Task-driven source–detector trajectories in cone-beam computed tomography: I. Theory and methods

J. Webster Stayman
Sarah Capostagno
Grace J. Gang
Jeffrey H. Siewerdsen

Task-driven source–detector trajectories in cone-beam computed tomography: I. Theory and methods

J. Webster Stayman,^{a,*†} Sarah Capostagno,^{a,†} Grace J. Gang,^a and Jeffrey H. Siewerdsen^{a,b}

^aJohns Hopkins University, Department of Biomedical Engineering, Baltimore, Maryland, United States

^bJohns Hopkins University, Department of Radiology and Radiological Science, Baltimore, Maryland, United States

Abstract. We develop a mathematical framework for the design of orbital trajectories that are optimal to a particular imaging task (or tasks) in advanced cone-beam computed tomography systems that have the capability of general source–detector positioning. The framework allows various parameterizations of the orbit as well as constraints based on imaging system capabilities. To accommodate nonstandard system geometries, a model-based iterative reconstruction method is applied. Such algorithms generally complicate the assessment and prediction of reconstructed image properties; however, we leverage efficient implementations of analytical predictors of local noise and spatial resolution that incorporate dependencies of the reconstruction algorithm on patient anatomy, x-ray technique, and geometry. These image property predictors serve as inputs to a task-based performance metric defined by detectability index, which is optimized with respect to the orbital parameters of data acquisition. We investigate the framework of the task-driven trajectory design in several examples to examine the dependence of optimal source–detector trajectories on the imaging task (or tasks), including location and spatial-frequency dependence. A variety of multitask objectives are also investigated, and the advantages to imaging performance are quantified in simulation studies. © 2019 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.JMI.6.2.025002](https://doi.org/10.1117/1.JMI.6.2.025002)]

Keywords: task-driven imaging; cone-beam computed tomography; robotic C-arm; imaging task; detectability index; task function; model-based image reconstruction; optimization; image quality; interventional imaging.

Paper 18251R received Nov. 12, 2018; accepted for publication Mar. 29, 2019; published online May 2, 2019.

1 Introduction

Cone-beam computed tomography (CBCT) finds application in many clinical interventions to provide up-to-date imaging (often in conjunction with hardware/instrumentation) for improved localization and assessment of treatment delivery.^{1–5} Intraoperative CBCT imaging platforms are often based on C-arm devices with flat-panel detectors that are capable of radiography and fluoroscopy in addition to the three-dimensional (3-D) CBCT. Advanced versions of these systems are motorized on several axes to permit nonisocentric source–detector orbits, to expand the field of view (FOV) in 3-D image reconstructions, and to quickly reposition between specific radiographic/fluoroscopic views. Such systems include floor- and ceiling-mounted C-arms as well as robotic C-arms⁶ with many degrees of motion freedom. In addition to C-arm systems, some CBCT mammography⁷ and body imaging systems⁸ are also capable of complex source–detector trajectories.

The additional flexibility provided by these systems permits more general orbits beyond the traditional circular and helical source–detector trajectories that have been the norm for CT for decades. To date, these flexible orbits have mainly been used to address the FOV and the sampling issues in interventional CBCT. For example, noncircular trajectories have been used to provide extended axial⁹ and elliptical¹⁰ FOVs and to improve 3-D sampling and data completeness^{11–14} to avoid

cone-beam artifacts that arise from traditional circular cone-beam orbits.

Tilted circular orbits are commonly used for their ability to positively impact image quality and/or localization. For example, tilting the CT gantry relative to the patient's longitudinal axis improves the image quality adjacent to the skull base,¹⁵ reduces eye lens dose,¹⁶ improves localization in CT-guided biopsies,^{17,18} and reduces metal artifacts associated with prostheses.¹⁹ These examples suggest that modifications of the orbit beyond simple tilts may also provide clinical advantages. However, selection of an optimal trajectory presents many challenges. For example, the above-mentioned simple tilt examples depend on the patient anatomy and/or the interventional procedure, e.g., aligning the gantry along the canthomeatal line, to the axis of the biopsy needle, etc. Thus, optimal trajectories are likely patient- and imaging-task-dependent. Moreover, data acquired from noncircular trajectories can be difficult to reconstruct, as the sampling conditions for traditional filtered-backprojection methods no longer apply.

There is a growing trend in the use of task-based measures in performance assessment.^{20–22} Such measures have also been used in prospective task-driven optimization of system design,^{23–25} regularization of model-based reconstruction,²⁶ and CT data acquisition parameters like dual-energy imaging,²⁷ tube current modulation,²⁸ and fluence-field modulation.²⁹ Thus, we expect that task-based performance models similarly provide a basis for optimizing the source–detector orbit in CBCT.

*Address all correspondence to J. Webster Stayman, E-mail: web.stayman@jhu.edu

†J. Webster Stayman and Sarah Capostagno contributed equally to this work.

Interventional imaging presents an ideal opportunity to customize orbits to the patient and diagnostic task for a number of reasons: (1) most patients in interventional imaging have previous diagnostic imaging studies, meaning that a detailed representation of the patient anatomy is available; (2) additional information, including surgical plans, the location and sizes of implants or tools, and particular anatomical targets, is known prior to the intervention; and (3) interventional imaging tasks tend to be well defined, including the particular volume of interest and specific image features that need to be identified or localized.

In this work, we present a mathematical model that leverages prior knowledge, including anatomical dependencies, to predict the imaging task performance for different trajectories. This predictor is integrated into an optimization framework that seeks the scan trajectory that maximizes performance. This framework is applied in simulation studies illustrating behavior and performance in simple digital phantoms. This work is built on previous reports on task-driven orbits,^{30–32} with a comprehensive introduction to the framework as well as an improved implementation and application to optimization for multiple tasks varying in location and/or spatial-frequency content.

2 Methods for Task-Driven Trajectory Optimization

2.1 Overview and Proposed Imaging Workflow

CBCT image quality can be highly dependent on patient size, anatomical site, and the presence of interventional hardware in the FOV. Even within a scan for a single patient, the data fidelity can vary widely with orders of magnitude differences in the noise for different measurements. This work seeks to develop a new imaging workflow in which CBCT scans are driven by the particular patient anatomy, selecting projections that maximize data fidelity for a specific imaging task. In general, this requires some knowledge of patient anatomy for prospective trajectory design. Such information is often available in interventional imaging but is typically not used directly by the imaging device.

Figure 1 illustrates a proposed imaging workflow that leverages preoperative imaging and planning data and contrasts the proposed methodology with a conventional workflow. Conventionally, a patient's diagnosis is obtained via CT (or other modality) to define and plan the particular interventional approach. In many procedures involving implantation of exogenous devices, this includes specification of the particular hardware required for the procedure (e.g., device size, location for deployment, and model number). Minimally invasive treatment increasingly relies on intraoperative imaging or immediate postoperative imaging for assessment of the surgical product and identification of possible complications. Intraoperative imaging has the advantage of detecting complications while the patient is still in the operating room with better opportunity to revise if necessary. Image quality for intraoperative and postoperative assessment is often challenged by the surgical tools and implants delivered during the procedure—often metallic and/or high density. Most importantly, image quality often suffers most in the vicinity of the implant, which, unfortunately, is the part of the image where complications are most likely to be found. In the conventional workflow, the imaging system ignores the wealth of knowledge about the patient anatomy, planned hardware delivery, and specifics of the imaging task. In contrast, a task-driven imaging workflow leverages this information to improve performance for pertinent imaging tasks associated with the particular intervention.

