

A learning-based method for online adjustment of C-arm Cone-beam CT source trajectories for artifact avoidance

Journal Article**Author(s):**

Thies, Mareike; Zaech, Jan-Nico; Gao, Cong; Taylor, Russell; Navab, Nassir; Maier, Andreas; Unberath, Mathias

Publication date:

2020-11

Permanent link:

<https://doi.org/10.3929/ethz-b-000438990>

Rights / license:

[Creative Commons Attribution 4.0 International](#)

Originally published in:

International Journal of Computer Assisted Radiology and Surgery 15(11), <https://doi.org/10.1007/s11548-020-02249-1>



A learning-based method for online adjustment of C-arm Cone-beam CT source trajectories for artifact avoidance

Mareike Thies^{1,2} · Jan-Nico Zäch³ · Cong Gao¹ · Russell Taylor¹ · Nassir Navab¹ · Andreas Maier² · Mathias Unberath¹

Received: 23 March 2020 / Accepted: 14 August 2020 / Published online: 25 August 2020
© The Author(s) 2020

Abstract

Purpose During spinal fusion surgery, screws are placed close to critical nerves suggesting the need for highly accurate screw placement. Verifying screw placement on high-quality tomographic imaging is essential. C-arm cone-beam CT (CBCT) provides intraoperative 3D tomographic imaging which would allow for immediate verification and, if needed, revision. However, the reconstruction quality attainable with commercial CBCT devices is insufficient, predominantly due to severe metal artifacts in the presence of pedicle screws. These artifacts arise from a mismatch between the true physics of image formation and an idealized model thereof assumed during reconstruction. Prospectively acquiring views onto anatomy that are least affected by this mismatch can, therefore, improve reconstruction quality.

Methods We propose to adjust the C-arm CBCT source trajectory during the scan to optimize reconstruction quality with respect to a certain task, i.e., verification of screw placement. Adjustments are performed on-the-fly using a convolutional neural network that regresses a quality index over all possible next views given the current X-ray image. Adjusting the CBCT trajectory to acquire the recommended views results in non-circular source orbits that avoid poor images, and thus, data inconsistencies.

Results We demonstrate that convolutional neural networks trained on realistically simulated data are capable of predicting quality metrics that enable scene-specific adjustments of the CBCT source trajectory. Using both realistically simulated data as well as real CBCT acquisitions of a semianthropomorphic phantom, we show that tomographic reconstructions of the resulting scene-specific CBCT acquisitions exhibit improved image quality particularly in terms of metal artifacts.

Conclusion The proposed method is a step toward online patient-specific C-arm CBCT source trajectories that enable high-quality tomographic imaging in the operating room. Since the optimization objective is implicitly encoded in a neural network trained on large amounts of well-annotated projection images, the proposed approach overcomes the need for 3D information at run-time.

Keywords Tomographic reconstruction · Metal artifact reduction · Deep learning · Image-guided surgery

M. Thies, J-N. Zäch contributed equally and are listed in alphabetical order.

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s11548-020-02249-1>) contains supplementary material, which is available to authorized users.

✉ Mathias Unberath
unberath@jhu.edu

¹ Laboratory for Computational Sensing + Robotics, Johns Hopkins University, Baltimore, MD, USA

² Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany

Introduction

The number of patients undergoing spinal fusion surgery in the USA has been increasing rapidly over the last years. From 204,000 cases in 1998, the number of interventions has grown to 457,000 cases in 2011. This growth also resulted in a steep increase in hospitalization charges of more than 750% as each

³ Computer Vision Laboratory, Eidgenössische Technische Hochschule Zürich, Zurich, Switzerland

intervention is relatively expensive with an average hospital bill of more than \$34,000 [2,7]. During spinal fusion surgery, two or more vertebrae are fused. Usually, this is achieved by implanting screws into the affected vertebrae in a transpedicular approach. These screws are then interlocked with metal rods to inhibit movement and induce the formation of new bony structure. This stabilizes the spine at the given location, which can reduce chronic back pain when non-operative treatment failed [1]. Despite the high number of interventions, spinal fusion surgery remains a high-risk operation. In 2010, 6.8% of the patients undergoing spinal fusion interventions in the US were rehospitalized within the first 30 days after surgery [32]. Misplacement of pedicle screws has been reported in up to 55% of the cases using traditional free-hand technique for screw placement [17]. While this number can be decisively reduced when the intervention is performed under fluoroscopy-guidance [10], even in navigated approaches the incidence of misplaced screws remains high [14]. Screw misplacement intrinsically carries the danger of cortical breach which can result in nerve damage and severe neurological impairment of the patient [9,26] and has been found in up to 8% of the cases [17]. Consequently, there is a need for an imaging modality capable of precisely capturing the anatomy in direct proximity of metal implants to assess the adequacy of implant placement during the intervention, allowing for immediate revision in case of misplacement. Tomographic reconstructions available through C-arm cone-beam CT (CBCT) have the potential to provide such information intraoperatively. CBCT has been deployed for this purpose [13,23]. However, even with one of the most recent CBCT devices and compared to conventional postoperative CT, 23% of the cases of cortical breach were missed on intraoperative CBCT images primarily due to much stronger metal artifacts around the screw [5]. Improving the quality of intraoperative CBCT reconstructions for the task of pedicle screw placement in spinal fusion surgery, consequently, has a great potential to identify cortical breach during the operation, allowing for immediate revision, and therefore, reduction of both neurological complications and need for revision surgery.

