Available online at www.sciencedirect.com

ScienceDirect

Procedia CIRP 00 (2026) 000–000

www.elsevier.com/locate/procedia

19th CIRP Conference on Computer-Aided Tolerancing (CAT 2026)

Fast Industrial Computed Tomography for Remanufacturing: Evaluating Measurement Time savings with 3D UNETR Sparse-view Segmentation

Edwin Blum[a,*], Tom Hild[a], Dominik Koch[a], Martin Benfer[a], Gisela Lanza[a]

[a] *wbk Institute of Production Science, Karlsruhe Institute of Technology (KIT), Kaiserstraße 12, 76131 Karlsruhe, Germany*

*Corresponding author. Tel.: +49 721 608-44011. *E-mail address:* edwin.blum@kit.edu

---

Abstract

Industrial computed tomography (CT) provides non-destructive insight into internal structures and full 3D geometry of complex assemblies, but its use in production is limited by long acquisition times. Sparse-view CT can reduce measurement time, yet introduces artefacts and geometric bias that may compromise metrology. We evaluate sparse-view, machine-learning-based component extraction on an assembled angle grinder containing a bevel gear, using five gear variants with wear and damage. A 3D UNETR model is trained for signed distance field (SDF) regression to enable robust surface extraction. Across five leave-one-gear-out runs, we assess segmentation robustness under reduced acquisition time and translate the results into metrology-relevant outcomes via surface deviations and a tolerance-oriented tooth-profile evaluation. Results show that acquisition time can be substantially reduced within the trained regime while maintaining high segmentation quality and largely stable feature-level deviations. For stronger reductions, errors increase non-linearly and become spatially localized, dominating tail statistics and tolerance failures. Beyond demonstrating time savings, the proposed tolerance-centered evaluation links sparse-view segmentation performance to feature-level deviation distributions and pass rates, providing a transferable basis for selecting acquisition-time settings and motivating adaptive inspection strategies that trade measurement time against decision reliability.

Keywords: Industrial computed tomography; sparse-view segmentation; UNETR; remanufacturing

---

1. Introduction

Industrial computed tomography (CT) provides non-destructive insight into internal information (e.g., defects) and full 3D geometry in a single measurement, making it attractive for the quality assurance of complex products and assemblies. In a CT scan, an object is rotated between an X-ray source and a detector while a series of radiographs (projections) is acquired at different angles; a 3D volume is then reconstructed from these projections. However, broader deployment in production is still constrained by acquisition time and computational effort, which limit throughput and scalability [1]. These limitations become particularly critical in circular production, where value creation relies on processing used products rather than homogeneous new-part streams. In remanufacturing, used products or components are restored to a defined functional and quality level so they can be returned to service [2]. A key challenge is that incoming products exhibit uncertain, instance-specific conditions and high variant diversity, which turns inspection into a decision-making bottleneck: components must be classified (e.g., reuse, refurbish, replace, scrap), downstream operations must be planned, and disassembly effort must be scheduled efficiently [3]. In this setting, CT is valuable not only for dimensional verification but also for *pre-disassembly diagnostics*, because it can capture internal wear and damage before dismantling and thereby support targeted routing and processing decisions.

Since real products typically contain multiple parts in close proximity, extracting the component of interest from the reconstructed volume is a prerequisite for building an instance-specific component model. Segmentation is therefore the key enabler: it isolates individual components and provides the basis for surface determination and feature extraction, which in turn underpin defect assessment and dimensional measurements.

Achieving this at scale requires shorter acquisitions. A direct lever is reducing the number of projection views (sparse-view CT), which lowers acquisition time approximately in proportion to the view count. However, fewer views make reconstruction more ill-posed and can introduce structured artefacts as well as geometric bias that propagate to segmentation and translate into measurement deviations on function-critical features [4, 5]. This motivates an *"as accurate as needed"* perspective: rather than maximizing overlap metrics, the goal is to minimize acquisition effort while maintaining reliable, tolerance-relevant decisions [3].