An overview of the task-driven optimization is illustrated in Fig. 2. The approach combines an anatomical model of the patient (as well as planning information), which is important for predicting the fidelity of the projection data, with an imaging system model, which accounts for the entire imaging chain, including any parameters to be optimized (e.g., source–detector trajectory). These models incorporate prediction of imaging performance, including local spatial resolution and noise in the reconstructed image volume. These measures of imaging performance may then be used to compute task performance using a particular observer model. With the ability to model the end-to-end system from data collection to observer performance, various parameters may be tuned in an iterative process to find the optimal source–detector trajectory that maximizes the imaging

Fig. 1 Illustration of prospective task-driven imaging using a robotic C-arm. Preoperative data from a diagnostic CT scan may be used for both preoperative planning (e.g., defining the interventional approach) and prospective design of the intraoperative scan. That is, an initial CT may be used to define the location of interest and the anatomical model used in predicting data characteristics like noise in projection data. The proposed task-driven workflow (white arrows) integrates knowledge, which is conventionally ignored (gray arrows), directly into acquisition design with the goal of optimized performance.

Fig. 2 Illustration of the task-driven optimization framework.

performance. The modeling and predictive framework in Fig. 2 are detailed in the following sections.

2.2 Task-Based Performance Prediction and Optimization

Various mathematical observer models have been used for performance prediction. The authors have previously demonstrated basic agreement between human observer performance and a non-prewhitening observer model over a broad range of imaging conditions³³ and have therefore elected to use this model to evaluate imaging tasks. Alternative metrics could be considered in future work, e.g., a prewhitening observer model to examine fundamental signal and noise content or a channelized Hotelling observer model to potentially capture aspects more closely related to a human observer. A metric based on large-area transfer characteristics (e.g., contrast-to-noise ratio) captures only the low-frequency performance and would likely miss aspects related to spatial resolution and frequency response. The non-prewhitening model chosen here derives detectability in terms of the local spatial resolution and noise properties of the reconstructed image as well as a task function that specifies the spatial frequencies of interest. The detectability index for the non-prewhitening observer model may be written as

$$d_j'^2(\Omega) = \frac{\left\{ \iiint [\text{MTF}_j(\Omega) \cdot H_{\text{Task}(j)}]^2 df_x df_y df_z \right\}^2}{\iiint \text{NPS}_j(\Omega) \cdot [\text{MTF}_j(\Omega) \cdot H_{\text{Task}(j)}]^2 df_x df_y df_z}, \quad (1)$$

where MTF_j denotes the local modulation transfer function, NPS_j is the local noise-power spectrum, and $H_{\text{Task}(j)}$ is the task function describing the location of interest (j) and the spatial frequencies of interest (e.g., formed by the difference of Fourier transforms between two stimuli in a binary hypothesis test, such as signal-present versus signal-absent, or discrimination between two stimuli). The subscript j indicates the locality of the various measures, centered at voxel j within the reconstructed image. The various quantities associated with image acquisition and/or reconstruction technique are denoted generally as Ω , which in this work refers to the source-detector trajectory.

The detectability index provides an objective for the design of an optimal source-detector orbit. The most straightforward objective that seeks to optimize detectability index for a single location in the imaging volume may be written as

$$\hat{\Omega} = \arg \max_{\Omega \in \Omega_{\text{feasible}}} d_j'^2(\Omega), \quad (2)$$

where we seek the parameter set $\hat{\Omega}$ that yields the maximum detectability index constrained by the physical or practical limitations ($\Omega \in \Omega_{\text{feasible}}$). This constraint permits specification of maximum tilt angles achievable by a particular system as well as limits to avoid contact of the gantry with the patient, patient table, or other obstacles.

Although single-location optimization may be appropriate for some imaging tasks, this objective does not consider performance at any other location in the image, thus leading to solutions that may be highly optimized to a single point and sacrifice image quality everywhere else in the image. As such, we also consider multilocation objectives. There are many possible choices for a multilocation objective; however, the principal concern is how to weigh the relative importance of performance at different locations in the FOV. We explore the following three choices in this work:

Maximum mean detectability (maxi-mean)—in which the average detectability index over an ensemble of locations within a specified region of interest (ROI), j_{ROI} , is computed and maximized. Mathematically, this objective is written as

$$\hat{\Omega} = \arg \max_{\Omega \in \Omega_{\text{feasible}}} \text{mean}_{j \in j_{\text{ROI}}} \{d_j'^2(\Omega)\}. \quad (3)$$

This objective treats all performance gains equally throughout the ROI, including potential solutions where detectability is decreased in one region for a larger gain in another region.

Maximum median detectability (maxi-median)—in which the median detectability index over the regional ensemble is maximized. Mathematically, it can be written as

$$\hat{\Omega} = \arg \max_{\Omega \in \Omega_{\text{feasible}}} \text{median}_{j \in j_{\text{ROI}}} \{d_j'^2(\Omega)\}. \quad (4)$$

This is similar to the maxi-mean objective, except that it permits larger outliers. For example, if overall detectability can be increased at the cost of a significant decrease at a few locations, this would be deemed an advantage.

Maximum minimum detectability (maxi-min)—in which one seeks to achieve the highest minimum detectability over an ensemble of locations. Mathematically, it can be written as

$$\hat{\Omega} = \arg \max_{\Omega \in \Omega_{\text{feasible}}} \text{minimum}_{j \in j_{\text{ROI}}} \{d_j'^2(\Omega)\}. \quad (5)$$

This objective indicates that it is the location of minimum performance that drives the design, and that imaging

performance cannot be sacrificed in one location (in the designated region of interest) for an improvement in another location.

Both the single-location and multilocation objectives are investigated below. To estimate the solution to any of these objective functions, we used the covariance matrix adaptation evolution strategy (CMA-ES) algorithm.³⁴ In CMA-ES, a population sample is randomly drawn according to a multivariate normal distribution at each iteration. The best solutions in the population are used to estimate the local covariance matrix of the objective function in an adaptive manner. The mean, covariance matrix, step size, and evolutionary paths are then updated to generate the next population with the goal of maximizing the number of successful samples in each successive population. This is repeated until convergence, defined as the iteration beyond which changes in function evaluation are negligible. This stochastic approach is attractive, as only function evaluations are required and populations of solutions are used, which helps to find global optima in nonlinear and nonconvex objectives. Another attractive feature of CMA-ES is its improved robustness with increased function evaluations (e.g., increasing population size), which can be tuned to help avoid local optima.