Background

Artifacts in CT reconstructions are a consequence of discrepancies between the assumed mathematical forward model of X-ray image formation and the real physics of data acquisition [21]. The process of reconstruction aims at finding a volumetric representation that optimally explains all measured X-ray images which is the inverse problem to the image formation process. Usually, this is performed by backprojecting the images into the volume thereby inverting the idealized mathematical assumption of the forward process [8]. This is a well-posed problem as long as the mentioned discrepancies

remain small. High discrepancies, however, render this an ill-posed inverse problem resulting in artifacts in the reconstructed volume. Other than noise, one of the most prominent discrepancies is beam hardening. It is characterized by a shift between incident and recorded energetic spectrum of the photons due to energy-dependent attenuation in dense objects, e.g., in our case titanium screws. This affects measurements on the detector and results in the overestimation of certain pixels during backprojection [22]. Existing approaches to artifact reduction in CBCT reconstructions usually rely on postprocessing of the acquired data in projection domain [16,18,35] or artifact suppression in volume domain [15]. These methods work on the corrupted data and oftentimes carry the risk of introducing new sub-optimal image content which compromises the quality of the reconstruction in a different way. Therefore, we propose to begin artifact reduction one step earlier by adjusting the CBCT protocol to the scene and directly acquire better data. Specifically, our approach automatically predicts adjustments to the C-arm trajectory in real-time to actively exploit views onto the anatomy which are most consistent with the assumptions made in the tomographic reconstruction process. Selection of good views is performed by a convolutional neural network that regresses a view-dependent quality index from the current projection of an ongoing scan. We hypothesize that this finally leads to artifact reduction and improved quality of the reconstruction. Ultimately, such approach could allow for intraoperative tomographic imaging with clinically acceptable quality in applications such as spinal fusion surgery.

Related work

Conceptually related ideas to ours have been proposed for real-time user guidance in free-hand ultrasound probe motion. In [19] the ultrasound image of each time point is interpreted by a deep reinforcement learning agent which predicts an incremental update on the probe motion. Similarly, incremental and real-time user feedback can be provided in the case of SPECT imaging with mobile freehand detectors based on the numerical condition of the system matrix corresponding to the reconstruction problem [31]. However, analyzing the entire system matrix is not feasible for CT due to its memory footprint. Instead, different approaches have been proposed to select the most valuable next projection for CT: The method in [36] favors views that sample rays tangential to edges of the 3D object to maximize the edge information in the reconstruction, obviously relying on precise knowledge of the object at optimization time. In [6] angular steps are selected such that with each additional projection the set of solutions that are consistent with the set of already measured projections is minimized. This, however, has only been applied to the 2D reconstruction case and is computationally very expensive as it needs several

preliminary reconstructions per step. Recently, finding an optimal sinusoidal trajectory which avoids metal parts of the imaged object while still ensuring a high coverage in Radon space for its direct vicinity has been studied in [11]. While all these approaches do not directly consider the X-ray imaging physics, [25] proposes an index to analyze the quality of different projections based on the local point spread function and noise power spectrum of the imaging device. Similarly, [33] calculates a quality map of possible views from the expected amount of spectral shift due to beam hardening depending on different path lengths of the photons through metal objects. Both methods were successfully applied to CBCT trajectory optimization. Yet, all previous approaches calculate optimal parameters in a (semi-)offline manner and rely on knowledge about the 3D object at optimization time, which is usually provided by a preoperative scan. This requirement is problematic since, during interventions, the anatomy is altered in an unpredictable way, e.g., by screw insertion.

Contribution

In this work, we expand on our MICCAI 2019 submission that introduced machine learning-based algorithms to predict on-the-fly adjustments for task-based C-arm trajectories [34]. Our contributions are twofold. First, in carefully controlled experiments on synthetic data we characterize the algorithm’s behavior and robustness (1) in the presence of varied noise levels, and (2) with varying initial poses of the C-arm gantry with respect to anatomy. Second, we substantially expand our experiments on real data acquired from a semianthropomorphic phantom. To this end, we acquire 17 CBCT short-scans at different swivel and tilt angles of the gantry and align all projection images to a common 3D object space via image-based registration. This produces a set of calibrated X-ray images that allows for validating the proposed C-arm servoing algorithm in a retrospective manner. This strategy enables feasibility studies on real X-ray

data but avoids the need for (1) a fully robotized and freely steerable C-arm device, and (2) flexible and robust online calibration methods that accurately estimate the C-arm imaging geometry without prior knowledge on the 3D scene. Exploring solutions to these challenges is important and will be the subject of our future work.

Methods

Online trajectory adjustment pipeline

Trajectory optimization is a problem with many degrees of freedom because recent scanners can realize very different motion patterns. Following the ideas in [25], we choose to parameterize the problem in terms of an in-plane angle φ and an out-of-plane angle θ . The in-plane angle is defined according to a traditional circular trajectory where source and detector move in one plane for the entire scan, whereas the out-of-plane angle is associated with tilting the C-arm relatively to this plane. Each trajectory consists of a set of pairs (φ_t, θ_t) , $t = 0, \dots, T$ where T is the total amount of projections images. The general pipeline we propose is illustrated in Fig. 1. An X-ray image is captured at a position (φ_t, θ_t) and processed by a VGG-type convolutional neural network, which regresses a detectability index (see “Projection-dependent detectability index” Section) for the next possible projections. The projection with the highest predicted value is identified and the out-of-plane angle θ_{t+1} is updated accordingly while the in-plane angle φ is always incremented by a fixed amount: $\varphi_{t+1} = \varphi_t + \Delta\varphi$. The new target $(\varphi_{t+1}, \theta_{t+1})$ is sent to the robotic C-arm and its position adjusted accordingly to acquire the next projection.