We study this question in a representative use case: an assembled angle grinder containing a bevel gear. For circular-use inspection, the gear must be extracted from the surrounding housing and neighboring components to derive an instance-specific component model and assess it under varying wear and damage states. This reflects remanufacturing scenarios in which products arrive as assemblies and diagnostic decisions are required *before* disassembly to plan efficient downstream processing. Accordingly, we investigate *instance-level component extraction under acquisition constraints*: we segment the gear from sparse-view industrial CT using a 3D *"UNEt TRansformers"* (UNETR) [6] backbone with a signed distance field (SDF) regression target, i.e., the network predicts a continuous signed distance to the component surface per voxel, whose zero level set defines the extracted geometry. Additionally, we evaluate whether the resulting surface remains metrologically usable as acquisition time is reduced. Tooth-profile measurements serve as a representative example of a function-critical feature.

The paper contributes (i) a sparse-view industrial CT workflow for instance-level component extraction from an assembled product using SDF-based UNETR segmentation, (ii) a metrology-centered evaluation that links acquisition-time reduction to surface deviations and tolerance-oriented feature outcomes rather than overlap metrics alone, and (iii) an empirical study on five gear instances showing that substantial time reduction is feasible within the trained regime, whereas stronger reductions lead to localized deviation hotspots that dominate tail errors and tolerance failures.

2. Related Work

Sparse-view CT has been studied extensively as a direct means to reduce acquisition time. In cone-beam CT, analytic filtered backprojection methods such as the Feldkamp - Davis - Kress algorithm remain a widely used baseline, but sparse angular sampling makes reconstruction increasingly ill-posed and typically introduces structured streak artefacts [1]. Evaluations are therefore commonly reported in terms of reconstruction fidelity, such as visual sharpness, noise/artefact level, or reconstruction error with respect to a reference volume [7, 8]. For production inspection, however, the decisive question is different: whether reduced-view acquisition still supports downstream tasks such as reliable tolerance-relevant measurement.

In industrial CT, segmentation has traditionally relied on rule-based procedures (thresholding, edge- and region-based methods), which can be sufficiently precise in controlled settings but often break down in the presence of noise, artefacts, and limited contrast, especially in multi-material assemblies [1]. Recent industrial CT work increasingly adopts supervised deep learning with voxel-wise classification, frequently targeting internal structures (e.g., pores) rather than extracting full components from complex assemblies [9]. Instance-level component extraction in industrial CT assemblies remains comparatively underrepresented [10, 11, 12].

Related challenges are well studied in medical CT, where instance-level segmentation in crowded 3D scenes is more established. In this field, architectures such as UNETR provide strong context-aware baselines by combining transformer encoders with CNN decoders [6]. In addition, medical segmentation research increasingly explores SDF as regression targets. Instead of predicting binary voxel labels, SDF formulations regress a continuous signed distance to the object surface (Fig. 1d), which typically yields smoother and metrically more consistent surfaces than voxel masks and is attractive when downstream objectives rely on surface-based analysis [13, 14, 15]. These developments are particularly relevant for industrial inspection workflows in which the desired output is not primarily a label map, but a geometrically usable component surface for subsequent measurement. The trade-off between acquisition speed and dimensional accuracy in CT metrology is well documented: reducing the number of projections directly affects reconstruction quality and measurement deviations [4, 5]. Recent work has shown that the degree to which an ML model generalizes across different scan conditions directly affects the dimensional accuracy of ML-enhanced CT reconstructions [16]. While such approaches apply ML at the reconstruction level and subsequently extract surfaces using conventional methods, the present work operates at the segmentation level: we train a model to extract a metrologically usable component surface directly from sparse-view input and evaluate the result through feature-level deviations and tolerance compliance.

Building on these foundations, we focus on *instance-level component extraction from an assembled industrial product under reduced acquisition time*. We assess whether the resulting component surface remains fit for metrology by linking segmentation output to surface deviations and tolerance-relevant outcomes in a measurement-centric evaluation chain.

3. Methodology

3.1. CT data acquisition and sparse-view protocol

For our study, we use an assembled angle grinder containing a bevel gear and acquire industrial CT volumes on a metrology-grade CT system (Zeiss Metrotom 1500). To reflect the variability expected for circular-use products, we employ five gear variants with synthetically generated wear patterns, ranging from intact teeth to abrasive wear/pitting and severe tooth-break fatigue. Defect magnitudes range from wear features on the order of a few μm up to approximately 5 mm in the case of tooth breakage.