2.3 Imaging System and Reconstruction Model

To use the performance objectives defined in the previous section, one must be able to predict imaging properties (viz., local spatial resolution and noise). This necessitates an imaging system model that includes all of the pertinent imaging physics and acquisition parameters as well as the reconstruction process. Starting with data acquisition, we adopt the following forward model for a CBCT system for which mean measurements are modeled as

$$\bar{y}_i = g_i \exp(-[\mathbf{A}(\Omega)\mu]_i), \quad (6)$$

where the subscript i denotes values associated with the i 'th measurement in vector y , g_i denotes a measurement gain (e.g., unattenuated x-ray fluence, and detector sensitivity), μ is a vector of attenuation values specifying the image volume, and $\mathbf{A}(\Omega)$ represents the projection operation for a particular trajectory, parameterized by the vector Ω . The specific parameterization of $\mathbf{A}$ is discussed below.

Traditional analytic reconstruction methods are challenged by orbits that deviate from standard designs (e.g., circular and helical). However, model-based reconstruction methods are straightforward to apply on unusual and even incomplete data orbits—providing the “best” possible estimates, given the data that was collected. Thus, we adopt a statistically motivated model-based reconstruction for this work. Presuming a Poisson noise model for the measurements, we may write the following penalized-likelihood estimator:

$$\hat{\mu} = \arg \max L(y; \mu) - R(\mu), \quad (7)$$

where $L(y; \mu) = \sum_{i=1}^N y_i \log(g_i \exp(-[\mathbf{A}(\Omega)\mu]_i)) - g_i \exp(-[\mathbf{A}(\Omega)\mu]_i)$ and denotes the log-likelihood function, which is a function of the pre-log-transformed measurements y , and $R(\mu) = \beta \mu^T \mathbf{R} \mu$ represents a roughness penalty used in controlling the noise in the reconstruction with a strength parameter β . In this work, we focus on a quadratic penalty that is specified by the matrix $\mathbf{R}$, a constant matrix that defines how voxels are combined and penalized such that $R(\mu) = \frac{1}{2} \sum_j \sum_k w_{j,k} (\mu_j - \mu_k)^2$, where $w_{j,k} = 1$ for the six nearest neighbors in 3-D space and

0 otherwise. Although there are many more sophisticated regularization schemes, this particular choice of roughness penalty is well suited to previously developed imaging performance predictors.³⁵

Another advantage of the penalized-likelihood framework is that arbitrary source–detector trajectories may be reconstructed without modifying the underlying algorithm used in solving Eq. (2). In contrast, many direct reconstruction approaches will not implicitly handle noncircular or nonhelical trajectories without substantial modification or rederivation of the algorithm. In this work, we have solved Eq. (7) iteratively using the ordered subsets, separable quadratic surrogate approach discussed in Ref. 36. At each iteration, all voxels are updated simultaneously, requiring one forward-projection and one back-projection to compute the likelihood gradient. The penalty gradient and curvature are computed directly from the image. In experiments described below, the penalized-likelihood estimator is run to convergence using a specific number of iterations and a zero image has been used for initialization.

2.4 Parameterization of the Source–Detector Trajectory

As described above, the system geometry associated with a particular source–detector trajectory was parameterized by the vector Ω . Parameterization of the trajectory can take many possible forms that depend on the particular capabilities of the imaging system. In this work, we considered two parameterizations of the orbit and concentrated on orbits that sample a sphere around a common center of rotation whose x-ray source positions are specified by the coordinate pair (θ, ϕ) with rotation angle θ and tilt angle ϕ . Other geometrical parameters, such as source–detector distance and translations, remained fixed in the current studies. We chose orbits that are continuous functions of the rotation angle θ such that the gantry tilt ϕ is defined by the rotation angle.

A simple parameterization involves *periodic basis functions* using constant, sine, and cosine terms such that

$$\begin{aligned} \phi(\theta) &= \sum_{i=1}^K \Omega_i b_i(\theta) \quad b_1(\theta) = 1, \quad b_2(\theta) = \sin \theta, \\ b_3(\theta) &= \cos \theta, \quad b_4(\theta) = \sin 2\theta, \dots, \end{aligned} \quad (8)$$

as shown in Fig. 3. Both short scans and full 360 deg orbits may be defined using this scheme with, for example, uniform sampling of the rotation angle θ . This parameterization provides a low-dimensional space to perform the orbital optimization. Because there are practical limitations to the orbit (e.g., collision with the table or patient), we apply hard constraints on the orbit by limiting the maximum tilt angle to ± 50 deg.

A second parameterization of the source–detector trajectory uses *B-spline basis functions*, where the individual parameters Ω_i define a limited set of knot locations. Each knot is fixed to a single rotation angle with equal spacing throughout the orbit and is allowed to vary in tilt angle. This parameterization may more easily admit nonperiodic designs while maintaining relatively low dimensionality. The orbit may be similarly constrained to those feasible with a given C-arm gantry and to avoid table collision. Using either the periodic or B-spline basis functions as presented here not only reduces the dimensionality of the parameterization but also imposes the additional constraint that the trajectory be smoothly changing, which is beneficial for

Fig. 3 Illustration of a short-scan orbit parameterized by constant, sine, and cosine basis functions.

practical implementation of a task-driven orbit from a mechanical point of view.

2.5 Imaging Performance Prediction and Anatomical Modeling

For prospective design of source–detector trajectories, the imaging performance must be estimated for various orbits. Although exhaustive simulation of projection data, reconstruction, and assessment is possible, it is more practical to estimate the imaging properties of the reconstruction directly. Previous work derived closed-form approximations for local spatial resolution³⁵ and noise³⁷ in penalized-likelihood reconstructions of the type in Eq. (7). Specifically, the local impulse response, l_j , and local covariance, c_j , may be approximated as

$$l_j \approx [\mathbf{A}(\Omega)^T \mathbf{D}\{\bar{y}\} \mathbf{A}(\Omega) + \beta \mathbf{R}]^{-1} \mathbf{A}(\Omega)^T \mathbf{D}\{\bar{y}\} \mathbf{A}(\Omega) e_j, \quad (9)$$

$$c_j \approx [\mathbf{A}(\Omega)^T \mathbf{D}\{\bar{y}\} \mathbf{A}(\Omega) + \beta \mathbf{R}]^{-1} [\mathbf{A}(\Omega)^T \mathbf{D}\{\bar{y}\} \mathbf{A}(\Omega)] \times [\mathbf{A}(\Omega)^T \mathbf{D}\{\bar{y}\} \mathbf{A}(\Omega) + \beta \mathbf{R}]^{-1} e_j, \quad (10)$$

where e_j denotes the j 'th unit vector which is all zeros except for the j 'th location which is unity, and $\mathbf{D}\{\cdot\}$ denotes the operator that places its vector argument on the diagonal of a matrix. Note that Eqs. (9) and (10) capture the various dependencies of the reconstructed image on system geometry $[\mathbf{A}(\Omega)]$, regularization ($\beta \mathbf{R}$), location (j), and the patient anatomy via projections (y) of the attenuation distribution (μ_{prior}) via Eq. (6). A recent work³⁸ extended these predictors to nonideal detectors and validated the predictions in flat-panel CBCT reconstructions.

Thus, we may predict local noise and resolution properties prospectively, given knowledge of μ_{prior} and a system model through which y may be simulated. As previously discussed, preoperative CT (or CBCT acquired earlier in the procedure) can provide an anatomical model for prediction, recognizing that the model may be mismatched to the intraoperative data for a number of reasons. First is registration of the previous image to the current measurements. For many anatomical sites (e.g., in neurointerventions), a rigid registration may be sufficient, allowing the designed trajectory to be transformed into the intraoperative patient coordinates using, for example, 3-D to two-dimensional (2-D) registration, as shown in Ref. 39. In this work, we will presume that an accurate registration is

achieved and will focus on the subsequent improvements in imaging performance gained from a task-driven design of the orbit.

