

Biomedical Physics & Engineering Express



PAPER

OPEN ACCESS

RECEIVED

27 June 2025

REVISED

29 September 2025

ACCEPTED FOR PUBLICATION

6 October 2025

PUBLISHED

14 October 2025

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



SCIMITAR: optimising chest digital tomosynthesis devices using geometric simulations and genetic algorithms

Alexander D Hill^{1,2,*} , Daliya Aflyatunova^{1,2} , Aquila Mavalankar³ , Stephen Wells³ , D Keith Bowen³, Fraser Holloway¹ , Lauryn Eley^{1,2} , Ishbel Jamieson³ , Matteo Contino³ and Carsten P Welsch^{1,2} 

¹ Department of Physics, University of Liverpool, Liverpool, United Kingdom

² The Cockcroft Institute, Sci-Tech Daresbury, Warrington, United Kingdom

³ Adaptix Ltd, Oxford University Begbroke Science Park, Oxford, United Kingdom

* Author to whom any correspondence should be addressed.

E-mail: a.d.hill@liverpool.ac.uk

Keywords: radiography, medical devices, genetic algorithms, digital tomosynthesis, medical imaging, simulation-based optimisation, chest imaging

Supplementary material for this article is available [online](#)

Abstract

Objective. Digital tomosynthesis (DT) bridges the gap between planar x-rays and computed tomography, offering rapid, low-dose 3D imaging. A mobile chest DT device could transform procedures such as nasogastric tube placement and early cancer detection. Adaptix Ltd. has developed 3D imaging systems using cold-cathode x-ray emitter arrays on flat panel source (FPS) units for veterinary and orthopaedic applications. Designing a chest DT device using multiple FPSs presents new challenges, requiring simulations that can efficiently explore the large design space and rapidly identify optimal configurations. *Approach.* We developed SCIMITAR, a geometry-based simulation framework that models x-ray radiation coverage in chest DT systems. It evaluates design viability and performance using irradiation uniformity metrics and integrates a genetic algorithm to optimise key system parameters. SCIMITAR further facilitates the evaluation of collimator designs, FPS arrangements, engineering constraints, and dynamic adaptation to different patient volumes. *Main results.* Square collimators generally outperformed circular designs due to better alignment with the cuboid target volume. Across FPS configurations, optimisation consistently yielded maximum source-to-image distances, minimal emitter spacing, and x-ray cone angles near 30°. A four-panel cross arrangement achieved highest uniformity. Imposing engineering constraints such as increased emitter spacing led to approximately linear reductions in uniformity. Introducing vertical offsets to central panels yielded modest gains, though still underperformed compared to configurations without central panels. Dynamic cone angle adjustment enabled device adaptation to different patient sizes, with the four-panel cross consistently delivering the best results. *Significance.* SCIMITAR efficiently optimises chest DT designs under various constraints and assumptions. This work identifies promising configurations, highlights design trade-offs, and demonstrates adaptability across patient sizes. As understanding of system requirements evolve, SCIMITAR's adaptability will enable it to remain a valuable tool in guiding the development of clinically effective, low-dose, mobile 3D imaging devices.

1. Introduction

When selecting the appropriate modality of medical imaging based on ionising radiation, clinicians and radiologists follow the 'as low as reasonably achievable' principle (ALARA), which balances apparent

clinical need against the risks associated with radiation exposure (Armao and Smith 2014). The imaging modalities commonly under consideration are planar 2D x-rays and 3D computed tomography (CT) scans. The former are often preferred due to their wider availability, lower radiation dose, and their ability to

be brought to the bedside. Improvements have been made in developing CT protocols that reduce radiation dosage from previous levels which were between 100 and 1,000 times greater than that received during a chest radiograph (Semelka *et al* 2007, Vonder *et al* 2021), however clinicians must consider the compounding effect of repeat measurements (Sodickson *et al* 2009), which particularly impacts paediatric patients, who are found to have an elevated lifetime cancer risk following early in life CT scans (Brenner *et al* 2001, Brenner and Hall 2007).

While such dosage and practicality considerations often lead to the clinical preference for planar x-rays for routine procedures, this approach presents other limitations. Placement of a nasogastric tube is a routine procedure, however misplacement in the lungs can be fatal (Andresen *et al* 2016). Confirmation of correct placement commonly occurs via planar x-ray radiography (Metheny *et al* 2019), however misinterpretation is a common cause of nasogastric tube incidents (Healthcare Safety Investigation Branch 2020). Although 3D imaging from CT scans would mitigate these errors, its routine is impractical due to the high radiation dose, cost, and lack of mobility. A low-dose and mobile 3D imaging device could offer a safe and reliable alternative to planar x-rays, improving patient outcomes without compromising practicality.

Planar x-rays are also preferred to CT in cases where symptoms are not severe or persistent. A comparison of planar chest radiography (CXR) and CT found that CXR correctly identified lung nodules in only 28% of cases later confirmed by CT (Vikgren *et al* 2008). As lung nodules are often an early indicator of lung cancer, such diagnostic limitations likely contribute to the high mortality rate of lung cancer in the UK. Between 2017 and 2019, lung cancer was the leading cause of cancer-related deaths in the UK despite being less common than prostate and breast cancer in males and females, respectively (Cancer Research UK 2025a). Late-stage diagnosis is a leading factor, with 43.1% of diagnoses being Stage IV (Cancer Research UK 2025b). Only 10% of people survive beyond 10 years after diagnosis, despite 79% of cases being preventable (Cancer Research UK 2025a). Given the limitations of CXR in early lung cancer detection, there is a pressing need for alternative imaging techniques that provide higher accuracy without significantly increasing radiation dose.

Digital tomosynthesis (DT) is a promising technology that bridges the gap between planar x-rays and CT, offering rapid, low-dose imaging while retaining 3D capabilities. While CT acquires projections from a full 360-degree rotation around the patient, DT utilises a smaller number of projections in a narrower angular range about the patient. As a result, DT offers a lower radiation exposure and faster total scan time compared to CT. While DT's depth resolution is lower than CT, studies suggest that DT's image quality

is sufficient for many clinical applications. DT is most commonly used for breast imaging, multiple meta-analyses have found that incorporating digital breast tomosynthesis (DBT) alongside 2D mammography improves patient outcomes (e.g. Yun *et al* 2017, Marinovich *et al* 2018, Alabousi *et al* 2021). As a result, DBT has become utilised as standard in some North American and European clinical settings (Sardanelli *et al* 2017, Boroumand *et al* 2018, Alsheik *et al* 2019). DT has also shown promising results for lung imaging, for example Vikgren *et al* (2008) found that DT demonstrated an accuracy of 92% in detecting lung nodules later confirmed by CT, while Choo *et al* (2016) found that DT offers a marked improvement over standard x-rays in the detection of airway lesions.

DT devices typically generate x-rays via a conventional x-ray tube that moves linearly through multiple firing positions to acquire the required projections for 3D reconstruction (e.g. Machida *et al* 2010, Johnsson *et al* 2014, Gange *et al* 2024). However, this approach introduces engineering and software challenges, including machine vibrations, uncertainty in source positioning after movement, and motion-induced blur, all of which can degrade image resolution (Zheng *et al* 2019). Additionally, patients are required to hold their breath for the scan duration, which can further impact image quality and even limit accessibility for certain individuals (Gay *et al* 1994). An alternative approach is to use multiple fixed-position x-ray sources that are sequentially fired in a rapid sequence. This has become viable in recent years with the advent of 'cold-cathode' x-ray sources, which are more compact than traditional x-ray tubes and can be arranged as a linear array or on a flat panel source (FPS) (Bowen *et al* 2023). Studies have shown that a 2D grid arrangement of emitters on an FPS, rather than the conventional linear distribution (Gunnell *et al* 2019, Inscoe *et al* 2024, Billingsley *et al* 2025), can improve image quality (Acciavatti *et al* 2019, Wells *et al* 2020). Additionally, this setup enables a reduced source-to-image distance, lowering power requirements and overall system cost (Phillips *et al* 2019).