All assembly scans are acquired at 220 kV, 2200 μA, 250 ms integration time with a 1 mm Sn filter (voxel size $\approx$ 111 μm). Scans are performed with identical fixture, scan geometry, and

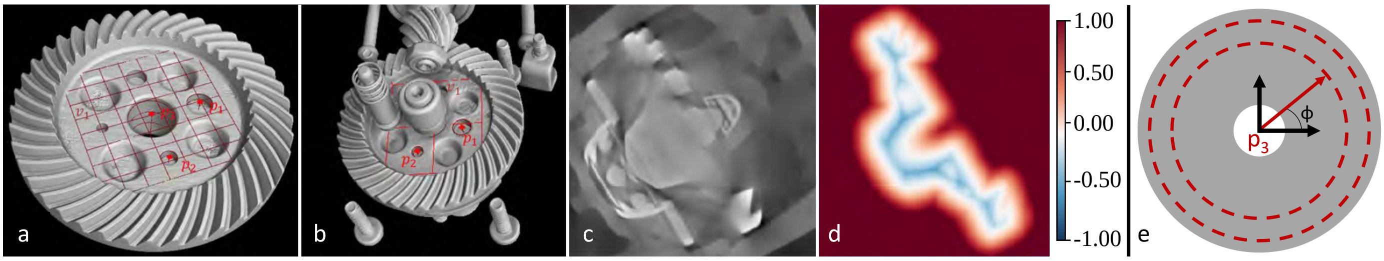

Fig. 1. **Workflow and coordinate definitions for component extraction and feature evaluation.** (a) Reference gear geometry with landmark points p_1 and p_2 used for reproducible alignment. (b) Same landmarks identified in the full assembly scan, enabling transfer of the reference and definition of a part coordinate system. (c) Example CT region of interest (ROI) of the gear as input to the learning pipeline. (d) SDF representation of the target surface within the ROI (red/blue indicate positive/negative distance; the zero level set corresponds to the surface). (e) Schematic of the polar feature-evaluation setup with origin at p_3, angular coordinate ϕ, and a radial measurement band (dashed) used for rotationally consistent sampling of tooth-profile features.

reconstruction settings; only the number of projection views (PZ) is varied to control sparse-view artefacts. We reconstruct a projection ladder from 3000 to 60 PZ (14 settings in total). Throughout the paper, we report each setting as *relative acquisition time* $\tau \in (0, 1]$, normalized to the baseline setting at 3000 PZ ($\tau = 1$). For reference, $\tau = 1$ corresponds to 12.5 minutes acquisition time in our setup.

For each wear variant, we additionally acquire a high-quality assembly scan (same scan parameters as above, with the machine's noise-suppression mode enabled) and a high-quality single-component reference scan (220 kV, 106 μA, 2000 ms, 0.75 mm Cu filter, voxel size $\approx$ 22 μm). The high-quality assembly scan allows robust identification of landmark points p_1 and p_2 (Fig. 1b), and the single-component reference scan provides a low-noise reference geometry for annotation and quantitative evaluation. Since the part remains clamped and its pose is kept constant throughout the projection ladder, the resulting rigid alignment can be propagated to all reduced-PZ reconstructions. In total, we obtain 70 assembly CT volumes (5 variants × 14 projection settings).

3.2. Reference geometry, preprocessing, and registration

For each gear, the high-quality single-component reference scan is used to obtain a low-noise volume suitable for robust surface extraction. We derive a reference surface via threshold-based segmentation and export the resulting mesh as an STL file to serve as reference geometry (Fig. 1a).

To transfer this reference geometry to each assembly reconstruction, we perform marker-based rigid registration and estimate 4×4 transformations. Specifically, the landmark points p_1 and p_2 (Fig. 1a–b) define a reproducible alignment between the reference STL and the product-scan coordinate system via a transformation $T_{\mathrm{ref}\to\mathrm{prod}}$. In addition, we define a part coordinate system (PCS) centered at p_3 to ensure reproducible downstream evaluation across reference and predictions (Fig. 1e).

In practice, reliable landmark identification is feasible down to 900 PZ. For settings below 900 PZ, sparse-view artefacts increasingly corrupt the local appearance such that mapping the reference geometry into these reconstructions yields visibly distorted SDF targets. We therefore use all settings with PZ $\geq$ 900 as the *trained range* (i.e., included during training and used for the main quantitative analysis), while settings with PZ $<$ 900 are treated as an *extreme sparse-view regime* beyond the training range and are reported separately as stress tests.