Other mismatches in patient anatomy include the delivery of hardware and changes in anatomy that might be found in the intraoperative data. Again, the workflow presented in Fig. 1 offers a means to model implanted hardware. As the preoperative scan is often used for planning, e.g., to determine the size and location of an implant, it is relatively straightforward to include these attenuation changes in a modified anatomical model. Modeling of significant attenuation changes like metal implants is important as these changes have a significant impact on the statistics of the data (i.e., the diagonal weighting $\mathbf{D}\{\bar{y}\}$ in the predictors). In contrast, soft-tissue differences like hemorrhage (an important complication one would like to detect) have a relatively small effect on noise in the projection data. This suggests that those more subtle changes do not need to be modeled explicitly for the proposed trajectory design.

2.6 Approximate Predictors and Practical Implementation

Although the predictors in Eqs. (9) and (10) describe the basic imaging performance metrics required for evaluation of detectability index, these expressions contain large matrix inverses that challenge efficient computation of the local spatial resolution and noise for such a large optimization space. Previous work used local Fourier approximation^{40,41} to yield approximate forms for the local MTF _{j} and local NPS _{j} as

$$\text{MTF}_j = \mathcal{F}\{l_j\} \approx \frac{\mathcal{F}\{\mathbf{A}(\Omega)^T \mathbf{D}\{\bar{y}\} \mathbf{A}(\Omega) e_j\}}{\mathcal{F}\{\mathbf{A}(\Omega)^T \mathbf{D}\{\bar{y}\} \mathbf{A}(\Omega) e_j + \beta \mathbf{R} e_j\}}, \quad (11)$$

$$\text{NPS}_j = \mathcal{F}\{c_j\} \approx \frac{\mathcal{F}\{\mathbf{A}(\Omega)^T \mathbf{D}\{\bar{y}\} \mathbf{A}(\Omega) e_j\}}{|\mathcal{F}\{\mathbf{A}(\Omega)^T \mathbf{D}\{\bar{y}\} \mathbf{A}(\Omega) e_j + \beta \mathbf{R} e_j\}|^2}, \quad (12)$$

where $\mathcal{F}$ denotes the 3-D discrete Fourier transform and the division is element by element. Although Eqs. (11) and (12) eliminate the computationally expensive matrix inverse, these expressions still require repeated computations involving projection, backprojection, and Fourier transforms. Following Ref. 42, a number of observations permit additional speedups for practical implementation. In particular, one may compute the regularization term once for a shift-invariant penalty,

and only local volumes within a $N \times N \times N$ region of interest of the voxel j are required. Repeated calculation of the Fourier transform of the weighted projection-backprojection ($\mathcal{F}\{\mathbf{A}(\Omega)^T \mathbf{D}\{\bar{\mathbf{y}}\} \mathbf{A}(\Omega) e_j\}$) represents a significant computational burden in direct computations of Eqs. (11) and (12). However, efficient calculation is possible by leveraging the observation, as described in Ref. 42, that the projection-backprojection term is linear in the diagonal weighting. Specifically, we may write

$$\mathcal{F}\{\mathbf{A}(\Omega)^T \mathbf{D}\{w\} \mathbf{A}(\Omega) e_j\} = \mathbf{L}_j w, \quad (13)$$

which means that a linear operator, $\mathbf{L}_j$, may be precomputed and stored for fast application. In addition, noting that the projection of a single point $\mathbf{A}(\Omega) e_j$ is very sparse with nonzero values located at about one point per projection view, one may precompute a much smaller $\mathbf{L}_j$ that is only $N^3 \times N_{\text{views}}$ for each location j . One may form $\mathbf{L}_j$ by substituting $w = e_i$ for $i = 1$ to N_{views} to build up $\mathbf{L}_j$ column by column. Thus, as long as the number of potential view angles is sufficiently small (e.g., one needs to store N_{views} Fourier volumes of $N \times N \times N$), this *precomputation approach* becomes tractable without additional modification.

For a large number of potential views, memory limitations may challenge the precomputation approach. In the spirit of Refs. 43 and 44, one can also recognize that we may replace $\mathbf{L}_j$ with an analytic form. That is, each column of $\mathbf{L}_j$ is the Fourier transform of the backprojection of a single projection view consisting of a single point projection. The backprojection of a point is a line in 3-D space through the original location j connecting both the source and detector. The Fourier transform of this line is a Fourier plane centered at the origin at the same angle as the line. However, for a discrete system model with finite-sized detector elements, voxels, etc., this Fourier plane will not be infinitely thin. Whereas Schmitt derived closed-form analytic expressions that specify the form and profile through this plane, we opted to use a plane whose profile is Gaussian with a width specified by fitting an empirical calculation of the Fourier-domain projection-backprojection. This results in a very fast analytic form for $\mathbf{L}_j$ that may be computed, given the location j and the coordinates of the x-ray source. This *on-the-fly computation approach* is another efficient alternative for calculating Eqs. (11) and (12). The difference between the precomputation and on-the-fly approaches is largely computational (without significant difference in noise and resolution estimates) and represents a classic computing trade-off between storage and speed. As both approaches are potentially useful, both are presented in studies below.

3 Experimental Methods

Three experiments were conducted to investigate various aspects of the proposed task-driven trajectory design process, summarized in Fig. 4. Each experiment is detailed in the following subsections.

Two simulation configurations were used in these studies. Studies that investigated basic trajectory design behavior used a compact system geometry to better illustrate dependencies on location and task, specifically 700-mm source-to-detector distance and 350-mm source-to-axis distance. A flat-panel detector with 560×1000 pixels at 1-mm pitch was simulated. This geometry is referred to as the *compact geometry*. A second system geometry emulated a *C-arm geometry* with 1200-mm

Fig. 4 Summary of the three simple scenarios studied. Cases are distinguished by the stimulus and anatomical model (top row) and the frequency-domain task function (bottom row). Different task functions are explored, including Gaussian detection, line-pair discrimination, and midfrequency discrimination.

source-to-detector distance, 800-mm source-to-axis distance, and a 960×1240 pixel detector at 0.308-mm pitch.

3.1 Location Dependence: A Sphere in a Cylinder

The first experiment used a simple object to illustrate location dependence. Specifically, a 20-cm diameter cylinder was simulated with 1-mm voxels and an attenuation of 0.05 mm^{-1} [Fig. 4(a)]. A relatively high attenuation was used to exaggerate the location-dependent effects. Two 3-mm spheres with 0.03 mm^{-1} contrast relative to background were added to the cylinder, centered in the axial slice, at the central slice and at 15 cm below the central slice. This study used the compact geometry, and bare-beam fluence was set at 10^5 photons per pixel.