Projection-dependent detectability index

To assess how single projections contribute to perceived reconstruction quality, we follow existing approaches based

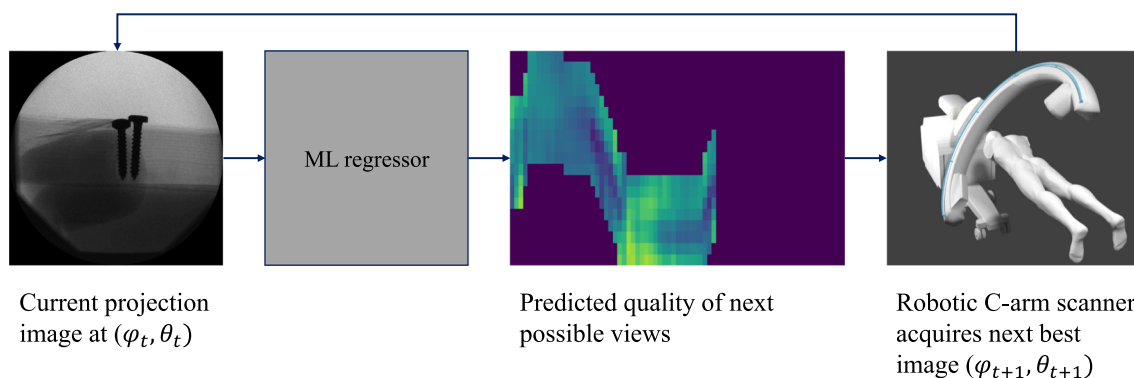


Fig. 1 High-level overview of the envisioned pipeline for online trajectory adjustment

on the non-prewhitening matched filter observer model which allows to find a so-called detectability index as per Eq. 1 [12].

$$d^2(\varphi, \theta) = \frac{\left[\int \int \int |\text{MTF}(\varphi, \theta)|^2 |W_{\text{task}}|^2 d f_x d f_y d f_z \right]^2}{\int \int \int \text{NPS}(\varphi, \theta) |\text{MTF}(\varphi, \theta)|^2 |W_{\text{task}}|^2 d f_x d f_y d f_z} \quad (1)$$

MTF is the local modulation transfer function, NPS is the local noise power spectrum and W_{task} is a task-function describing the properties of the object to be imaged with highest quality in Fourier space. For the case of iterative penalized-likelihood reconstruction, it is possible to derive analytic expressions for both MTF and NPS [12]. These equations rely on forward projecting voxels into all views contained in a trajectory, comparing the projected value with the measured values and back projecting this information into the volume. Using these calculations for MTF and NPS, the final detectability index d^2 thus depends on the 3D structure of the imaged object as well as the set of images in a trajectory. This means that, if accurate 3D information is available, Eq. 1 can be maximized with respect to φ and θ to find an optimal trajectory. Note that the local MTF and NPS are very general measures which can be calculated for any imaged object. This work is centered around metal artifacts suppression as these are usually the most severe artifacts during interventions. In a different setting, the same index could potentially also be used for improving, e.g., soft-tissue contrast.

Network for detectability prediction

During an intervention, the volume to be imaged is altered compared to preoperatively acquired information. Therefore, offline trajectory optimization approaches (e.g., the one outlined in “Projection-dependent detectability index” Section) usually cannot succeed in these cases. As introduced in our previous work [34], we instead propose to regress the detectability index in Eq. 1 on-the-fly during an ongoing scan using a convolutional neural network (CNN) using only fluoroscopic images as input. In this approach, knowledge about the task is encoded in the weights of the machine learning model, thereby overcoming the need for explicit 3D information at CBCT acquisition time. We rely on an architecture that is similar to the VGG architecture [24], but adapted to perform regression instead of classification because we believe that a highly parameterized CNN is well suited to implicitly capture the underlying 3D structure in a learning-based manner. From an input X-ray projection, the network is trained to predict the detectability of those projections with an increment of $+5^\circ$ in in-plane angle and a range of $[-25^\circ, +25^\circ]$ in out-of-plane angle relative to the current position. The out-

of-plane interval is discretized in steps of 5° which leads to 11 values to be predicted from each input image. For training, two different datasets were generated by forward projecting 3D volumes using the open-source physics-based X-ray simulator DeepDRR [27,28]. The resulting digitally reconstructed radiographs (DRRs) were created on a uniform grid with step size 5° in both φ and θ . For each position (φ, θ) , one clean image and one image with additional realistic noise injection were generated. The former was used to calculate ground truth detectability for each projection using Eq. 1 and the corresponding 3D scan while the latter was used as actual network input during training. The first dataset is based on five publicly available chest CT scans from the Cancer Imaging Archive (TCIA) [4]. Screw positions were manually annotated in six different vertebrae per scan. During the generation of projection data, only one vertebral level was considered at a time and a titanium screw was virtually inserted at the annotated position into the corresponding anatomy. Additionally, the isocenter of the simulated C-arm was varied randomly between the different simulations. 212 simulations were performed on 30 different anatomical sites, each resulting in 1368 images on a 5° grid with a whole rotation ($[0^\circ, 360^\circ]$) for the in-plane angle and an interval of $[45^\circ, 135^\circ]$ for the out-of-plane angle. The resulting images of one chest CT scan are held out as a test set. Data generation for the second dataset is identical to the first dataset, but based on a semianthropomorphic representation of a human chest that is composed of a long box-like object, two cylinders, and two screws. The position of these objects was randomly varied within reasonable bounds to account for different anatomy from patient to patient finally leading to 75 simulations again consisting of 1368 images each, distributed over the same interval as above. The test set consists of three simulations. The first dataset is used for the experiments on synthetic data, while the second one is used to train the network in the real data case. Additionally, batch normalization and data augmentation using random rotations were included in the network for the real data experiments as we observe that it helps to improve generalization.