3.3. Learning problem and training data generation

We extract the gear from the assembly CT scan by predicting a *metrically usable surface*. We formulate component extraction as SDF regression: each voxel stores the signed distance to the nearest point on the gear surface (negative inside, positive outside), and the surface is given by the zero level set (SDF= 0) (Fig. 1d). Compared to binary voxel masks, SDF regression provides a continuous target and supports stable surface extraction. Voxel-based SDF ground truth is computed from the aligned reference STL via point-to-triangle distances with an inside/outside sign convention. Training samples are generated by extracting a fixed-size region of interest (ROI) around the gear (Fig. 1c): an oriented bounding box (OBB) is placed around the aligned reference geometry and resampled into an axis-aligned cube using trilinear interpolation, yielding 224^3 voxel inputs at an isotropic voxel size of $\approx$ 253 μm. Random OBB rotations increase sample diversity while preserving the fixed physical pose of the part in the CT measurements. CT intensities are z-score normalized; SDF values are clipped to $\pm$5 mm and scaled to $[-1, 1]$.

3.4. Network and training

We use a 3D UNETR backbone [6] to predict an SDF volume at the same resolution as the input ROI. All experiments are trained on a single NVIDIA RTX A6000 GPU. Since geometric accuracy is tied to voxel spacing, we choose a high-resolution ROI of 224^3 voxels, which is the largest volume that fits into GPU memory for training. This constraint results in a batch size of 1; we therefore use gradient accumulation over 10 iterations to stabilize optimization.

Because only a small number of fully annotated industrial CT volumes is available, we employ a compact UNETR configuration to improve generalization in the low-data regime. Concretely, we use an embedding size of 384, 6 attention heads, an MLP dimension of 1536, and a decoder feature size of 32. Training uses AdamW with a warmup cosine-annealing schedule over 250 epochs and a peak learning rate of 1×10^{-4}. Robustness across instances is assessed with five leave-one-gear-out runs: one physical gear is held out for validation/test and the remaining four gears are used for training.

Augmentations are applied consistently to CT and SDF for geometric transforms, while intensity-only transforms are applied to CT only. Optimization is driven by a surface-aware loss: for downstream measurement we ultimately extract the zero level set (SDF= 0), hence errors close to the surface translate most directly into geometric surface deviations, whereas errors far away from the surface are less critical and mainly act as a shape regularizer. We therefore minimize a weighted L_1 error $\mathcal{L}(\hat{d}, d) = w(d)\,|\hat{d} - d|$ with $w(d) = w_s$ for $|d| \leq s$ and $w(d) = w_i$ otherwise. Inside a narrow band $|d| \leq s$ around the surface we up-weight the loss ($w_s = 3$) to prioritize accurate localization of the zero crossing, while outside the band we use a lower weight ($w_i = 1$) to keep the SDF globally consistent. We set $s = 0.2$ in the normalized SDF scale, corresponding to an emphasis on the near-surface region most relevant for later surface extraction and metrology.

3.5. Surface extraction and feature evaluation

We extract an explicit surface from the predicted SDF using marching cubes at SDF= 0 (Fig. 1d), yielding a polygonal mesh for evaluation. We report Dice (after binarizing at SDF= 0) and HD95 between predicted and reference surfaces to capture overlap and tail-sensitive boundary errors.

For metrology-centered evaluation, we analyze (i) global surface-to-surface deviations and (ii) a local, tolerance-oriented tooth-profile measurement by *virtual probing*. For the global analysis and the heatmaps, we compute point-to-point deviations by associating each vertex of the predicted mesh with its closest point on the reference surface and storing the resulting Euclidean distance; these distances are then mapped onto the reference mesh for visualization. A nominal measurement pose T_{target} is defined in the PCS, and probe rays are generated by sampling the angular coordinate ϕ within a predefined radial band around p_3 (Fig. 1e). For each ray, the surface point is determined robustly by searching within a cylindrical tolerance region of radius R_{tol} around the ray. The resulting profile deviations are summarized using distribution statistics (e.g., p50, p95) and tolerance pass rates $P(\text{dev} \leq T)$ for representative tolerances T.