Adaptix Ltd.⁴ develops DT devices for 3D imaging in the veterinary, orthopaedic, non-destructive testing, and dentistry sectors. These devices utilise an FPS containing an array of cold-cathode x-ray emitters embedded on a silicon wafer. The emitters are arranged in a 2D grid, providing sufficient coverage of source locations for 3D imaging. Building on this, there is now an effort to develop a low-dose, mobile chest DT device that could enable bedside and off-site imaging, reduce CT load, and support public health screening campaigns (e.g. Ferrari *et al* 2018). The larger volume of the human chest compared to existing applications necessitates research into multiple FPS configurations to ensure full coverage, given that from an engineering perspective a single, large FPS

⁴ <https://adaptix.com/>

would be impractical (Primidis *et al* 2022). Using Monte Carlo simulations, Primidis *et al* (2021) found that such an array can achieve radiation doses comparable to existing clinical DT systems, supporting its feasibility for widespread use.

Tomosynthesis with the proposed device involves sequential acquisition of projection images from a fixed 2D array of x-ray emitters. Unlike conventional systems that rely on a mechanically swept single source or a linear stationary array, this design uses a stationary two-dimensional emitter configuration. Each source is activated in turn, eliminating mechanical delays and enabling faster image acquisition. Traditional systems can require up to 10 seconds to complete a sweep across tens of source positions (Quaia *et al* 2013). By removing the need for physical movement, the proposed system aims to reduce total acquisition time or use it more efficiently. Each emitter operates via cold cathode emission, with electrons accelerated to 90–120 keV striking a metal target to produce x-rays (Primidis *et al* 2021). The entire device, including the detector, is being designed to fit within a 1 m³ volume, which constrains the angular span that the source array can cover. The equivalent gantry angle shall depend on the final optimised design.

Constructing a chest DT device using multiple FPSs presents a complex engineering challenge due to the numerous factors affecting performance and clinical viability. The fundamental challenge is to maximise imaging quality while minimising patient dose and stray radiation. Key considerations include individual x-ray emitter properties, emitter distribution on each FPS wafer, panel position and orientation, and patient size. These aspects cannot be optimised in isolation. For example increasing the angular spread of x-ray emission may improve coverage, but beyond a certain threshold it leads to excessive stray radiation. Likewise increasing the distance between the device and the patient may enhance irradiation uniformity within the target volume, but it may also raise the power requirements prohibitively. A deeper understanding of these trade-offs is essential for effective chest DT design. To address these challenges, various simulations have been developed to address different aspects of the design process (Primidis 2022a). In particular, there is a need for efficient simulation and optimisation frameworks that can rapidly evaluate design solutions based on imaging performance, radiation safety, and engineering constraints.

This work introduces SCIMITAR, a geometry-based simulation designed to model x-ray radiation coverage for a given FPS arrangement. Based on specified design criteria, SCIMITAR first determines whether a design is physically viable and accords with radiation safety protocols, ensuring that the highest percentage of dose is used for imaging and enabling adherence to the ALARA principle. SCIMITAR evaluates the design's performance in terms of irradiation

uniformity throughout a given patient volume. Unlike previous approaches, SCIMITAR integrates genetic algorithms for design optimisation, enabling the rapid and efficient evaluation of chest DT system configurations. Optimal configurations generated by SCIMITAR can serve as a foundation for more detailed Monte Carlo simulations and experimental validation. Additionally, SCIMITAR offers greater flexibility than previous methods, allowing newly identified engineering constraints or desirable design features to be incorporated as input parameter bounds, new optimisation metrics, or acceptance criteria.

This article is structured as follows: section 2 describes SCIMITAR and the numerical methods employed in this work; section 3 presents key findings and insights into the design of a chest DT system; section 4 discusses these findings in a wider context; section 5 summarises the work; and the Electronic Supplementary Materials (ESM) presents extra methodology and results, discussed in the main body of the work.