Trajectory design was conducted using the single-location objective in Eq. (2) using nine periodic basis functions given in Eq. (8), constrained to tilts in the range of $\phi = -50$ deg to 50 deg, and with rotation angles $\theta = 1$ deg to 360 deg with 1 deg increments. The task function in the objective corresponded to the 3-mm spherical stimulus [Fig. 4(d)] and optimization was performed for each stimulus location individually. For computation of detectability, the on-the-fly method was applied with statistical weights (w , calculated as described in Sec. 2.6) sampled over 110 equally spaced rotation angles and 51 equally spaced tilt angles. To estimate the solution of Eq. (2), the CMA-ES optimization was applied using a population size of 40 without restarts due to the simplicity of the search space.

Penalized-likelihood reconstruction using the optimized trajectory found for each sphere location (central slice and 15 cm below the central slice) was performed using dynamically relaxed ordered subsets with the number of subsets changing every five iterations through the sequence $\{54, 24, 12, 6, 4, 2, 1\}$ for a total of 50 iterations. This schedule was chosen to accelerate the convergence of the simple object. Quadratic regularization with a regularization strength of $\beta = 10^6$ was applied using 1-mm isotropic voxels on a $240 \times 240 \times 500$

grid. This choice of regularization strength was empirically selected based on visual assessment of noise and resolution in the images.

3.2 Task Dependence: Line Pairs in a Cylinder

The second study used the same 20-cm cylinder as the first experiment—in this case involving a line-pair stimulus placed in the center of the central slice. Specifically, the line pairs were a $20 \times 20 \times 20 \text{ mm}^3$ cube with 2 lp/mm and a contrast of 0.015 mm^{-1} relative to background. The cube was rotated 20 deg around the x-axis to angulate the features relative to the coordinate axes [Fig. 4(b)].

Trajectory design was performed using the single-location objective, nine periodic bases, and the same angular sampling, constraints, geometry, and bare-beam fluence as in the first experiment. Two task functions were examined: the 3-mm sphere detection task from the first experiment and a line-pair discrimination task. The discrimination task was defined as the difference between the true rotated line-pair stimulus and a stimulus consisting of an identically rotated solid cube of the same dimensions and attenuation, i.e., a binary hypothesis of line pairs “present” or “absent” [Fig. 4(e)]. As in the first experiment, on-the-fly detectability computations were adopted and the same CMA-ES parameters were used. Penalized-likelihood ordered-subset reconstructions were computed as in the first experiment, using the {54, 24, 12, 6, 4, 2, 1} schedule, 50 iterations to reach convergence, $\beta = 10^6$, and 1-mm isotropic voxels on a $240 \times 240 \times 500$ grid.

3.3 Multiple Locations: Elliptical Cylinder with Needle

The third experiment examined task-driven imaging in a situation for which the precise task location was unknown. We therefore chose several potential task locations and solved for a single trajectory that maximized a multilocation objective. A digital elliptical cylinder phantom was modeled with major and minor axis diameters of 25 and 17.5 cm, respectively, and with a height of 25 cm. The elliptical cylinder had a 12-mm outer shell with attenuation coefficient of 0.04 mm^{-1} and the core was filled with low-contrast spheres of diameter 80 mm and attenuation coefficients ranging from 0.0175 to 0.0225 mm^{-1} . A high-contrast cylinder (attenuation coefficient 0.2 mm^{-1}) with a diameter of 5 mm and a length of 75 mm was added at a 10-deg angle in the central coronal plane to simulate a needle entering the body [shown in Fig. 4(c)]. Nine 8-mm spheres with 0.08 mm^{-1} contrast were added along the length and tip of the cylinder and their nine locations were used for multilocation optimization. We used the C-arm geometry with bare-beam fluence lowered to 10^4 photons per pixel to mimic scanning at lower dosage.

Trajectory design was first performed for each task location using the single-location objective in Eq. (2), followed by the three multilocation objectives in Eqs. (3)–(5). The B-spline basis functions were used with eight equally spaced knots, and the trajectories were constrained to tilt angles in the range of $\phi = -30$ deg to 30 deg and rotation angles of $\theta = 1$ deg to 360 deg. The task function at all nine locations consisted of midfrequency content [Fig. 4(f)] to discriminate the presence of a single object or two separate objects (needle and sphere), as described in the ICRU Report 54.⁴⁵ The precomputation approach was used with w sampled over 72 equally

spaced rotation angles and 13 equally spaced tilt angles to compute detectability. The CMA-ES optimization algorithm was used to estimate $\hat{\Omega}$ using a population size of 200 with six restarts and random initialization to ensure that the optimal detectability values were achieved due to the more complicated search space compared to the previous two experiments.

Quadratic penalized-likelihood reconstructions for the nominal circular orbit, nine independent single-location optimizations, and three variations of multilocation optimization were performed using 10 ordered subsets over 200 iterations. The total number of iterations was increased and the number of subsets was decreased from previous experiments to achieve stability in convergence of the structurally more complicated object. Regularization strength β was set at $10^{5.5}$ to decrease noise in the highly attenuating object, and 0.5-mm isotropic voxels were used in a $512 \times 512 \times 512$ grid.

To illustrate the convergence of the CMA-ES algorithm, the optimization for a single sphere was performed as described above, using a circular orbit for initialization as opposed to random initialization. This allowed comparison of the detectability index and the reconstructed images associated with the initial circular orbit, an intermediate (suboptimal) orbit, and the resulting optimal orbit.

4 Results

4.1 Location Dependence: A Sphere in a Cylinder

Results from the first experiment are summarized in Fig. 5. Task-driven source trajectories are shown for the cylindrical object and two optimizations: maximum detectability of the spherical stimulus at locations #1 [Fig. 5(b)] and #2 [Fig. 5(c)]. Source orbits are shown for each case in magenta as well as a sampled sphere of all possible source locations (blue dots). Note that the designed orbits suggest that a simple equatorial source trajectory around each stimulus is optimal for detection of that stimulus. Reconstructions using the two orbits confirm that conspicuity of each sphere is maximized for such an equatorial orbit, with reduced detectability for the stimulus in the location that is not optimized. That is, the stimulus at location #2 is more difficult to detect when using the orbit designed to maximize detectability at location #1, and vice versa.

This simple, idealized simulation offers basic intuition to the optimization by considering the fluence through each stimulus location for all potential views. This fluence is equal to the statistical weighting— w in Eq. (13)—and we see that the orbit (dashed magenta line) in each case maximizes fluence through the target location, thereby maximizing data fidelity by selecting projection views with the shortest path length through the object. This experiment illustrates the importance of location in task-driven designs and shows that using a single-location objective can maximize task performance at a given location but may do so at the cost of decreased performance at other locations.

4.2 Task Dependence: Line Pairs in a Cylinder

The results for a task with strong directional dependence (a line-pair stimulus within a cylinder) are summarized in Fig. 6. Single-location task-driven designs are performed for two tasks: the sphere detection task from Sec. 3.1 [Fig. 6(b)] and a discrimination task corresponding to the frequency content and angulation of the line-pair stimulus [Fig. 6(c)]. For the

Fig. 5 Location dependence. (a) Spherical stimuli are placed at the center (location #1) and 15 cm below (location #2) in a cylindrical object. (b) Task-driven trajectory design for the sphere detection task at location #1. (c) Task-driven design for the same task at location #2. Sample reconstructions and the fluence through each location for all potential views (analogous to the statistical weighting) are shown. Optimal orbits are identified in magenta.

task function corresponding to the line-pair stimulus, the task-driven orbit (orbit A) is tilted to match the angulation of the line-pair stimulus. Example reconstructions illustrate that the task-driven design outperforms the nominal orbit (orbit S), thus

showing increased noise between the line pairs for the trajectory optimized for the spherical task [Fig. 6(b)] and improved discrimination of line pairs for the orbit optimized according to the line-pair task [Fig. 6(c)].

