While deviations increase slightly for more complex shapes, such as curved and twisted wires, the overall accuracy remains acceptable and consistent. The most critical factor influencing performance is the quality of segmentation, which strongly depends on parameter settings such as Gaussian filtering and thresholding. Currently, the algorithms parameters must be fine-tuned manually, making the process sensitive to image noise and variability in CT scan quality.

Future improvements should focus on enhancing robustness against poor or noisy segmentation. Incorporating advanced thresholding strategies, automatic parameter tuning, and exploring machine learning techniques for improved segmentation would be viable steps toward achieving a fully automated and generalizable pipeline. These developments would not only enhance the algorithm's applicability in more variable environments but also pave the way for its integration into automated production workflows for in-line quality assessment.

Acknowledgement

The authors gratefully acknowledge financial funding from the Ministry of Economic Affairs, Labor and Tourism of Baden-Württemberg (grant no. BW1_1478/02).

References

[1] L. Wierda and A. Zanuttini, 'Ecodesign Impact Accounting Overview Report 2024', European Commission, Directorate General for Energy, Delft, The Netherlands, Technical Report Specific Contract No 2021/OP/0004/ENER/B3/FWC 2020-708/LOT 1/04/FV2022-531, Oct. 2024. [Online]. Available: https://www.vhk.nl/eia.html

[2] A. Kampker, Elektromobilproduktion. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014. doi: 10.1007/978-3-642-42022-1.

[3] S. Carmignato, W. Dewulf, and R. Leach, Eds., Industrial X-Ray Computed Tomography. Cham: Springer International Publishing, 2018. doi: 10.1007/978-3-319-59573-3.

[4] M. J. Emerson, K. M. Jespersen, A. B. Dahl, K. Conradsen, and L. P. Mikkelsen, 'Individual fibre segmentation from 3D X-ray computed tomography for characterising the fibre orientation in unidirectional composite materials', Composites Part A: Applied Science and Manufacturing, vol. 97, pp. 83–92, Jun. 2017, doi: 10.1016/j.compositesa.2016.12.028.

[5] J. Pillon et al., 'Three-Dimensional Topological Reconstruction of the Sensing Coil of a Fiber-Optic Gyroscope Using X-Ray Computed Tomography', J. Lightwave Technol., vol. 39, no. 14, pp. 4861–4872, Jul. 2021, doi: 10.1109/JLT.2021.3068605.

[6] I. Tiseanu et al., 'Multi-scale 3D modelling of a DEMO prototype cable from strand to full-size conductor based on X-ray tomography and image analysis', Fusion Engineering and Design, vol. 146, pp. 568–573, Sep. 2019, doi: 10.1016/j.fusengdes.2019.01.025.

[7] J. Hoole, D. North, N. Simpson, and P. H. Mellor, 'Verifying Strand Transposition in Stator Windings via X-ray Computed Tomography derived Three-Dimensional Models', in 2023 IEEE Workshop on Electrical Machines Design, Control and Diagnosis (WEMDCD), Newcastle upon Tyne, United Kingdom: IEEE, Apr. 2023, pp. 1–6. doi: 10.1109/WEMDCD55819.2023.10110906.

[8] T. C. Lee, R. L. Kashyap, and C. N. Chu, 'Building Skeleton Models via 3-D Medial Surface Axis Thinning Algorithms', CVGIP: Graphical Models and Image Processing, vol. 56, no. 6, pp. 462–478, Nov. 1994, doi: 10.1006/cgip.1994.1042.

[9] M. G. Forero and C. A. M. Rulz, A Note on 'Building Skeleton Models via 3-D Medial Surface/Axis Thinning Algorithms'. in SYMPOSIUM ON SIGNAL PROCESSING, IMAGES AND COMPUTER VISION (STSIVA), no. 20. 2015. Accessed: Aug. 10, 2024. [Online]. Available: https://www.researchgate.net/publication/304251296

[10] S. Suzuki and K. Be, ‘Topological structural analysis of digitized binary images by border following’, Computer Vision, Graphics, and Image Processing, vol. 30, no. 1, pp. 32–46, Apr. 1985, doi: 10.1016/0734-189X(85)90016-7.