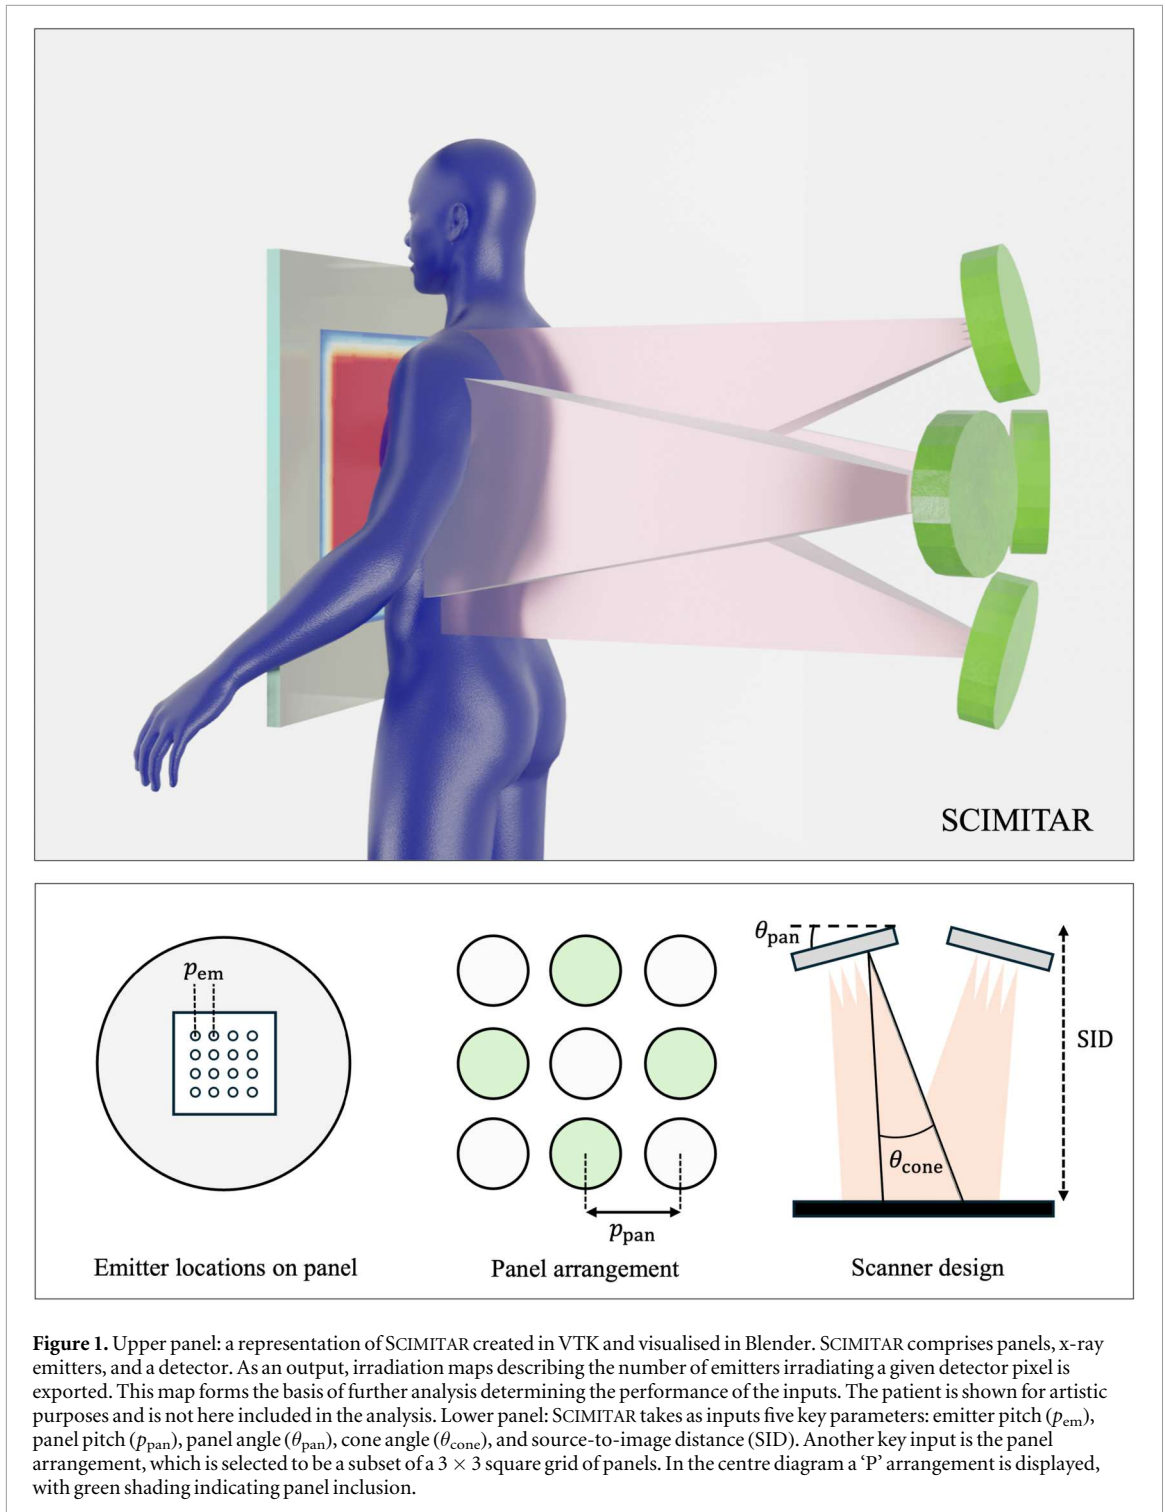
2. Methodology

This section introduces and describes the SCIMITAR simulation, a geometry-based simulation designed to model x-ray radiation coverage for a given flat panel source (FPS) arrangement. Section 2.1 details SCIMITAR's construction, in particular its input parameters and output evaluation metrics, while section 2.2 describes the optimisation procedure. Sections 1 to 4 of the ESM provide a more thorough description of SCIMITAR's design and optimisation. A publicly available version of the SCIMITAR software is available on GitHub (Hill *et al* 2025).

2.1. SCIMITAR simulation

We introduce the 'Simulating Complete Irradiation Maps and Improving Tomosynthesis in x-ray Radiography' (SCIMITAR) project, a geometric simulation that models the construction of a chest digital tomosynthesis (DT) device. This simulation builds on the Iponi simulation (Primidis *et al* 2022) by incorporating greater flexibility in system design and genetic algorithm-based optimisation. Figure 1 displays a visualisation of SCIMITAR in operation, as well as diagrams depicting some of the key input parameters. SCIMITAR is constructed using the Visualisation Toolkit (VTK Schroeder *et al* 2006), an open-source software system utilised for 3D modelling in the context of computer graphics, scientific research, and many other sectors. VTK has been widely used to support medical research, including in the context of prostheses (Cutti *et al* 2024), quantitative imaging (Fedorov *et al* 2012), and practitioner training (Wheeler *et al* 2018).

SCIMITAR creates a 3D environment that incorporates sources of x-ray radiation, a detector and a



patient volume. Irradiation maps are computed as the intersection of the x-rays and the detector, and form the basis for further analysis. X-rays emanate from individual emitters which are themselves embedded on a series of flat panel sources (FPSs), per the basic design of the chest DT device under development. In SCIMITAR a single FPS is modeled as a cylindrical monolith of diameter 21 cm and thickness 5 cm, with a 100 cm^2 wafer embedded on the underside. The emitters are etched onto the wafer, arranged in a 4×4 square grid. X-rays in SCIMITAR are treated as a cone

of uniform intensity propagating from emitters at an angle perpendicular to the face of their host FPS. We note that what we term here 'X-rays' are in fact simple geometric objects, unaffected by absorption, attenuation, or any other physics associated with real x-rays. SCIMITAR is designed to rapidly assess the viability of a large number of possible design configurations of a chest DT device, the learnings of which will then inform more realistic simulations and experimental campaigns. The x-ray cones in this work are modelled with either circular or square cross-sections to mimic

possible choices of collimator shape. The cones are described by an opening angle (θ_{cone}) and a cone resolution (N_{res}), which is the number of vertices describing the cross-section of the cone. For circular cones θ_{cone} determines the diameter of the cross-section, while for square cones θ_{cone} determines the distance between corners of the cross-section.

In SCIMITAR, the detector is modelled as a regular $43 \times 43 \text{ cm}^2$ square grid with a pixel density of n_{pix} . Irradiation maps are generated not only at the detector surface but also at a series of parallel planes positioned above it. Each irradiation map maintains the same resolution as the detector. To compute the irradiation map at a given height h , we define a plane at h parallel to the detector. We export irradiation maps up to a height of 28 cm above the detector. The intersection of this plane with each x-ray cone is then determined, producing a set of 2D intersection regions. Each intersection consists of n_{pix} triangular elements whose vertices are defined by the cone's geometry, including its orientation relative to the plane. These intersection regions are then mapped onto the square grid by determining which pixels fall within the triangular areas, forming the final irradiation map. Further details of SCIMITAR's design, including numerical convergence tests of resolution, justification of the measured irradiation volumes, and checks for physical realism and stray radiation safety are detailed in sections 1 and 2 of the ESM.

2.1.1. Simulation inputs and outputs

Several components of SCIMITAR are tested in this work in the context of designing an optimal chest DT device. These inputs are varied during optimisation, while other remain fixed. These key input parameters are: panel pitch (p_{pan}), measured as the distance between the centres of adjacent panels; panel angle (θ_{pan}), the angle of inclination for each panel with respect to the detector⁵; emitter pitch (p_{em}), measured as the distance between the centres of adjacent emitters on a panel; and the source-to-detector distance (SID): defined as the normal distance between the detector surface and the base of the central panel. We also test various base panel arrangements. Starting from a symmetric 3×3 square grid of panels, we selectively 'turn-off' individual panels to create a given panel arrangement. We consider four panel arrangements in this work: a five-panel diagonal cross (DC), a five-panel 'plus-sign' configuration (PC), and two four-panel arrangements with the central panel removed (D and P). In our nomenclature, the arrangement in figure 1 is 'P'.

The primary output of SCIMITAR are irradiation maps, as visualised in figure 1. These are constructed by determining whether a given x-ray cone is irradiating a given pixel. For the four- and five-panel

arrangements there are 64 and 80 total emitters, respectively. The maximum pixel values for each of these cases are therefore 64 and 80. SCIMITAR by default exports 29 irradiation maps, one at the detector surface ($h = 0 \text{ cm}$) and 28 linearly spaced between $h = 0.5$ and 27.5 cm. Example irradiation maps can be seen throughout this work. We expect radiation coverage throughout the patient volume and a high level of detector illumination to be correlated with DT image quality.