Fig. 6 Task dependence. (a) Line-pair stimuli are placed at the center of a cylindrical object. (b) Task-driven trajectory design for the sphere detection task (orbit S, not optimal for this object). (c) Task-driven design for an angulated line-pair discrimination task (orbit A). Sample reconstructions and the fluence through each location for all potential views are shown. Orbits are identified in magenta.

4.3 Multiple Locations: Elliptical Cylinder with Needle

Figure 7 summarizes the task-driven trajectories resulting from optimization with respect to a single-location objective (in which d' is optimized for each location independently), and for the three variations of a multilocation objective [maxi-mean, maxi-median, and maxi-min shown in Eqs. (3)–(5), in which d' is optimized over all locations simultaneously]. The fluence through the stimulus location for all potential views is shown for each of the nine task locations with the optimal orbits shown in magenta. The orbits resulting from a single-location objective are shown in the left-most column, and the single orbit resulting from a multilocation objective is shown in the three right columns for the maxi-mean, maxi-median, and maxi-min objectives for all nine locations. The single-location objective shows a different orbit for each location, and these can be considered the best possible trajectory for each location. Note that at some of the locations the tilt angle differs for the starting and ending views (i.e., locations 4, 5, and 7). The 2-D B-spline basis function representing the source trajectory is not constrained to be equal at the start and end vertices, allowing a discontinuity. In these cases, an orbit with a discontinuity produces a higher d' value at the task location. We may theorize that allowing discontinuities at any point in the orbit (not only at start and end vertices) might increase d' ; however, doing so loses the advantage of a low-dimensional parametrization and increases the difficulty of implementation on a robotic C-arm.

When comparing the three multilocation objectives (in which a single trajectory is generated for all stimulus locations), each objective function yields a different result, as expected. The trajectory generated when using a maxi-median objective is dictated by the median d' value, which in this case corresponds to location 3 for the optimal trajectory. Note that the orbit is

not the same as that for location 3 achieved when using a single-location objective. This is because further changes to the orbit would cause all d' values to shift, which would generate a new median value that does not maximize the objective function. The maxi-min objective is similarly driven by a single location, i.e., location 1. Location 1 is at the tip of the needle and is subsequently obscured by the angulated needle to some extent at all possible views. This results in the lowest fluence, compared to all other locations for all possible views. As the d' value at location 1 is significantly smaller than the values at locations 2 to 9 due to the lower fluence, it remains the minimum d' value for all trajectories examined during optimization. This results in an equivalent optimization as when using a single-location objective at location 1, and the resulting orbit is almost identical to that case. When using a maxi-mean objective, all locations jointly drive the orbit, and no single location skews the result.

Figure 8(a) shows detectability index for the nine task locations corresponding to a circular trajectory, using a single-location objective, and using the three multilocation objective functions. The single-location objective is seen to provide the highest d' value and represents the upper limit of improvement in d' . The respective advantage of each multilocation objective function is evident, with maxi-mean having the highest mean d' value at 3.1, maxi-median having the highest median d' value at 3.7, and maxi-min having the highest minimum d' value at 1.3. Figure 8(b) shows the percentage increase (or decrease) in d' compared to the location-matched d' value for a circular orbit. For a single-location objective, all d' values increase, ranging from 5.5% at location 8 to 121.1% at location 1. In contrast, for all three multilocation objectives there are locations that exhibit a decrease in detectability. For maxi-mean, the change in d' ranges from –8.4% at location 4 to +37.7% at location 1. For maxi-median, the range is from –46.1% at

Fig. 7 Comparison of multilocation and single-location objectives. The task-driven trajectory resulting for each case is shown overlaid on the statistical weights for all source locations at each task location (1 to 9). Using a single-location objective shows the ideal task-driven trajectory for each location.

Fig. 8 Optimization for multiple task locations. (a) Boxplots showing detectability index for the nine locations for a circular orbit, a single-location objective, maxi-mean, maxi-median, and maxi-min objective functions. On each box, the central line indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. The whiskers indicate the range of the data excluding outliers, which are indicated with the “+” symbol. (b) Boxplots showing the percentage change in detectability from a nominal circular orbit after optimization at all the nine locations for the single-location objective function and the three multilocation objective functions. (c) The resulting image for each orbit in a region around each stimulus location, demonstrating improved visualization of small spheres placed adjacent to a high-contrast cylindrical “needle.” Corresponding images from a nominal circular orbit are shown for comparison in the top row.

location 4 to +45.4% at location 9. For maxi-min, the range is from −35.4% at location 3 to +120.9% at location 1.

The reconstructed image in the region around each stimulus is shown for each orbit in Fig. 8(c), as well as for a nominal circular orbit for comparison (top row). Either the axial, coronal, or sagittal plane is chosen at each location for visual representation of improvements, although the reconstructed plane could

be arbitrary. Improvements are particularly apparent for locations 2, 7, and 9.

Figure 9 demonstrates convergence of the CMA-ES algorithm for the single-location objective performed for the task defined at location 7. Images reconstructed for a circular orbit, suboptimal intermediate orbit, and optimal task-driven orbit are shown in Figs. 9(a)–9(c). The d' values for each are

Fig. 9 Convergence of the optimal orbit at location 7. (a) Image reconstructed from a circular orbit (inset below the image, plotted in magenta overlaid on a map of statistical weights). (b) Image reconstructed from a suboptimal orbit representing an intermediate solution between a circular orbit and the optimal orbit. (c) Image reconstructed from the optimal orbit for this task and object model. (d) Convergence of the CMA-ES optimization showing detectability (d') versus iteration. The d' value for the orbits shown in (a)–(c) are indicated in the plot.

2.85, 3.35, and 4.08, respectively. Visualization of the spherical stimulus alongside the highly attenuating needle is progressively improved as the algorithm converges on the optimal orbit, consistent with the increase in d' . A plot of detectability index versus iteration is shown in Fig. 9(d), which shows a fairly smooth convergence for the CMA-ES optimizer when initialized with the circular orbit. The d' values corresponding to the three orbits (A–B–C) in Figs. 9(a)–9(c) are indicated in the plot.

5 Discussion

In this work, we have presented a framework for task-driven trajectory design for advanced CBCT imaging systems that leverages knowledge of both the patient anatomy and imaging task and exercises the motion capabilities of motorized, multi-axis CBCT systems to maximize imaging performance. The task-driven approach provides a strategy to overcome traditional imaging limitations via orbital flexibility, e.g., challenges associated with highly attenuating anatomy or implants that can be mitigated through intelligent data collection, selecting the best projection views to accomplish a particular task.