Experiments and results

Simulation experiments

The network was first trained on the TCIA chest dataset. The training objective was to predict the detectabilities of the 11 projections with an offset of $+5^\circ$ in in-plane angle and an interval of $[-25^\circ, +25^\circ]$ in out-of-plane angle relative to the position of the input projection discretized in steps of 5° . During inference, an out-of-plane increment was chosen to be a step toward the highest predicted detectability. Additionally, the whole trajectory was restricted to an


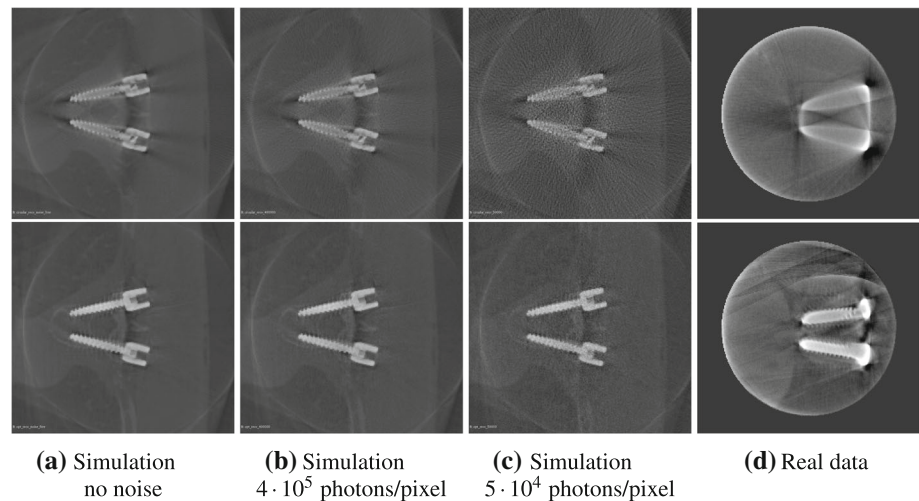


Fig. 2 Spatial distribution of the angular and detectability error. The X-axis shows the full 360° in in-plane angle ϕ and the Y-axis possible out-of-plane angles θ between 45° and 135°

Fig. 3 Slices through the reconstructions of synthetic and real data from a circular scan (upper row) and the task-aware trajectory (lower row) at different noise levels. Note that the simulated screws are not identical to the screws of the real phantom in size or shape



interval of $[-45^\circ, +45^\circ]$ concerning the out-of-plane angle relative to the starting position. In a purely simulated environment without realistic noise injection, the algorithm achieves $8.35^\circ \pm 11.61^\circ$ angular distance and $13.69 \pm 18.92\%$ relative difference in detectability of the predicted trajectory compared to the ground truth [34]. Their angular distributions are shown in Fig. 2. In the following, the influence of different levels of noise and varying initialization poses of the C-arm on the prediction quality will be analyzed. Eight different 200° short scan protocols are simulated for each screw pair in the test-set. Half of the protocols employ a circular trajectory, each with 200 X-ray projections in total. These serve as the baseline protocols. The other half of the scans are generated on trajectories optimized with the proposed pipeline. For each of the two trajectory types, scans without noise and with a noise level corresponding to $5 \cdot 10^4$, $1 \cdot 10^5$, and $4 \cdot 10^5$ photons per pixel are generated. Each X-ray projection is acquired with 620×480 pixels and a pixel-size of 0.31×0.31 mm. This image corresponds to the central part of a standard flat-panel detector in 4×4 binning mode. A figure showing the predicted trajectories in the presence of different noise levels can be found in the supplementary material. Only small angular changes are observed which proves robustness against noise. Furthermore, the network did not overfit to a single detectability map, as the trajec-

tories generated from different vertebral levels show major differences. Defining the trajectory in the noise-free case as the ground truth prediction allows calculation of the sensitivity to noise. The sensitivity is calculated as the angular mismatch, averaged over all angles and trajectories for a single noise level. For $4 \cdot 10^5$, $1 \cdot 10^5$, and $5 \cdot 10^4$ photons per pixel, the mean angular error reads 0.83° , 1.13° , and 1.64° , respectively. The standard deviation of the predictions is 1.56° , 1.63° , and 1.73° in the same order. Besides robustness against noise, it is desirable that the optimal trajectory is largely independent of the starting angle. This property holds for the proposed algorithm, as a prediction only depends on the last acquired image. Therefore, two trajectories that intersect at any point will merge and continue as the same trajectory, given the noise is identical. To show this property on data, the trajectories predicted from different starting angles, but the same anatomy were simulated (see plot in supplementary material). After few angle increments, the trajectories merge into two main bands that represent local maxima, which then merge into a single trajectory at $\phi = 50^\circ$. The initial differences of the trajectories can be explained by the limitations of the slope. The predicted trajectories were reconstructed using a GPU implementation of the iterative conjugated gradient least squares algorithm for cone-beam geometry provided by the ASTRA toolbox [29,30]. Figure 3

shows axial slices through the reconstructions from projections at different noise levels for qualitative analysis. For quantitative assessment, comparison is performed by computing the full width half maximum (FWHM) of the screws of one vertebral level averaged over two different positions which quantifies the amount of blooming artifact. Further, we investigate the intensity of the Fourier spectrum of a small normalized image patch containing the screw thread at the frequency of the thread itself. For comparison, the ground truth value for each of these measures is listed which is obtained by reconstructing mono-energetic, noise-free simulated projections without any physics-based artifacts. We also report the structural similarity (SSIM) of a slice containing both screws between the ground truth and the noisy reconstructions. Results are reported in Table 1. Both the FWHM and the thread frequency height are closer to the true value for the task-aware trajectories compared to the circular ones. Also the image slices extracted from the reconstructions are more similar to the ground truth slices as indicated by higher SSIM values. Noise in general deteriorates the reconstruction performance, but this seems to be less severe for the task-aware trajectories.