As in Primidis (2022b), the performance of a given simulation run is assessed using output metrics. In this work, we focus on the quality of x-ray illumination across the detector surface and within an approximate patient volume by analysing outputted irradiation maps. This is motivated by the aim to construct a device that provides consistent image quality while minimising radiation dose. At this stage, we assume that uniform image quality is associated with many emitters delivering well-distributed irradiation throughout the target volume. However, we acknowledge that angular coverage and gaps between source positions also directly impact image quality, and that engineering constraints will ultimately shape the final system configuration. The design space for a stationary chest DT system with a 2D array of emitters embedded on separate panels is large, and the precise influence of emitter distribution on image quality is not yet established in this context. We therefore elect not to impose a priori constraints in this regard. The current version of SCIMITAR represents a first step in an iterative design pipeline, where promising geometries are rapidly generated and subsequently evaluated through more computationally intensive simulations. These evaluations will inform the ongoing refinement of optimisation criteria. The explicit incorporation of angular span, emitter distribution, and dose-weighted coverage into the optimisation process, alongside improved image quality metrics, will be a focus of future work.

We define the 2D overlap quality metric as

$$Q_{2D} = \frac{\sum_{p=1}^{N_{\text{pix}}} I_p}{N_{\text{pix}} N_{\text{em}} N_{\text{pan}}} \times 100, \quad (1)$$

where I_p is the number of emitters irradiating a given pixel, N_{pix} is the total number of pixels on the detector surface, N_{em} is the number of emitters per panel, and N_{pan} is the number of panels in the arrangement. The optimal case, $Q_{2D} = 100$, would see every emitter irradiating every pixel. This metric can be computed for the detector surface (Q_{2D}^{det}) or at an arbitrary height X for which an irradiation map has been exported (Q_{2D}^X).

We further introduce the 3D overlap quality metric, which describes the extent to which a volume above the detector surface is irradiated. This is computed by calculating Q_{2D} at a series of heights above the detector surface. These heights start at 0.5 cm and are separated by 1 cm. Formally, the metric is defined

⁵ The rotation axes for each panel are described by Primidis (2022a), their table 6.3.

as

$$Q_{3D} = \frac{\sum_{j=1}^{N_{\text{lay}}} \sum_{p=1}^{N_{\text{pix}}} I_{p,j}}{N_{\text{pix}} N_{\text{em}} N_{\text{pan}} N_{\text{lay}}} \times 100, \quad (2)$$

where N_{lay} is the number of layers above the detector surface. Again, the ideal case would see $Q_{3D} = 100$, with every emitter irradiating every ‘pixel’ at each layer.

The default $43 \times 43 \times 28 \text{ cm}^3$ volume considered in this work in practice would be appropriate for larger adults, however imaging smaller patients by irradiating such a volume would be excessive and potentially harmful. We therefore also consider smaller volumes to represent different patient physiologies and assess how a device optimised for a larger irradiation volume may be adapted by a user to suit a smaller patient. We refer to Kleinman *et al* (2010), who report on paediatric chest dimensions as measured by a CT scanner as a function of age. Based on these results, in this work we consider secondary volumes of $32 \times 32 \times 23 \text{ cm}^3$ and $23 \times 23 \times 18 \text{ cm}^3$ to represent 20 and 10 year olds, respectively. We note that Kleinman *et al* (2010) find that different chest measures do not increase at the same rate with age (their figures 4 and 5), and as such their ratios are not fixed. It is therefore important to test a chest DT device’s ability to adapt to not only scaled volumes, but volumes with different height/width ratios.

2.2. Optimisation of SCIMITAR

In this work we aim to evaluate the performance of a given chest DT device design based on its ability to uniformly irradiate a patient throughout their volume, i.e. by maximising the Q_{3D} metric. However, calculating this for our default $43 \times 43 \times 28 \text{ cm}^3$ volume is computationally time consuming and therefore impractical for the many thousands of simulation runs required to optimise SCIMITAR. A ‘full-run’ (i.e. computing both Q_{2D}^{det} and Q_{3D}) requires calculating 29 irradiation layers at a spatial resolution of 4 cm^{-2} , taking between 7 and 12 seconds⁶, depending on the number and shape of the emitters (square or circular). In contrast, computing only Q_{2D}^{det} and Q_{2D}^{mid} involves just two irradiation layers, reducing the runtime by an order of magnitude (0.7 to 1.1 seconds in our tests). It is therefore appropriate to consider a suitable proxy for Q_{3D} that can return approximately the same findings at a lower computational cost.

We find that Q_{2D}^{mid} , Q_{2D} , computed at a height midway through the default volume, demonstrates strong correlation with Q_{3D} , returning a Spearman’s rank correlation coefficient of $\rho \sim 1$. This is relevant for the purposes of optimising SCIMITAR, as it indicates that comparing and ranking runs with Q_{2D}^{mid} is functionally equivalent to doing so with Q_{3D} . Optimisation therefore proceeds with Q_{2D}^{mid} as the target

metric, while Q_{3D} is computed with full-runs to validate the performance of promising candidates, and Q_{2D}^{det} remains important for sanity checks, as only attenuated radiation reaching the detector contributes to image reconstruction. Further details on the correlation between the output metrics are shown in section 3 of the ESM.

2.2.1. Genetic algorithms

We optimise the design of a chest DT device in SCIMITAR by incorporating a single-objective genetic algorithm (GA, e.g. Goldberg 1989, Katoch *et al* 2021), which acts to maximise Q_{2D}^{mid} by identifying optimal values for a set of input parameters: SID, p_{pan} , θ_{pan} , p_{em} and θ_{cone} . While gradient descent methods would suffice for this single-objective optimisation, we employ a GA anticipating future SCIMITAR iterations involving multi-objective challenges, such as maximizing gantry angle for improved z-resolution while minimising emitter spacing to reduce image artifacts, problems for which GAs are well-established solutions (e.g. Deb *et al* 2002). The allowed ranges for these parameters in optimisation are specified in table 1 of the ESM, with the first column representing the default arrangement. The use of a GA enables the quick identification of promising design modalities within the large input parameter space, as can be seen in the efficiency contrast between section 3.1 and subsequent sections.

The GA generates an initial population (P_1) of 500 candidate solutions using Latin hypercube sampling (LHS) to ensure better parameter space coverage than random sampling. Each chromosome is evaluated based on its Q_{2D}^{mid} value, with candidates failing physical realism checks discarded. The top 50% of valid chromosomes are retained and serve as parents for creating an equal number of children through tournament selection, which selects the best from random groups of three. Offspring inherit three genes from one parent and two from the other, with 50% undergoing single-gene mutation. This process repeats for 30 generations, with the highest-performing final candidate selected as the optimised design. The ALARA principle is maintained by maximisation of irradiation uniformity throughout the patient volume while ensuring that radiation safety standards are maintained for any chest DT device design.

3. Results

In this section we present the results of experiments in chest DT device design with SCIMITAR. In section 3.1 we compare the performance of square and circular collimators broadly. In section 3.2 we present GA optimisation for various design constraints and introduce potential improvements with novel device design. Finally, section 3.3 details a method of device adaptation to different patient sizes.

⁶ Measured on a single core of an Apple M1 chip.

Table 1. Feature importance for inputs of SCIMITAR in setting the value of the target, Q_{2D}^{mid} . A linear regression methodology is employed. 30,000 feature and target pairings are obtained via pseudo-random LHS for square and circular x-ray cones.