The importance of both the location and spatial-frequency dependence of the imaging task has been investigated. Each of these elements can contribute significantly to which projection views carry the greatest information for a specific task. In general, the task-based approach balances the data fidelity of a view (i.e., noise) with the signal content (frequency response) that each view provides toward the imaging task. Because the trajectory design for a single stimulus location in the image volume can optimize performance at that location to the detriment of performance elsewhere, a number of multilocation design objectives are examined to obtain optimal performance over regions of interest.

The proposed framework is general with respect to various options for parameterization of the source–detector trajectory, including constraints based on system geometry, degrees of freedom, and collision avoidance—permitting application to robotic C-arms and other CBCT systems with additional flexibility to gantry tilt, etc. This includes short-scan geometries, which are popular in CBCT imaging. A short-scan can be similarly optimized as the above-described full-scans with additional parameters such as start and end rotation angles included in the optimization to seek the views that are best for the object and task. Moreover, it is straightforward to further generalize these methods with additional geometric degrees of freedom (e.g., axial translation and variation in source–detector distance), acquisition parameters (e.g., tube current modulation²⁸ and fluence-field modulation²⁹), and reconstruction parameters (e.g., regularization strength, β ²⁶). Such factors certainly carry interdependency in a full multivariate optimization over all parameters—well beyond the scope of the current work, which focuses specifically on optimization of the source–detector orbit (holding other parameters fixed).

We note that Fourier approximation of both the MTF and NPS is one potential source of error in orbital design. Specifically, this presumes that the term $\mathbf{A}^T \mathbf{D} \mathbf{A} \mathbf{e}_j$ is locally shift-invariant. For heterogeneous objects, the diagonal weighting is variable with location. However, we note that the data-dependence of this term falls between projection and backprojection operations, which have the effect of strong smoothing even for highly nonuniform weights. The impact of this smoothing is that the Fourier approximation tends to hold well even for very nonuniform objects; however, there is the potential for

breakdown of the approximation when shift variance is pushed to the same scale as the spatial resolution of the system.

Consideration should also be given to challenges of accurate CBCT image reconstruction. Physical factors that are not accounted for include x-ray scatter, beam hardening, image lag, detector glare, and patient motion, which may produce artifacts that confound visualization of the task. Recent work has investigated the optimization of orbits based on scatter and scatter-to-primary ratio for imaging the weight-bearing spine using a robotic CBCT system, indicating that these effects are small and can be included in the overall optimization.⁴⁶ Effects that occur from the sampling pattern itself, such as streakiness and cone-beam artifacts, are included in the model of the local MTF and are therefore represented in the task-driven image in such a way as to improve the detectability of the task (although they may not be eliminated from the image).

Accurate geometric calibration of noncircular orbits is also required for accurate image reconstruction. The motion of the C-arm gantry as it moves through complex task-driven orbits may result in gantry wobble. This can be accounted for using a calibration technique like the “self-calibration,” as described in Ref. 47, that performs the geometric calibration subsequent to the scan using the set of acquired projections.

Because the design objectives are generally nonconvex and require many predictions of image quality, different strategies for efficient evaluation and solution of the design objective are presented. Image quality predictors in the current work focus on quadratically penalized model-based image reconstruction; however, future work aims to extend this methodology to non-quadratic regularization.^{44,48}

Task-driven trajectory design has the potential for application in a number of imaging scenarios. This work has focused on theoretical underpinnings of the framework and has illustrated the method in a series of experiments ranging from simple to complex. Although the phantoms used in these scenarios had somewhat uniform backgrounds, the method is designed to incorporate complex anatomic variations in the optimization via the patient model. In fact, the complexity of the model drives the selection of optimal trajectory. This is demonstrated in the Part II companion paper⁴⁹ titled “Task-driven source–detector trajectories in cone-beam computed tomography: II. Application to neuroradiology,” where we investigate application to imaging scenarios in neuroradiology with complex surrounding anatomy. We expect that the methodology also has potentially broad application in other interventional imaging scenarios (e.g., orthopedic procedures). The above-detailed theoretical foundations suggest a new paradigm for interventional imaging wherein preoperative information is included explicitly within a rigorous definition of the imaging task to prospectively drive customized data acquisition that maximizes performance. The framework is an important first step in realizing advanced CBCT capabilities and more fully leveraging the wealth of information available in interventional imaging scenarios.

Disclosures

No conflicts of interest, financial or otherwise, are declared by the authors.

Acknowledgments

This research was supported by the National Institutes of Health Grant Nos. R01-EB-017226 and U01-EB-018758, academic-

industry partnership with Siemens Healthineers (AX Division, Forchheim, Germany), and a Johns Hopkins University Catalyst Grant.