Real data experiments

One central challenge when implementing non-circular orbits on any CBCT scanner—robotized or conventional C-arm—is the calibration required for precise reconstruction. Usually only few reproducible circular short-scan trajectories are calibrated in advance using phantoms specifically designed for this purpose. The trajectories aimed for here, however, cannot be precalibrated because they are scene specific, and thus, not known in advance. To overcome this challenge, in this work, multiple pre-calibrated CBCT short-scans of the phantom were acquired at various swivel and tilt angles prior to trajectory prediction. The CBCT reconstructions, and via pre-calibration also all projection images, acquired in this way were then aligned in a common 3D object space using image-based registration. This procedure aims at providing a sufficient sampling of all possible views (φ, θ) and is explained in detail in the following paragraph. During inference, the sampled view closest to the predicted optimal view is identified and used as subsequent input for the network instead of an image acquired in real-time by a robotic device. This allows to predict non-circular trajectories from real data using the proposed method in a retrospective manner and avoids the need for a fully robotic C-arm. Instead, data acquisition was performed on a conventional CBCT scanner (Siemens Arcadis Orbic 3D). For the experiments, a phantom was built in line with the simulated training data for this case. It consists of two screws drilled into a wooden rod and two cylinders filled with ballistic gel.

Sampling the (φ, θ) space using only circular scans can be achieved by scanning the phantom on tilted but circular orbits (see Fig. 4 left). However, tilting the scanner would require calibration of each of these tilted trajectories due to mechanical sagging and wobble. Instead, the position of the phantom itself was altered between successive scans while the scanner trajectory was kept identical. In this manner, 17 scans were acquired mimicking tilt as well as swivel of the C-arm. In terms of the in-plane and out-of-plane angle notation, each of these 17 scans results in a curve of sampled views in the (φ, θ) space (see Fig. 4 right). All 17 scans were reconstructed and 16 tilted scans were rigidly registered to one reference volume. Registration was performed by optimizing a normalized cross-correlation objective function using quadratic optimization (BOBYQA). The rigid transformation T_i aligning the i -th moving tilted volume with the reference volume was obtained and used to adjust the projection matrices as:

$$P_{\text{tilt}}^i = T_i^{-1} P_{\text{flat}} \quad (2)$$

Applying the inverse transformation to the projection matrices allows changing from several volumes reconstructed with the same set of matrices P_{flat} to a scan-specific set of matrices P_{tilt}^i such that all projections can be integrated into the same volume during reconstruction. The network was trained on the dataset created from a digital copy of the used phantom mentioned in “Network for detectability prediction” Section. Training ground truth was chosen to be identical to the setup described for the synthetic data experiments in “Simulation experiments” Section. During inference, increments in in-plane angle were fixed to $\Delta\varphi = 1^\circ$ and the out-of-plane angle step is computed from the predicted detectability and a regularization component that penalizes high directional changes and promotes a smoother trajectory. As the pipeline is targeted to be implemented on a robotic C-arm device, we need to account for the limited mechanical capabilities of such a system. Sudden directional changes would require high accelerations that cannot be realized safely. Therefore, we introduce the cosine of the angle between two subsequent steps as additional penalty term. With this term, sudden directional changes are traded off with best next steps as predicted by the network.


$$\theta_{t+1} = \theta_t + \max_i (\lambda(u \cdot v_i) + p_i) \quad (3)$$

Here, u denotes the previous trajectory direction in terms of $(\Delta\varphi, \Delta\theta)$, v_i is the i -th possible next direction and p_i is the corresponding predicted detectability. We heuristically find that $\lambda = 0.6$ is a suitable weighting factor and keep it constant for all experiments. The projection image which is closest to the optimal predicted view in terms of φ and θ is identified from the set of acquired projections, added to the trajectory,

Table 1 Evaluation of reconstruction quality based on screw FWHM, screw thread frequency peak height and SSIM for circular and task-aware trajectories and different noise levels on simulated data

	Screw FWHM [mm]	Thread frequency height	SSIM to ground truth
Ground truth	3.92	6.83	1.00
Circular no noise	6.38	9.05	0.83
Circular $4 \cdot 10^5$ photons/pixel	6.35	8.96	0.81
Circular $5 \cdot 10^4$ photons/pixel	6.19	7.52	0.69
Task-aware no noise	3.65	7.57	0.90
Task-aware $4 \cdot 10^5$	3.72	8.31	0.89
Task-aware $5 \cdot 10^4$	4.10	6.95	0.85