Feature	Square Cones	Circular Cones
Cone angle (θ_{cone})	0.530	0.479
Source-to-detector distance (SID)	0.432	0.472
Panel pitch (p_{pan})	0.032	0.030
Panel angle (θ_{pan})	0.004	0.008
Emitter pitch (p_{em})	0.002	0.010

3.1. Square and circular collimation of x-rays

The x-ray beams in SCIMITAR are modelled with either square or circular cross sections, reflecting popular collimator shapes in radiography. To determine which is generally more effective for a chest DT system irradiating cuboid volume, we employ a brute-force approach. In this, we use LHS to generate a population of 30,000 pseudo-random input parameter sets for each collimator-type. We present here the results of our analysis with respect to a ‘P’ arrangement of panels as an illustrative example, and further comment on the other panel arrangements in the text. Each design is then simulated in SCIMITAR, and the resulting distribution of Q_{2D}^{mid} values is shown in figure 2.

Of the 30,000 inputs each generated for the square and circular collimators, only 7,563 and 4,246, respectively, passed the checks for physical realism and acceptable levels of stray radiation. High values of Q_{2D}^{mid} are somewhat rare in this brute force approach, with 125 square collimator configurations returning $Q_{2D}^{mid} > 40$, reducing to only 13 for circular collimators. Moreover, only 10 of the square collimator designs result in $Q_{2D}^{mid} > 50$, while no circular collimator designs do.

We assess the relative importance of the varied inputs relative in determining Q_{2D}^{mid} by employing a linear regression-based methodology, the results of which are displayed in table 1. For both collimator types, source-to-image distance SID and cone angle θ_{cone} emerge as the most significant parameters, while in both cases panel angle θ_{pan} and the emitter pitch p_{em} are the least significant. Similar trends are found when determining feature importance via XGBoost and random forest-based methodologies.

Having found that in general square collimators return higher values of Q_{2D}^{mid} than do circular collimators, we further compare their performance by identifying each regime’s best performer. For these two configurations we re-run SCIMITAR to compute full irradiation maps and evaluate Q_{3D} . Figure 3 displays irradiation maps for these cases at three key heights: the top of the volume of interest; the midpoint where Q_{2D}^{mid} is calculated; and at the detector surface. The input parameters and key metrics of these best performers are indicated on the panels.

For both collimator types, irradiation uniformity improves as one draws closer to the detector surface, with only a central region of $\sim 10 \times 10\text{cm}^2$ irradiated by all 64 emitters at the top of the volume compared with an area of $\sim 30 \times 30\text{cm}^2$ at the bottom. This indicates that chest DT systems comprised of x-ray sources clustered in separate FPSs will yield the highest image quality directly beneath the device centre, and in general improve closer to the detector surface. However, image quality will also be influenced by the angular distribution of sources relative to the voxel of interest and the detector. For both optimal designs for the two collimator types, we find that both exhibit near-maximal SID and near-minimal p_{em} . The values of θ_{cone} in both cases ensure that the detector surface is maximally irradiated by a high number of overlapping beams, while ensuring a tolerable level of stray radiation.

From figures 2 and 3, it is evident that square cones offer superior irradiation uniformity, which is perhaps to be expected given that both the detector is square and our volume of interested is cuboid. We repeat this analysis for the D, DC and PC panel arrangements, each time generating a new set of 30,000 LHS-selected input parameters. We obtain comparable results to the P arrangement, however for the D and DC arrangements that circular collimation is slightly more competitive. While for the P arrangement the difference in Q_{2D}^{mid} between the best-performing square and circular collimator designs is 13.0, for PC, D and DC it is 13.8, 6.0, and 4.6, respectively. The relative importance of each input parameter in determining Q_{2D}^{mid} is largely the same for each panel arrangement, with the exception being that SID is the most significant parameter for a D panel arrangement with circular collimation.

Given the superiority of square collimators for each panel arrangement, we elect to focus exclusively on square collimators in subsequent optimisations of panel and emitter arrangements. We note, however, that in some clinical applications a rectangular, circular or custom irradiation field may be preferred. The flexibility of SCIMITAR allows for these scenarios, which will be the focus of further study if based on experimental, simulation, or medical guidance.

Finally note the inefficiency of this brute force approach, in which over two-thirds of the 30,000 simulations fail our checks for physical realism and acceptable stray radiation. Further, only a small subset produce high Q_{2D}^{mid} values. This justifies the transition to GA-based optimisation in our subsequent analyses of possible chest DT design.

3.2. Genetic algorithm optimisation

In figure 4 we display the results of our GA run as outlined in section 3.2. We compare four main panel arrangement configurations: P (blue), PC (orange), D (green), and DC (red). For each of these configurations, we run the GA

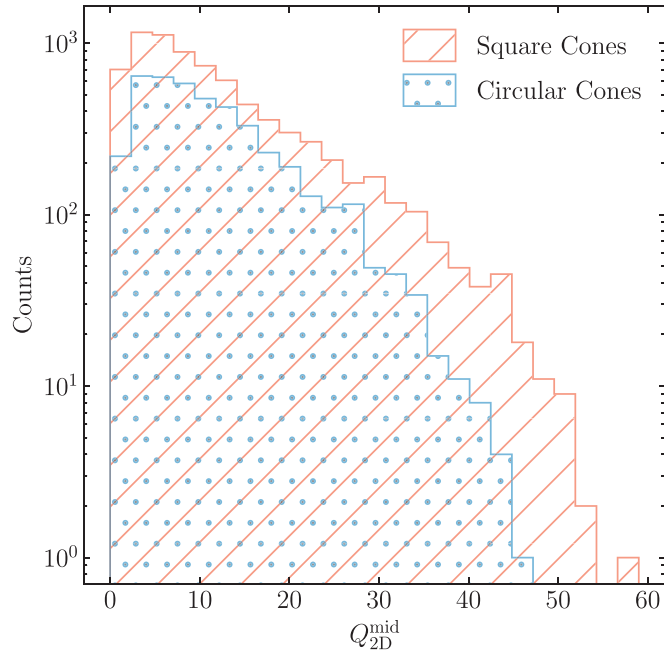


Figure 2. Distribution of Q_{2D}^{mid} values for pseudo-randomly generated simulation inputs. 30,000 sets of inputs each are considered for square and circular collimators. A larger proportion of the square collimator simulation runs pass checks for physical realism and stray radiation. Further, the largest Q_{2D}^{mid} values are associated with square collimators.