4. Results

4.1. Segmentation robustness under measurement-time reduction

Table 1 summarizes Dice and HD95 over five cross-validation runs for acquisition-time fractions down to $\tau = 0.3$, normalized to the baseline setting $\tau = 1$. Across this range, segmentation performance remains consistently high and reproducible (mean Dice 0.9363 - 0.9564; HD95 means 1.414 - 2.129 mm). Since each run holds out one physical gear instance (with its wear/damage state) from training, these results indicate that the model generalizes across previously unseen, instance-specific defects. This is a key requirement for remanufacturing scenarios with high variability across incoming products. Moreover, the cross-validation design inherently quantifies repeatability of the extraction pipeline: across all five runs and acquisition-time settings within the trained range, Dice standard deviations remain below 0.009 and HD95 values stay within 1.4–2.4 mm (Table 1), indicating that the pipeline produces consistent results across independently trained models evaluated on unseen gear instances with different wear and damage states.

Table 1. Run-wise segmentation performance in the in-distribution regime (baseline time fraction $\tau = 1.0$ down to $\tau = 0.3$). Dice is a dimensionless metric (0–1) reporting volumetric overlap (higher is better) and HD95 reports boundary error in mm (lower is better). Values are given as mean ± std and min - max.

Run	Dice (mean ± std)	Dice (min–max)	HD95 [mm] (mean ± std)	HD95 [mm] (min–max)
1	0.9564 ± 0.0036	0.9509–0.9614	1.414 ± 0.000	1.414–1.414
2	0.9502 ± 0.0088	0.9314–0.9642	1.647 ± 0.272	1.414–2.236
3	0.9364 ± 0.0077	0.9212–0.9471	2.129 ± 0.166	2.000–2.449
4	0.9450 ± 0.0061	0.9340–0.9567	1.961 ± 0.219	1.414–2.236
5	0.9363 ± 0.0040	0.9306–0.9444	1.984 ± 0.269	1.414–2.236

Figure 3 relates both metrics to the acquisition-time fraction over the full time ladder. The top panel reports the mean Dice loss across all runs relative to the best achieved Dice score as a baseline (ΔDice in percentage points), and the bottom panel reports the relative change in boundary error averaged over all runs (HD95 ratio, log scale). Within the range covered during training ($\tau \geq 0.3$), reducing acquisition time causes only minor changes in both metrics. For stronger reductions beyond this range ($\tau < 0.3$), ΔDice increases rapidly and HD95 rises by multiple factors. This non-linear degradation suggests that errors become increasingly dominated by spatially localized boundary outliers that drive tail statistics, rather than by a gradual global drift.

4.2. Spatial localization of geometric deviations

While Dice and HD95 quantify *how much* the segmentation changes, they do not reveal *where* deviations occur on the component. To localize geometric errors, we compute surface-deviation heatmaps by mapping point-to-point distances onto the reference geometry (Fig. 2). In the exemplar comparison (Run 1), the baseline acquisition ($\tau = 1.0$) exhibits predominantly low deviations, whereas a strongly reduced setting ($\tau \approx 0.02$) shows spatially coherent regions of elevated deviation.

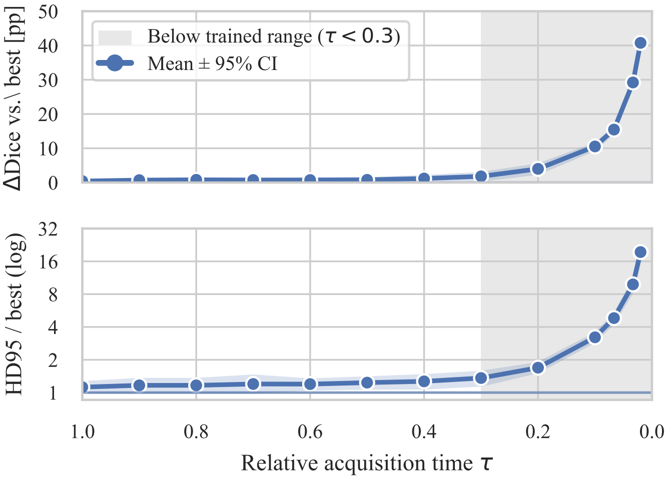

Fig. 3. Segmentation robustness under acquisition-time reduction. Top: Dice degradation relative to the best setting (ΔDice in percentage points). Bottom: boundary-error change relative to the best setting (HD95 ratio, log scale). Lines indicate the mean across runs and shaded bands show 95% confidence intervals; the shaded vertical region marks data points not in the training range.