Fig. 4 Left: Two exemplary tilted orbits and a non-circular trajectory with varying out-of-plane angle. Right: Sampling of the (φ, θ) -space using tilted orbits. Solid red represents the untilted reference scan, dashed black refers to scans acquired on tilted circular orbits



deleted from the set of available sampled views for all following steps, and used as next input for the network. Using the first projection of the reference scan which corresponds to an out-of-plane angle of 90° (scan plane intersecting long axis of both screws) as initialization of the algorithm, this procedure results in the trajectory depicted in Fig. 5. From the initial out-of-plane angle, the algorithm proposes to increase the tilt of the C-arm for the majority of the scan. The trajectory reaches the most extreme sampled out-of-plane angles in positive direction for in-plane angles 50° to 80° and the most extreme angles in negative direction toward beginning and end of the scan. In the central part, it exhibits a slightly alternating behavior. Note that in the real data case, only views that have been sampled can be part of the trajectory which limits the number of possible solutions considerably. Reconstructions were calculated using the same algorithm as for the synthetic data [29,30]. Projection images were masked prior to reconstruction based on forward projecting a centered sphere with 5 cm radius in 3D to reduce truncation artifacts and the algorithm was executed for 300 iterations. A slice through the reconstructed volume of the trajectory corresponding to Fig. 5 and the circular reference trajectory can be found in the last column of Fig. 3. While the overall shape of the two screws, as well as its threads, are only poorly recovered in the reconstruction from the circular trajectory, the task aware protocol is able to recover much finer structures. For quantitative evaluation, we additionally initialize

our algorithm with the first projection of the four swivel trajectories in our dataset, which each provide initialization with a different out-of-plane angle. We compare reconstructions obtained from all these trajectories to the reference circular trajectory and the two circular trajectories of our dataset associated with the highest tilt and swivel, respectively. Comparison is again performed using the FWHM of the screw and the thread frequency peak height. Results can be found in Table 2. Calculating the SSIM is not possible as no ground truth information is available. The circular reference scan performs worst by far, exhibiting the largest FWHM of all scans and revealing severe problems in visualizing the shape of the screw. While the trajectories corresponding to maximum tilt and swivel perform best when considering either FWHM or peak height, respectively, the task-aware trajectories can improve both measures decisively compared to the reference scan. Initializing with angles different from the reference scan (90° out-of-plane) additionally seems to improve the ability to reconstruct the screw thread.

Discussion

The presented results on simulated data help to understand strengths and limitations of the method in a controlled setting and serve as an upper bound of the ideally achieved performance. They show that predicting the detectability values

Fig. 5 Predicted trajectory on real data in black and network predictions relative to the current position. The crosses show optimal views based on network output, the black line is the final trajectory based on the closest sampled view




Table 2 Evaluation of reconstruction quality based on screw FWHM and screw thread frequency peak height for different trajectories

	Screw FWHM [mm]	Thread frequency height
Circular reference	12.67	2.11
Circular max. tilt	6.29	3.24
Circular max. swivel	9.36	9.79
Task-aware init 67°	7.35	5.98
Task-aware init 76°	7.29	5.99
Task-aware init 90°	6.93	4.33
Task-aware init 99°	6.92	4.38
Task-aware init 121°	6.79	9.42

of possible next views from the current projection is possible with reasonable accuracy and robustness against different noise levels and initialization angles. The resulting trajectories are in line with previously published concepts on the emergence of reconstruction artifacts introduced in “Introduction” Section. If possible, our algorithm avoids views with overlapping screws as well as views along the screws’ long axes, which cause the most severe inconsistencies (beam hardening up to photon starvation) based on the assumptions made during reconstruction. Therefore, the fine structures of the screws can be reconstructed with higher quality and metal artifacts can be reduced significantly. The performed real data experiments hint at the feasibility of the approach in real CBCT acquisitions. Benchmarks for the inference time of VGG-19 point out that it is generally feasible to use the network predictions for real-time adjustment of the C-arm [3], but as our real data evaluation was performed retrospectively, we did neither implement a real-time capable system including the mechanical components nor did we investigate whether the final trajectory can be realized by a scanner within reasonable scanning time. The retrospec-

tive evaluation still suggests that the task-aware trajectories lead to considerably improved reconstructions of the screws’ general shape as well as its thread. This holds true especially when comparing to the reference scan, the scan plane of which is parallel to both screws’ long axes. Unfortunately, this acquisition scheme is most predominantly employed in the operating room. This leads to the conclusion that slightly re-positioning the C-arm to acquire a short scan trajectory that is tilted with respect to the standard plane already avoids many of the worst views, and would thus already result in considerably improved reconstruction quality without changes to the routine acquisition protocol.

On real data, the predicted trajectory shows an increased alternating behavior between positive and negative out-of-plane angles compared to the simulations. Possible reasons for this are a sub-optimal generalization of the network from its training domain to the domain of real images which could potentially be mitigated by an increased amount of training data. Moreover, the network fails to disambiguate positive and negative increments for the out-of-plane angle in some cases while still clearly following the trend to favor high out-of-plane angles in our specific setup. This behavior results in trajectories that tend to jump between high positive and high negative out-of-plane angles and might be caused by the Markov property of the algorithm described here. As each prediction is only based on one preceding projection image, there is very little contextual information available that could be used to disambiguate predictions with similar detectability. The predicted trajectory still leads to remarkable improvements in the reconstruction results compared to the circular reference trajectory which can already improve the ability for accurate clinical assessment tasks. Still, there are many open challenges which need to be addressed to push the approach closer to the level of accuracy and robustness needed for clinical application. First, the retrospective calibration procedure presented here on a non-robotic C-arm

with only one actuated axis to enable CBCT is not applicable in a clinical setting because it would expose the patient to high doses of ionizing radiation. Instead, an online calibration procedure which does not require precise knowledge about the 3D structure would be desirable. Relying on the joint encodings of a fully robotic C-arm for initialization, further fine-tuning of the pose parameters could be performed in an image-based manner, e.g., using autofocus measures [20]. To ensure robust and precise network predictions in a clinical environment, important steps are the generation of synthetic training data, which is representative of the variety of different anatomies and tools as well as views onto these. The domain gap between the simulations used for training and the real fluoroscopy images during inference could be minimized using state-of-the-art domain adaptation techniques. Once a task-aware protocol is deployed in practice and a broader spectrum of fluoroscopy images from different views onto the anatomy becomes available, real data from predicted trajectories can be used to directly retrain the network parameters. Experimenting with different network architectures might also improve prediction performance. Finally, we envision a clinically applicable version of the pipeline to supersede existing CBCT protocols as they are already applied in the operating room.