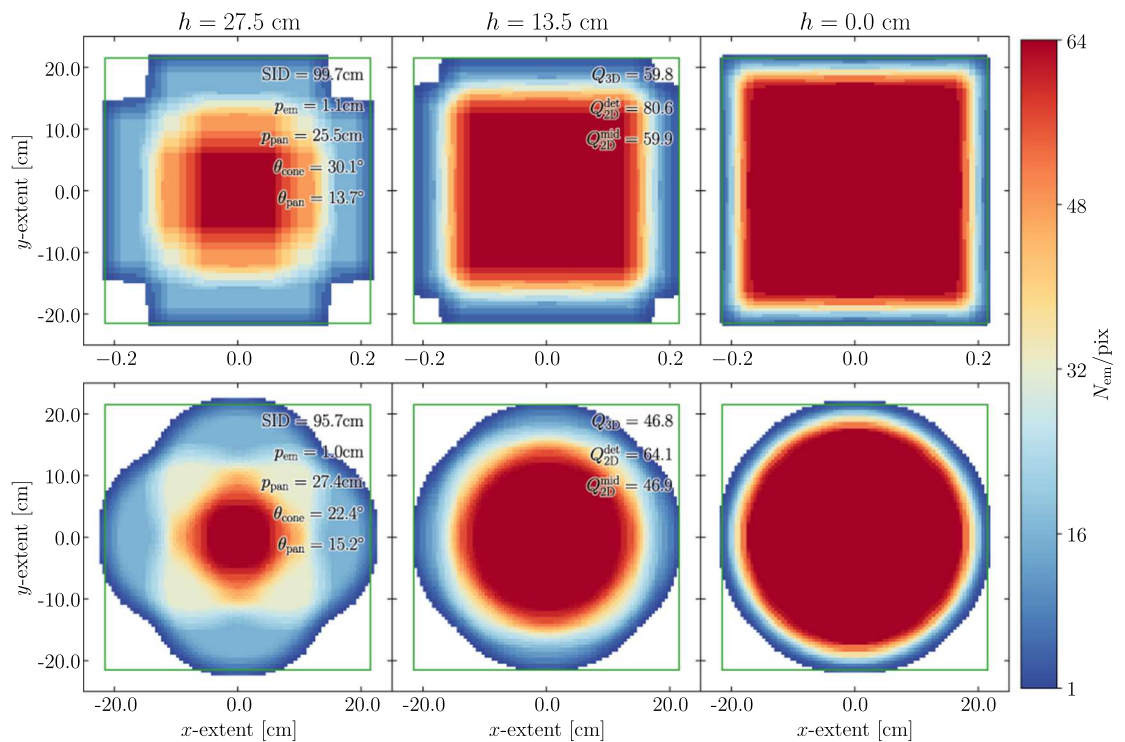


Figure 3. Irradiation maps of the top-performing chest DT device configurations simulated by SCIMITAR from 30,000 pseudo-random LHS-generated input parameter sets. Columns display irradiation at decreasing heights as one moves from the left to the right panels. The top row corresponds to the best performing square-cone arrangement, while the bottom corresponds to the best performing circular-cone arrangement. The corresponding input parameters and key metrics for these arrangements are indicated on the left and centre panels. The green boxes indicate the $43 \times 43 \text{ cm}^2$ region in the x - y plane within which the output metrics are computed. Radiation falling outside these bounds can be considered stray radiation, however the stray radiation criterion outlined in section 2 of the ESM is currently only computed at height $z = 0.0 \text{ cm}$, i.e. the detector surface. The number of emitters irradiating a given pixel area are indicated by the colour, with white areas receiving no radiation. By the metrics employed in this work, which assess irradiation uniformity in a cuboid volume and on square surfaces, square cones are clearly preferred.

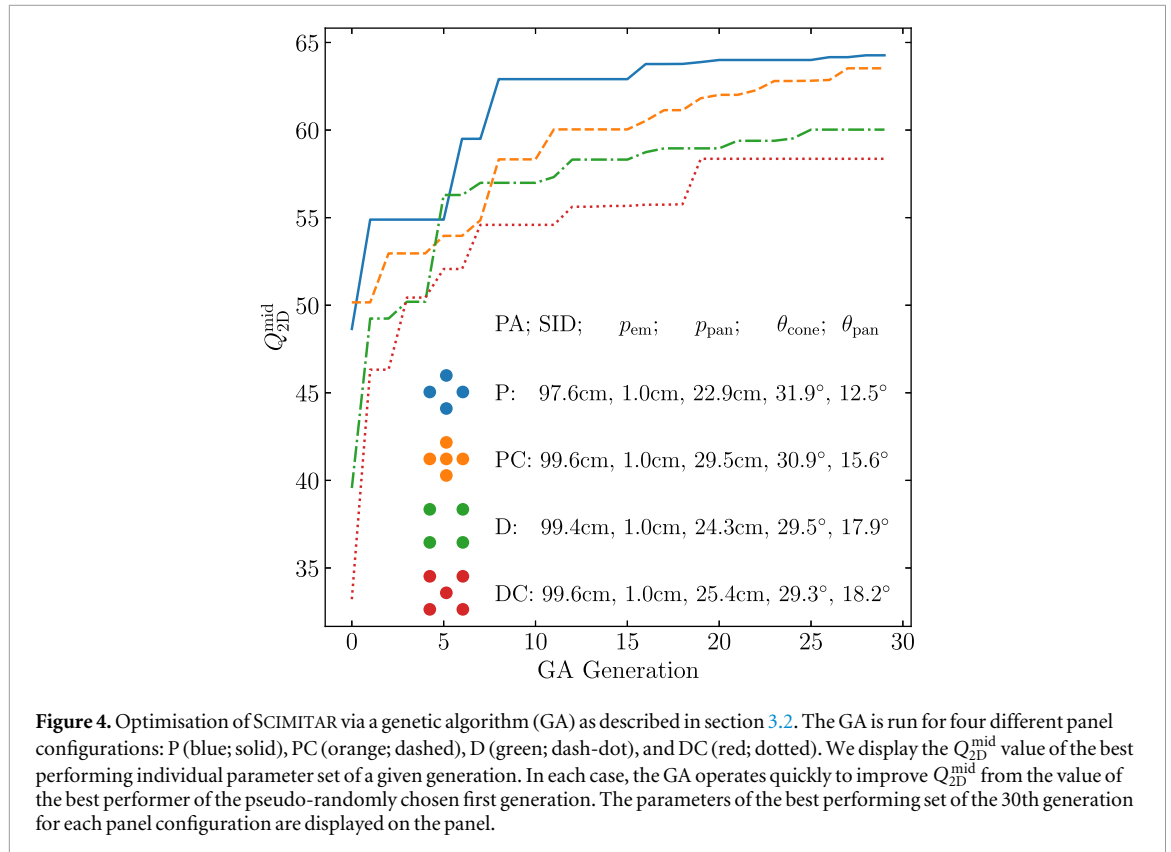


Figure 4. Optimisation of SCIMITAR via a genetic algorithm (GA) as described in section 3.2. The GA is run for four different panel configurations: P (blue; solid), PC (orange; dashed), D (green; dash-dot), and DC (red; dotted). We display the Q_{2D}^{mid} value of the best performing individual parameter set of a given generation. In each case, the GA operates quickly to improve Q_{2D}^{mid} from the value of the best performer of the pseudo-randomly chosen first generation. The parameters of the best performing set of the 30th generation for each panel configuration are displayed on the panel.

with parameter bounds as detailed in section 2.1.1. We find the best performing input parameter set to be a P arrangement with $SID = 97.6$ cm, $p_{em} = 1.0$ cm, $p_{pan} = 22.9$ cm, $\theta_{cone} = 31.0^\circ$, and $\theta_{pan} = 12.5^\circ$, resulting in $Q_{2D}^{mid} = 64.3$. The remaining rankings of the panel configurations are PC, D, and DC with Q_{2D}^{mid} values of 63.5, 60.0, and 58.4, respectively. The parameter inputs of the other configurations are detailed on the figure.