References

- J. H. Siewerdsen et al., "Volume CT with a flat-panel detector on a mobile, isocentric C-arm: pre-clinical investigation in guidance of minimally invasive surgery," *Med. Phys.* **32**(1), 241–254 (2005).
- S. Hirota et al., "Cone-beam CT with flat-panel-detector digital angiography system: early experience in abdominal interventional procedures," *Cardiovasc. Intervent. Radiol.* **29**, 1034–1038 (2006).
- M. J. Daly et al., "Intraoperative cone-beam CT for guidance of head and neck surgery: Assessment of dose and image quality using a C-arm prototype," *Med. Phys.* **33**, 3767–3780 (2006).
- R. C. Orth, M. J. Wallace, and M. D. Kuo, "C-arm cone-beam CT: general principles and technical considerations for use in interventional radiology," *J. Vasc. Interv. Radiol.* **19**, 814–820 (2008).
- W. Hohenforst-Schmidt et al., "Cone beam computer tomography (CBCT) in interventional chest medicine—high feasibility for endobronchial real time navigation," *J. Cancer* **5**(3), 231–241 (2014).
- N. Binder et al., "A robotic C-arm fluoroscope," *Int. J. Med. Rob. Comput. Assisted Surg.* **1**, 108–116 (2005).
- J. P. Shah et al., "Three dimensional dose distribution comparison of simple and complex acquisition trajectories in dedicated breast CT," *Med. Phys.* **42**, 4497–4510 (2015).
- A. Fieselmann et al., "Twin robotic x-ray system for 2D radiographic and 3D cone-beam CT imaging," *Proc. SPIE* **9783**, 97830G (2016).
- Z. Yu et al., "Extended ellipse-line-ellipse trajectory for long-object cone-beam imaging with a mounted C-arm system," *Phys. Med. Biol.* **61**, 1829–1851 (2016).
- M. Herbst, "Dynamic detector offsets for field of view extension in C-arm computed tomography with application to weight-bearing imaging," *Med. Phys.* **42**, 2718–2729 (2015).
- X. Tang and R. Ning, "A cone beam filtered backprojection (CB-FBP) reconstruction algorithm for a circle-plus-two-arc orbit," *Med. Phys.* **28**, 1042–1055 (2001).
- G. L. Zeng and G. T. Gullberg, "A cone-beam tomography algorithm for orthogonal circle-and-line orbit," *Phys. Med. Biol.* **37**, 563–577 (1992).
- F. Noo et al., "The dual-ellipse cross vertex path for exact reconstruction of long objects in cone-beam tomography," *Phys. Med. Biol.* **43**, 797–810 (1998).
- H. K. Tuy, "An inversion formula for cone-beam reconstruction," *SIAM J. Appl. Math.* **43**, 546–552 (1983).
- H. G. Menzel, H. Schibilla, and D. Teunen, "European guidelines on quality criteria for computed tomography European guidelines on quality criteria," pp. 1–71, 1999 http://www.msct.info/CT_Quality_Criteria.html.
- U. Nikupaavo et al., "Lens dose in routine head CT: comparison of different optimization methods with anthropomorphic phantoms," *AJR. Am. J. Roentgenol.* **204**, 117–123 (2015).
- T. Yamagami et al., "Percutaneous needle biopsy for small lung nodules beneath the rib under CT scan fluoroscopic guidance with gantry tilt," *Chest* **126**, 744–747 (2004).
- S. Hussain, "Gantry angulation in CT-guided percutaneous adrenal biopsy," *Am. J. Roentgenol.* **166**, 537–539 (1996).
- M. Lewis et al., "CT metal artefact reduction of total knee prostheses using angled gantry multiplanar reformation," *Knee* **17**, 279–282 (2010).
- J. H. Siewerdsen and L. E. Antonuk, "DQE and system optimization for indirect-detection flat-panel imagers in diagnostic radiology," *Proc. SPIE* **3336**, 546–556 (1998).
- I. Reiser and R. M. Nishikawa, "Task-based assessment of breast tomosynthesis: Effect of acquisition parameters and quantum noise," *Med. Phys.* **37**, 1591–1600 (2010).
- S. Park et al., "A statistical, task-based evaluation method for three-dimensional x-ray breast imaging systems using variable-background phantoms," *Med. Phys.* **37**, 6253–6270 (2010).
- J. H. Siewerdsen and D. A. Jaffray, "Optimization of x-ray imaging geometry (with specific application to flat-panel cone-beam computed tomography)," *Med. Phys.* **27**(8), 1903–1914 (2000).
- P. Prakash et al., "Task-based modeling and optimization of a cone-beam CT scanner for musculoskeletal imaging," *Med. Phys.* **38**, 5612–5629 (2011).
- J. Xu et al., "Modeling and design of a cone-beam CT head scanner using task-based imaging performance optimization," *Phys. Med. Biol.* **61**, 3180–3207 (2016).
- H. Dang, J. H. Siewerdsen, and J. W. Stayman, "Prospective regularization design in prior-image-based reconstruction," *Phys. Med. Biol.* **60**, 9515–9536 (2015).
- S. Richard and J. H. Siewerdsen, "Cascaded systems analysis of noise reduction algorithms in dual-energy imaging," *Med. Phys.* **35**(2), 586–601 (2008).
- G. J. Gang et al., "Task-driven optimization of CT tube current modulation and regularization in model-based iterative reconstruction," *Phys. Med. Biol.* **62**, 4777–4797 (2017b).
- G. J. Gang, J. H. Siewerdsen, and J. W. Stayman, "Task-driven optimization of fluence field and regularization for model-based iterative reconstruction in computed tomography," *IEEE Trans. Med. Imaging* **36**, 2424–2435 (2017a).
- J. W. Stayman and J. H. Siewerdsen, "Task-based trajectories in iterative reconstructed interventional cone-beam CT," in *Int. Meeting Fully Three-Dimensional Image Reconstr. Radiol. Nucl. Med.*, pp. 257–260 (2013).
- J. W. Stayman, G. J. Gang, and J. H. Siewerdsen, "Task-based optimization of source-detector orbits in interventional cone-beam CT," in *Int. Meeting Fully Three-Dimensional Image Reconstr. Radiol. Nucl. Med.*, Vol. 13 (2015).
- S. Ouadah et al., "Task-driven orbit design and implementation on a robotic C-arm system for cone-beam CT," *Proc. SPIE* **10132**, 101320H (2017).
- G. J. Gang et al., "Analysis of Fourier-domain task-based detectability index in tomosynthesis and cone-beam CT in relation to human observer performance," *Med. Phys.* **38**, 1754–1768 (2011).
- N. Hansen and S. Kern, "Evaluating the CMA evolution strategy on multimodal test functions," *Lect. Notes Comput. Sci.* **3242**, 282–291 (2004).
- J. A. Fessler and W. L. Rogers, "Spatial resolution properties of penalized-likelihood image reconstruction: space-invariant tomographs," *IEEE Trans. Image Process.* **5**, 1346–1358 (1996).
- H. Erdogan and J. A. Fessler, "Ordered subsets algorithms for transmission tomography," *Phys. Med. Biol.* **44**, 2835–2851 (1999).
- J. A. Fessler, "Mean and variance of implicitly defined biased estimators (Such as penalized maximum likelihood): applications to tomography," *IEEE Trans. Image Process.* **5**, 493–506 (1996).
- W. Wang, G. J. Gang, and J. W. Stayman, "Spatial resolution properties in penalized-likelihood reconstruction of blurred tomographic data," in *Int. Meeting Fully Three-Dimensional Image Reconstr. Radiol. Nucl. Med.*, pp. 702–707 (2017).
- A. Uneri et al., "3D–2D registration for surgical guidance: effect of projection view angles on registration accuracy," *Phys. Med. Biol.* **59**(2), 271–287 (2014).
- J. Qi and R. M. Leahy, "Fast computation of the covariance of MAP reconstructions of PET images," *Proc. SPIE* **3661**, 344–355 (1999).
- J. W. Stayman and J. A. Fessler, "Regularization for uniform spatial resolution properties in penalized-likelihood image reconstruction," *IEEE Trans. Med. Imaging* **19**, 601–615 (2000).
- J. W. Stayman and J. A. Fessler, "Efficient calculation of resolution and covariance for penalized-likelihood reconstruction in fully 3-D SPECT," *IEEE Trans. Med. Imaging* **23**, 1543–1556 (2004).
- S. M. Schmitt and J. A. Fessler, "Fast variance prediction for iteratively reconstructed CT with arbitrary geometries," in *13th Int. Meeting Fully Three-Dimensional Image Reconstr. Radiol. and Nucl. Med.*, Newport, Rhode Island, pp. 228–31 (2015).
- S. M. Schmitt, M. M. Goodsitt, and J. A. Fessler, "Fast variance prediction for iteratively reconstructed CT images with locally quadratic regularization," *IEEE Trans. Med. Imaging* **36**, 17–26 (2017).
- W. Vennart, "ICRU Report 54: medical imaging—the assessment of image quality," *Radiography* **3**(3), 243–244 (1997).

46. C. Zhao et al., “A robotic X-ray cone-beam CT system: trajectory optimization for 3D imaging of the weight-bearing spine,” *Proc. SPIE* **10948**, 109481L (2019).
47. S. Ouadah et al., “Self-calibration of cone-beam CT geometry using 3D-2D image registration,” *Phys. Med. Biol.* **61**, 2613–2632 (2016).
48. H. Zhang et al., “Prospective regularization analysis and design for prior-image-based reconstruction of X-ray CT,” in *Proc. Int. Meeting Fully Three-Dimensional Image Reconst. Radiol. and Nucl. Med.* (2017).
49. S. Capostagno et al., “Task-driven source–detector trajectories in cone-beam computed tomography: II. Application to neuroradiology,” *J. Med. Imaging* **6**(2), accepted (2019).

Biographies of the authors are not available.