Conclusion

We introduced a learning-based method for online CBCT trajectory adjustment that overcomes the need for volumetric information at imaging time. This is the first step toward high-quality intra-operative C-arm CBCT imaging which is based on the idea of directly acquiring better data for artifact avoidance. Such an approach might ultimately enable intraoperative verification of implant placement with high confidence, as is required for high volume procedures including spinal fusion surgery. Future work will address the lack of sequential modeling in the current approach and investigate whether the image quality delivered by a refined version of our approach is sufficient for clinical interpretation.

Acknowledgements We gratefully acknowledge the support of National Institutes of Health R01 EB023939, R01 EB016703, R21 EB028505, and JHU Internal Funds. We further acknowledge the support of the NVIDIA Corporation with the donation of the GPUs used for this research. We would like to thank Gerhard Kleinszig and Sebastian Vogt, both with Siemens Healthineers, for making the C-arm used in this research available. Mareike Thies was supported by a DAAD Promos Fellowship.

Funding The concepts and information presented in this paper are based on research and are not commercially available. Open Access funding provided by Projekt DEAL.

Compliance with ethical standards

Conflict of interest The authors have no conflict of interest to declare.

Ethical approval This article does not describe research on human subjects; therefore, ethical approval was not required.

Informed consent This article does not contain patient data.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Aebi M (2010) Indication for lumbar spinal fusion. In: Szpalski M, Gunzburg R, Rydevik BL, Le Huec JC, Mayer HM (eds) *Surgery for low back pain*. Springer, Berlin, pp 109–122
2. Andersson G, Watkins-Castillo SI (2014) Spinal fusion. The burden of musculoskeletal diseases in the United States (BMUS). <https://www.boneandjointburden.org/>. Accessed 17 Feb 2020
3. Canziani A, Paszke A, Cururciello E (2016) An analysis of deep neural network models for practical applications. arXiv preprint [arXiv:1605.07678](https://arxiv.org/abs/1605.07678)
4. Clark K, Vendt B, Smith K, Freymann J, Kirby J, Koppel P, Moore S, Phillips S, Maffitt D, Pringle M, Tarbox L (2013) The cancer imaging archive (TCIA): maintaining and operating a public information repository. *J Digit imaging* 26(6):1045–1057
5. Cordemans V, Kaminski L, Banse X, Francq BG, Cartiaux O (2017) Accuracy of a new intraoperative cone beam CT imaging technique (Artis zeego II) compared to postoperative CT scan for assessment of pedicle screws placement and breaches detection. *Eur Spine J* 26(11):2906–2916
6. Dabravolski A, Batenburg KJ, Sijbers J (2014) Dynamic angle selection in x-ray computed tomography. *Nucl Instrum Methods Phys Res Sect B: Beam Interact Mater Atoms* 324:17–24
7. Deyo RA, Nachemson A, Mirza SK (2004) Spinal-fusion surgery—the case for restraint. *Spine J* 4(5):138–142
8. Feldkamp LA, Davis LC, Kress JW (1984) Practical cone-beam algorithm. *JOSA A* 1(6):612–619
9. Fritzell P, Hägg O, Nordwall A (2003) Complications in lumbar fusion surgery for chronic low back pain: comparison of three surgical techniques used in a prospective randomized study. A report from the Swedish lumbar spine study group. *Eur Spine J* 12(2):178–189
10. Fu TS, Wong CB, Tsai TT, Liang YC, Chen LH, Chen WJ (2008) Pedicle screw insertion: computed tomography versus fluoroscopic image guidance. *Int Orthop* 32(4):517–521
11. Gang GJ, Siewerdsen JH, Stayman JW (2020) Non-circular CT orbit design for elimination of metal artifacts. In: Chen GH, Bosmans H (eds) *Medical imaging 2020: physics of medical imaging*, vol 11312. International Society for Optics and Photon-