This indicates that degradation is not spatially uniform but concentrates in specific surface areas. These localized "hotspots" can dominate tail-sensitive metrics such as HD95 and, more importantly, they are likely to drive tolerance failures on function-critical features. This motivates evaluating reduced-time CT in terms of feature-level measurement impact rather than overlap metrics alone.

4.3. *Tolerance-oriented feature evaluation on tooth profiles*

For industrial decision-making, segmentation quality must ultimately be translated into tolerance-relevant statements on measurement features rather than being reported solely via overlap metrics. We therefore perform a tolerance-oriented evaluation on the tooth profile as a representative, function-critical feature. Figure 4 summarizes the distribution of profile deviations versus relative acquisition time τ using the median (p50) and an upper-tail statistic (p95) (top), and converts these deviations into tolerance pass rates $P(\text{dev} \leq T)$ for representative tolerances $T \in \{0.2, 0.5, 1.5\}$ mm (bottom).

Across the higher-quality range ($\tau \geq 0.3$), both p50 and p95 remain comparatively stable, and tolerance compliance is largely preserved. When acquisition time is reduced below this range, the p95 increases sharply and pass rates deteriorate rapidly: failures appear first for tight tolerances (0.2–0.5 mm), while a looser threshold of 1.5 mm maintains high pass rates over a larger portion of the time ladder. This shows that time reduction primarily drives growth of the upper tail of the deviation distribution rather than a gradual increase in typical error. Together with the localized deviation patterns in Fig. 2, these results indicate that reduced-time CT can remain metrologically usable within this range, whereas extreme reductions are dominated by spatially localized geometric outliers that govern tail statistics and tolerance failures. As a practical interpretation, for tolerances in the millimeter range (e.g., $T = 1.5$ mm), the pass-rate curves indicate that acquisition time can be reduced to about $\tau \approx 0.4$ (i.e., ~60% time saving relative to the baseline) while maintaining pass rates close to 1.0 across runs.

5. Conclusion

This work demonstrates that machine learning can substantially reduce industrial CT acquisition time while still supporting metrology-relevant decisions. By combining sparse-view acquisition with learned component extraction and feature-level evaluation, we reduce acquisition to a fraction of a high-quality baseline and retain measurement usability within the studied range.

By evaluating all reduced-time reconstructions against a high-quality reference acquisition of the same clamped part, reported deviations and tolerance pass rates reflect the end-to-end pipeline including reconstruction-induced artefacts and geometric bias. We identify a practically stable range (down to $\tau \approx 0.4$) in which segmentation metrics degrade only mildly and feature-level deviations remain largely unchanged. For stronger