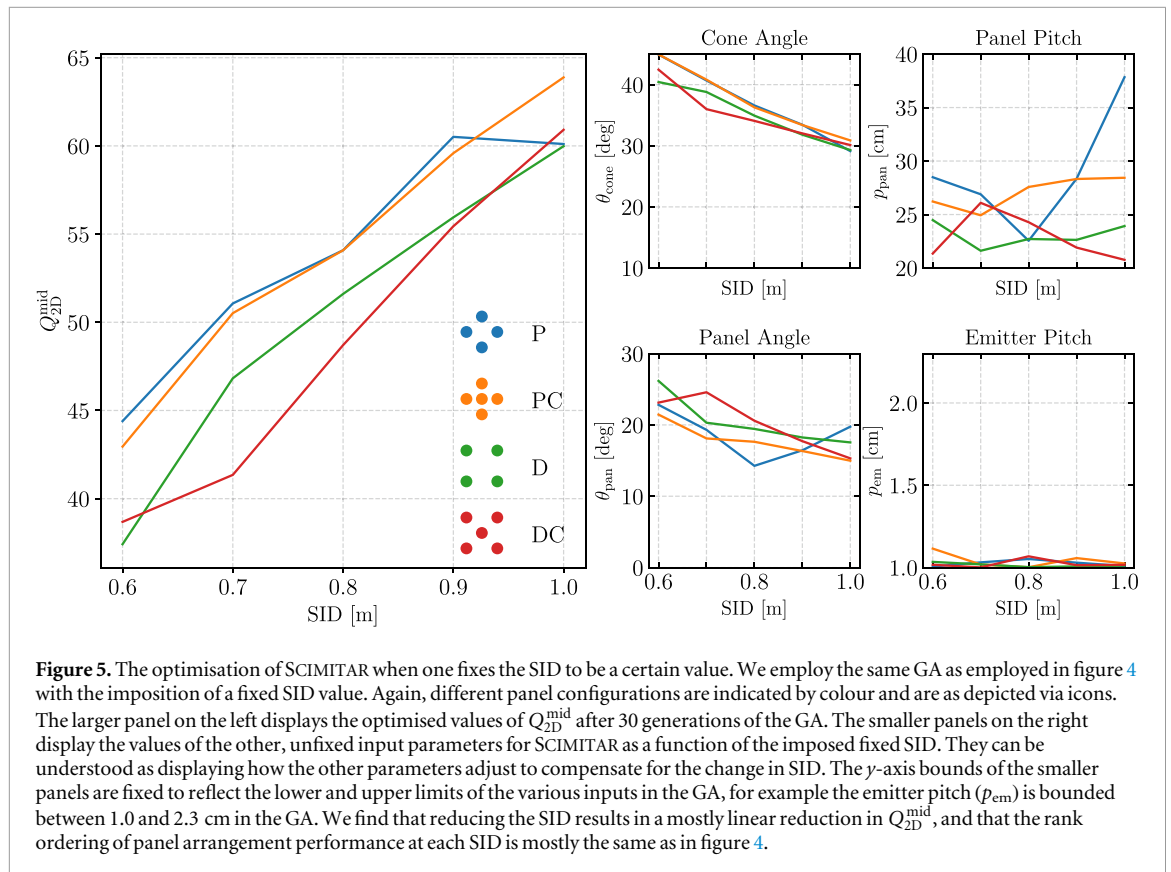
From these results, it appears consistent across all panel arrangements that an optimised chest DT system design as represented in SCIMITAR will exhibit a maximised SID, a minimised p_{em} , and that the value of θ_{cone} will be such that a typical x-ray will maximally fill the detector area given the SID and the panel locations. We note that the irradiation areas created by centre, edge and corner panels will be different for fixed θ_{pan} and θ_{cone} . These will be square, trapezoidal and kite-shaped, respectively. As such, the optimal θ_{cone} is such that the irradiation areas of all the panel types are well contained in the target area while ensuring strong beam overlap in the target volume. Corner panels appear to require smaller θ_{cone} values to strike this balance.

The maximising of the SID can be understood as allowing the system to approach the idealised design of perfectly parallel beams normal to the detector. The minimising of p_{em} can be seen as a means by which Q_{2D}^{mid} can be increased in a manner independent from the response of the other input parameters, as by bringing the origin of x-ray beams of a given panel as close together as possible, the overlap within the

patient volume will naturally increase. This appears to be the case in spite of the apparently modest boost to Q_{2D}^{mid} , given the low significance of p_{em} in determining Q_{2D}^{mid} , as seen in table 1.

Our findings indicate that P-based panel arrangements perform better than those D-based, and that P and D arrangements outperform PC and DC, respectively. This indicates that for a given arrangement removing the central panel results in an improved performance. While some variation due to initial population randomness exists, the final rank ordering of the panel configurations appears as early as the eight generation. Furthermore, we find that in running repeats of this experiment with different LHS-sampled initial conditions that the rank ordering presented here occurs three out of five times, and that the P-configuration each time achieves the highest Q_{2D}^{mid} . Together, these results support the conclusion that the GA effectively samples the input parameter space for each panel arrangement, albeit with a mild sensitivity to initialisation.

The better performance of P-based configurations may be due to their enabling the edges of the x-ray cones to trace the edge of the detector perfectly, as can be seen in the top row of figure 3. This is in contrast to the irradiation maps displayed in section 3.2.2, which are part of a separate analysis but display the performance of a typical DC panel arrangement. Here there are clear gaps in coverage at the detector edge, albeit small and likely outside of the practical imaging



region a radiologist may wish to have imaging capabilities.

The improved performance with no central panels may appear counter intuitive, as one may expect that having a central panel directly normal to the detector would be the best way of ensuring uniform radiation. However, it appears that removing the central panel allows the outer panels to come further in and increase the mutual overlap of their x-ray cones. This can be seen in the reduction of the p_{pan} parameter in figure 4 when comparing DC with D and PC with P.

Despite these findings, it should be noted that the difference between the best and worst performing panel arrangements is only on the order of 9%. This indicates some flexibility in determining optimal panel arrangement when considering other engineering constraints, as will be discussed in the proceeding section, or in imaging quality. Understanding the connection between image quality and beam overlap, among other factors, will be the topic of future work.

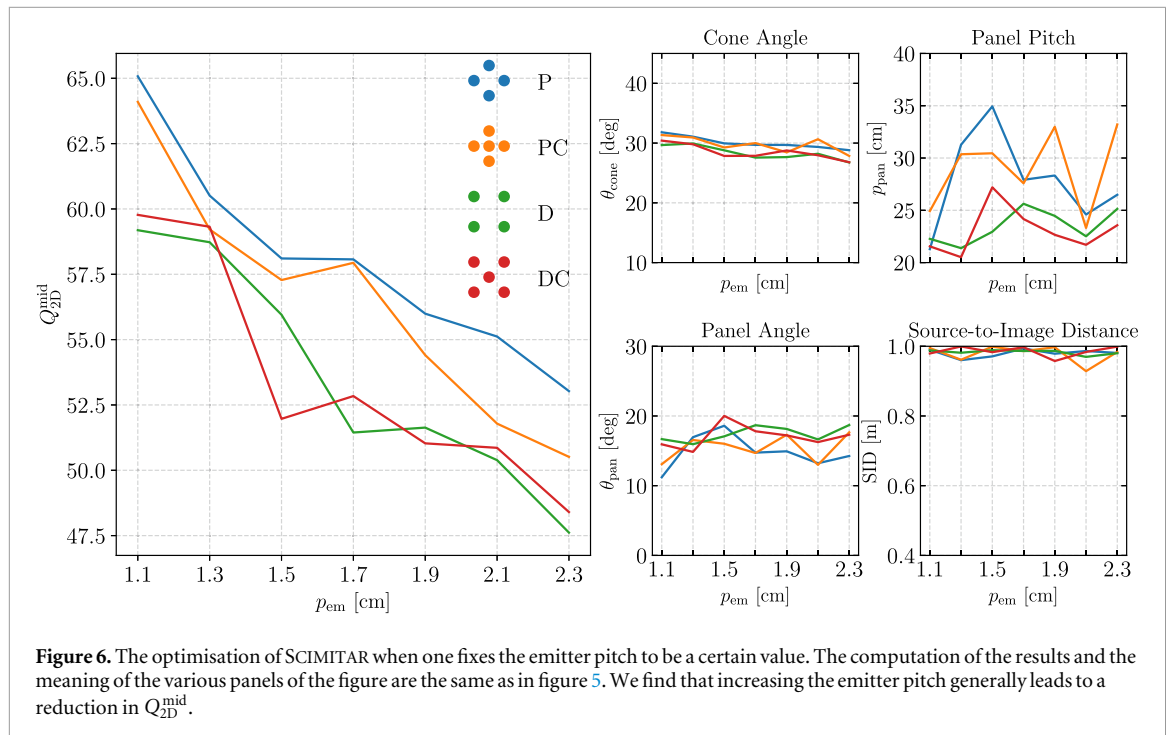
3.2.1. Imposition of engineering constraints