- ics (SPIE), Bellingham, pp 531–536. <https://doi.org/10.1117/12.2550203>
12. Gang GJ, Stayman JW, Zbijewski W, Siewerdsen JH (2014) Task-based detectability in CT image reconstruction by filtered backprojection and penalized likelihood estimation. *Med Phys* 41(8Part1):081902
 13. Garber ST, Bisson EF, Schmidt MH (2012) Comparison of three-dimensional fluoroscopy versus postoperative computed tomography for the assessment of accurate screw placement after instrumented spine surgery. *Glob Spine J* 2(2):095–098
 14. Gelalis ID, Paschos NK, Pakos EE, Politis AN, Arnaoutoglou CM, Karageorgos AC, Ploumis A, Xenakis TA (2012) Accuracy of pedicle screw placement: a systematic review of prospective in vivo studies comparing free hand, fluoroscopy guidance and navigation techniques. *Eur Spine J* 21(2):247–255
 15. Kondo A, Hayakawa Y, Dong J, Honda A (2010) Iterative correction applied to streak artifact reduction in an X-ray computed tomography image of the dento-alveolar region. *Oral Radiol* 26(1):61–65
 16. Liao H, Lin WA, Huo Z, Vogelsang L, Sehnert WJ, Zhou SK, Luo J (2019) Generative Mask Pyramid Network for CT/CBCT Metal Artifact Reduction with Joint Projection-Sinogram Correction. In: International conference on medical image computing and computer-assisted intervention. Springer, pp 77–85
 17. Manbachi A (2016) Towards ultrasound-guided spinal fusion surgery. Springer, Berlin
 18. Meilinger M, Schmidgunst C, Schütz O, Lang EW (2011) Metal artifact reduction in cone beam computed tomography using forward projected reconstruction information. *Z Med Phys* 21(3):174–182
 19. Milletari F, Birodkar V, Sofka M (2019) Straight to the point: Reinforcement learning for user guidance in ultrasound. In: Smart ultrasound imaging and perinatal, preterm and paediatric image analysis. Springer, pp 3–10. https://doi.org/10.1007/978-3-030-32875-7_1
 20. Preuhs A, Manhart M, Roser P, Stimpel B, Syben C, Psychogios M, Kowarschik M, Maier A (2020) Deep autofocus with cone-beam CT consistency constraint. In: Bildverarbeitung für die Medizin 2020. Springer, pp 169–174
 21. Schulze R, Heil U, Groß D, Bruellmann D, Dranischnikow E, Schwanecke U, Schoemer E (2011) Artefacts in CBCT: a review. *Dentomaxillofacial Radiol* 40(5):265–273
 22. Schulze RKW, Berndt D, d’Hoedt B (2010) On cone-beam computed tomography artifacts induced by titanium implants. *Clin Oral Implants Res* 21(1):100–107
 23. Sembrano JN, Polly DW, Ledonio CGT, Santos ERG (2012) Intra-operative 3-dimensional imaging (O-arm) for assessment of pedicle screw position: Does it prevent unacceptable screw placement? *Int J Spine Surg* 6:49–54
 24. Simonyan K, Zisserman A (2014) Very deep convolutional networks for large-scale image recognition. arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556)
 25. Stayman JW, Siewerdsen JH (2013) Task-based trajectories in iteratively reconstructed interventional cone-beam CT. In: Proceedings 12th international meeting on fully three-dimensional image reconstruction in radiology and nuclear medicine, pp 257–260
 26. Thomsen K, Christensen FB, Eiskjær SP, Hansen ES, Fruensgaard S, Bünger CE (1997) The effect of pedicle screw instrumentation on functional outcome and fusion rates in posterolateral lumbar spinal fusion: A prospective randomized clinical study. *Spine* 22(24):2813–2822
 27. Unberath M, Zaech JN, Gao C, Bier B, Goldmann F, Lee SC, Fotouhi J, Taylor R, Armand M, Navab N (2019) Enabling machine learning in X-ray-based procedures via realistic simulation of image formation. *Int J Comput Assist Radiol Surg* 14(9):1517–1528
 28. Unberath M, Zaech JN, Lee SC, Bier B, Fotouhi J, Armand M, Navab N (2018) DeepDRR—a catalyst for machine learning in fluoroscopy-guided procedures. In: International conference on medical image computing and computer-assisted intervention. Springer, pp 98–106
 29. Van Aarle W, Palenstijn WJ, Cant J, Janssens E, Bleichrodt F, Dabrovolski A, De Beenhouwer J, Batenburg KJ, Sijbers J (2016) Fast and flexible X-ray tomography using the ASTRA toolbox. *Opt Express* 24(22):25129–25147
 30. Van Aarle W, Palenstijn WJ, De Beenhouwer J, Altantzis T, Bals S, Batenburg KJ, Sijbers J (2015) The ASTRA toolbox: a platform for advanced algorithm development in electron tomography. *Ultramicroscopy* 157:35–47
 31. Vogel J, Lasser T, Gardiazabal J, Navab N (2013) Trajectory optimization for intra-operative nuclear tomographic imaging. *Med Image Anal* 17(7):723–731
 32. Weiss AJ, Elixhauser A, Steiner C (2013) Readmissions to US hospitals by procedure, 2010. In: Healthcare cost and utilization project (HCUP) statistical brief, vol 154. Agency for Healthcare Research and Quality (US). <https://www.ncbi.nlm.nih.gov/books/NBK154387/>
 33. Wu P, Sheth N, Sisniega A, Uneri A, Han R, Vijayan R, Vagdari P, Kreher B, Kunze H, Kleinszig G, Vogt S, Lo SF, Theodore N, Siewerdsen JH (2020) Method for metal artifact avoidance in C-Arm cone-beam CT. In: Chen GH, Bosmans H (eds) Medical Imaging 2020: Physics of Medical Imaging, vol 11312. International Society for Optics and Photonics (SPIE), Bellingham, pp 522–530. <https://doi.org/10.1117/12.2549840>
 34. Zaech JN, Gao C, Bier B, Taylor R, Maier A, Navab N, Unberath M (2019) Learning to avoid poor images: Towards task-aware C-arm cone-beam CT trajectories. In: International conference on medical image computing and computer-assisted intervention. Springer, pp 11–19
 35. Zhang Y, Zhang L, Zhu XR, Lee AK, Chambers M, Dong L (2007) Reducing metal artifacts in cone-beam CT images by preprocessing projection data. *Int J Radiat Oncol Biol Phys* 67(3):924–932
 36. Zheng Z, Mueller K (2011) Identifying sets of favorable projections for few-view low-dose cone-beam CT scanning. In: 11th International meeting on fully three-dimensional image reconstruction in radiology and nuclear medicine

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.