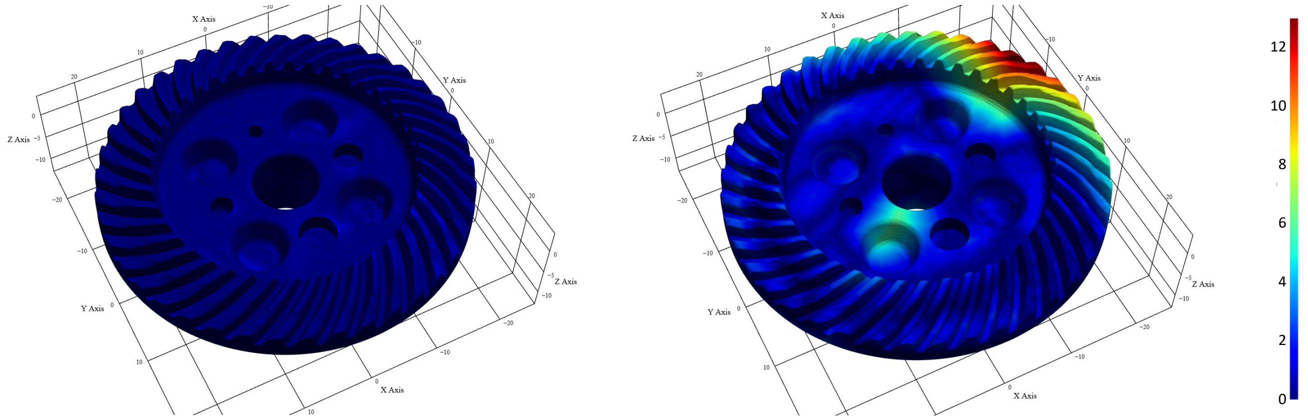

Fig. 2. Spatial localization of geometric deviations (Run 1). Point-to-point surface errors (in mm) mapped onto the reference gear mesh for the baseline acquisition ($\tau = 1.0$, left) and a strongly reduced acquisition ($\tau = 0.02$, right). Both heatmaps use an identical color scale (0–13 mm), highlighting that extreme time reduction produces spatially coherent error hotspots rather than a uniform increase across the surface.

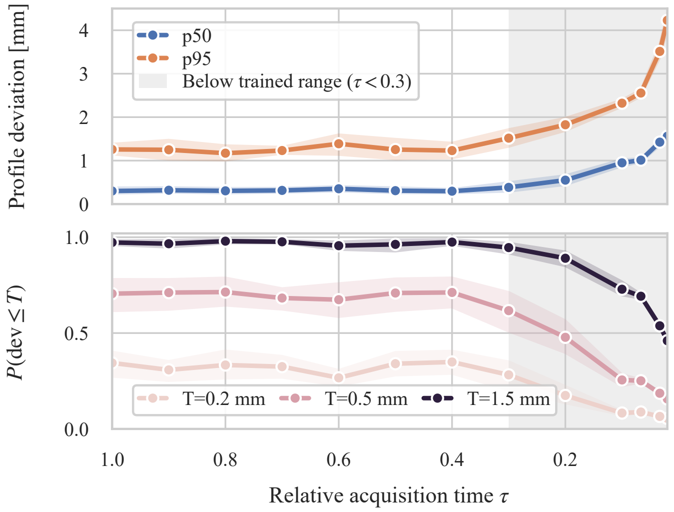

Fig. 4. Tolerance-oriented feature evaluation on a representative tooth-profile task. Top: deviation statistics (p50 and p95) versus relative acquisition time τ (normalized to the baseline scan). Bottom: tolerance pass rates $P(\mathrm{dev} \leq T)$ for $T \in \{0.2, 0.5, 1.5\}$ mm.

reductions, errors increase non-linearly and appear as spatially localized deviation hotspots that dominate tail statistics and can trigger tolerance failures even when most surface regions remain close to the reference.

At the same time, the locality of these hotspots means that very strong time reductions are not universally unusable. Depending on feature location and tolerance level, heavily reduced acquisitions may still suffice for screening. This is particularly relevant for circular-use inspection, where instance-specific decisions are required and adaptive strategies are attractive: a fast scan provides an initial component model and localized error indications, and acquisition is increased only if the relevant regions or tolerances demand it.

We note that all deviations are reported relative to a high-quality, low-noise single-component reference scan acquired at approximately five times higher resolution (voxel size $\approx 22\,\mu$m) than the assembly scans. For complex functional features such as bevel-gear tooth profiles embedded in a multi-component assembly, traceable tactile or optical reference measurements with comparable point density and coverage are practically infeasible. The chosen reference strategy therefore represents the most rigorous practically achievable baseline for this class of geometry. A formal uncertainty budget according to the GUM framework [17] remains a desirable extension but is beyond the scope of this work. Moreover, the present study uses a single 3D architecture operating at a voxel size of $\approx 253\,\mu$m, which is constrained by GPU memory. Comparing 2D and 3D segmentation strategies and their trade-off between spatial resolution, 3D context, and computational cost is an important direction for improving geometric fidelity beyond the current voxel-level limit. Finally, the specific time fractions reported here are setup- and task-dependent and should not be generalized. What transfers is the evaluation principle: linking acquisition-time reduction to feature-level deviations and tolerance compliance, rather than relying on segmentation metrics alone.

Acknowledgements

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SFB 1574–471687386.

References

[1] 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.

[2] DIN SPEC 91472:2023-06, “Remanufacturing (Reman) – Qualitätsklassifizierung für zirkuläre Prozesse,” 2023.

[3] G. Lanza, B. Deml, S. Matthiesen, M. Martin, O. Brützel, and R. Hörsting, ‘The vision of the circular factory for the perpetual innovative product’, at - Automatisierungstechnik, vol. 72, no. 9, pp. 774–788, Sep. 2024, doi: 10.1515/auto-2024-0012.

[4] A. Weckenmann and P. Krämer, ‘Computed tomography in quality control: chances and challenges’, Proc. IMechE Part B: J. Engineering Manufacture, vol. 227, no. 5, pp. 634–642, 2013, doi: 10.1177/0954405413479849.

[5] H. Villarraga-Gómez and S. T. Smith, ‘Effect of the number of projections on dimensional measurements with X-ray computed tomography’, Precision Engineering, vol. 66, pp. 445–456, Nov. 2020, doi: 10.1016/j.precisioneng.2020.08.006.

[6] A. Hatamizadeh et al., ‘UNETR: Transformers for 3D Medical Image Segmentation’, Oct. 09, 2021, arXiv: arXiv:2103.10504. doi: 10.48550/arXiv.2103.10504.

[7] E. Y. Sidky and X. Pan, ‘Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization’, Phys. Med. Biol., vol. 53, no. 17, pp. 4777–4807, Sep. 2008, doi: 10.1088/0031-9155/53/17/021.

[8] Z. Zhang, X. Liang, X. Dong, Y. Xie, and G. Cao, ‘A Sparse-View CT Reconstruction Method Based on Combination of DenseNet and Deconvolution’, IEEE Trans. Med. Imaging, vol. 37, no. 6, pp. 1407–1417, Jun. 2018, doi: 10.1109/TMI.2018.2823338.

[9] S. Bellens, P. Guerrero, P. Vandewalle, and W. Dewulf, ‘Machine learning in industrial X-ray computed tomography – a review’, CIRP Journal of Manufacturing Science and Technology, vol. 51, pp. 324–341, Jul. 2024, doi: 10.1016/j.cirpj.2024.05.004.

[10] R. Gruber et al., ‘Instance Segmentation XXL-CT Challenge of a Historic Airplane’, J Nondestruct Eval, vol. 44, no. 1, p. 1, Mar. 2025, doi: 10.1007/s10921-024-01136-y.

[11] R. Gruber, S. Rüger, and T. Wittenberg, ‘Adapting the Segment Anything Model for Volumetric X-ray Data-Sets of Arbitrary Sizes’, Applied Sciences, vol. 14, no. 8, p. 3391, Apr. 2024, doi: 10.3390/app14083391.

[12] J. Zheng, C. Tang, Y. Sun, M. Feng, and C. Wang, ‘An Enhanced U-Net Approach for Segmentation of Aeroengine Hollow Turbine Blade’, Mathematics, vol. 10, no. 22, p. 4230, Nov. 2022, doi: 10.3390/math10224230.

[13] Y. Xue et al., ‘Shape-Aware Organ Segmentation by Predicting Signed Distance Maps’, Dec. 09, 2019, arXiv: arXiv:1912.03849. doi: 10.48550/arXiv.1912.03849.

[14] K. A. Juhl et al., ‘Guiding 3D U-nets with signed distance fields for creating 3D models from images’, Aug. 28, 2019, arXiv: arXiv:1908.10579. doi: 10.48550/arXiv.1908.10579.

[15] K. A. Juhl, J. Slipsager, O. de Backer, K. Kofoed, O. Camara, and R. Paulsen, ‘Signed Distance Field based Segmentation and Statistical Shape Modelling of the Left Atrial Appendage’, Feb. 12, 2024, arXiv: arXiv:2402.07708. doi: 10.48550/arXiv.2402.07708.

[16] F. Zanini, N. Bonato, D. Pentucci, and S. Carmignato, ‘Investigating the effects of machine learning generalization for enhancing accuracy in fast X-ray computed tomography for industrial metrology’, CIRP Annals, vol. 74, pp. 719–723, 2025, doi: 10.1016/j.cirp.2025.04.025.

[17] J. P. Kruth, M. Bartscher, S. Carmignato, R. Schmitt, L. De Chiffre, and A. Weckenmann, ‘Computed tomography for dimensional metrology’, CIRP Annals, vol. 60, no. 2, pp. 821–842, 2011, doi: 10.1016/j.cirp.2011.05.006.