In the preceding section we demonstrated that maximal SID appears a consistent factor in optimising the design of the chest DT system. However, in the development of a practical device it may be desirable to fix or limit the SID below a certain height due to power constraints. This follows as increasing SID reduces the radiation intensity at the detector due to beam divergence. To maintain the same radiation flux

at the detector as SID increases, more power is required to generate a higher initial x-ray output, either by increasing photon quantity or energy. This can lead to increased heat dissipation, increasing operating costs, and increased wear on the x-ray source. It is important then to investigate how reducing the SID would impact the overall irradiation uniformity throughout the patient volume, and therefore the likely imaging quality.

In figure 5 we run GAs for each of the four panel arrangements while fixing the SID in turn to values of 60 cm, 70 cm, 80 cm, 90 cm, and 100 cm. We display how the optimised Q_{2D}^{mid} value changes for each configuration as a function of SID. We also display how the optimised values of the other unfixed parameters change to compensate for the forced reduction of SID. We find that reducing the SID reduces Q_{2D}^{mid} in a mostly linear fashion. The rank ordering of the various panel arrangements found in section 3.2 remains broadly consistent across fixed SID values, however at the lowest SID value (60 cm) the separation in performance between the P- and D-based configurations has increased, i.e. if one were forced to impose $SID = 60$ cm, there would be an even stronger argument for selecting a P-based arrangement.

Considering the response of the unfixed parameters to SID reduction, we find that θ_{cone} reliably increases with reduced SID for all panel configurations. This appears to be a result of simple geometry. For a given isosceles triangle, if one reduces the height while keeping the length of the base the same the



vertex angle must increase. Therefore an increase in θ_{cone} compensates for reduced beam spread at lower SID by maintaining full detector coverage. We do not find evidence of a meaningful, predictable trend in the response of p_{pan} to changing SID, however there is a modest increase in θ_{pan} as SID decreases. Finally, p_{em} appears decoupled from SID in that it is always minimised to ~ 1 cm.

Next, we examine the impact of restricting p_{em} . This is firstly a practical consideration due to potential engineering limitations on emitters spacing. There is also the probable impact of emitter spacing on the quality of eventual image reconstruction, either in that maximising emitter spacing will provide a wider angular distribution of viewpoints and therefore improved image quality, or in that tight clusters of emitters may lead to undesirable image artifacts.

In figure 6, we run GAs for each of the four panel arrangements while fixing the p_{em} to values between 1.1 and 2.3 cm. We find that at all p_{em} values P-based panel configurations exhibit higher Q_{2D}^{mid} values than to D-based, and again P is consistently the best performer. We observe a clear trend where increasing p_{em} leads to a decrease in Q_{2D}^{mid} . For the P configuration, Q_{2D}^{mid} is largest at $p_{\text{em}} = 1.1$ and smallest at $p_{\text{em}} = 2.3$, with values dropping by approximately 18.6%.

We find that the unfixed parameters do not in general display significant changes in their optimal values as p_{em} varies. This is as may be expected, given the low feature importance score of p_{em} in table 1. The parameters p_{pan} and SID do not display significant deviation from their optimal values when no constraints are applied, while the optimal θ_{cone} exhibits a modest decrease as p_{em} increase. This is to be expected, as at fixed θ_{cone} a given panel's irradiation area

will increase with p_{em} , potentially resulting in an undesirable level of stray radiation. Reduction of the cone angle mitigates this. The parameter p_{pan} does not demonstrate a consistent trend across panel configurations, which is perhaps significant. This suggests that an optimal value has not been found in the allowed range with our GA, or that there is a large degree in leniency in what values p_{pan} can take and still return an acceptable Q_{2D}^{mid} .

Finally, we assess the impact of jointly constraining SID (70 cm) and p_{em} (2.3 cm). We run GAs for each panel arrangement, evaluating Q_{2D}^{mid} and the corresponding optimal values of the remaining parameters. We find a further reduction in irradiation uniformity relative to each configuration's performance when only one parameter is fixed. P is again the best performer, with $[Q_{2D}^{\text{mid}}, \theta_{\text{cone}}, p_{\text{pan}}, \theta_{\text{pan}}]$ found to be $[42.0, 37.6^\circ, 24.7 \text{ cm}, 17.7^\circ]$. Next is PC with $[40.3, 36.3^\circ, 31.7 \text{ cm}, 22.3^\circ]$, followed by DC with $[36.9, 35.0^\circ, 21.5 \text{ cm}, 21.0^\circ]$, and lastly D with $[36.0, 33.7^\circ, 24.0 \text{ cm}, 23.2^\circ]$.

Our findings show that constraining system parameters directly affects radiation uniformity throughout the patient volume, but in a largely predictable manner. This will allow chest DT engineers to balance imaging quality, power constraints, and system cost. While the optimal values of other parameters change in response to the constraints on the fixed parameter, they are not able to effectively compensate for the reduction in Q_{2D}^{mid} .

3.2.2. Raising the central panel

We observe that a limitation of configurations including a central panel is that it forces outer panels

to be spread out, reducing the overlap of their x-ray emission. It further has been the case that panel arrangements with no central panel result in improved performance in terms of Q_{2D}^{mid} , however it may be the case that in practical image reconstruction it is desirable to include source positions directly above the patient. We therefore consider a new design that raises central panels above the plane of the outer panels, potentially allowing those outer panels to be brought further in and thus increasing Q_{2D}^{mid} . To achieve this, we introduce two new parameters unique to the central panel in SCIMITAR: vertical offset (z_{stag}), and central panel cone angle (θ_{cone}^{cent}). The former parameter controls the height of the central panel above the plane of the other panel, which remains set by the SID. The latter controls the opening angle of the cones of x-ray emission emanating from the central panel. We add z_{stag} and θ_{cone}^{cent} as parameters to be optimised in our GA, with bounds of [0 cm, 30cm] and [10°, 45°]. We increase the size of each population in the GA to 1000 to allow for the increased size of parameter space to be explored. We reduce the lower limit of the explored p_{pan} range to 5 cm, which enables the outer panels to be placed just outside the edges of x-ray radiation emanating from the raised central panel. SID is fixed to 70 cm, ensuring that $SID + z_{stag} \leq 1$ m. Designs with x-ray radiation incident on the panel bodies themselves are discarded.

We assess the effectiveness of this approach with respect to the panel arrangements that include a central panel: PC and DC. We find that when incorporating a central panel offset, the PC panel arrangement performs marginally better than DC, with Q_{2D}^{mid} values of 54.3 and 52.3, respectively. These can be compared with results presented in figure 5, specifically where SID is likewise fixed to 70 cm but no central panel alterations are included. We find that a central panel offset results in higher Q_{2D}^{mid} not only with respect to comparable PC and DC arrangements, but also to P and D. The improvement in performance for the D-based arrangements is particularly notable, which had previously exhibited values of Q_{2D}^{mid} less than 48. We next assess the potential improvement in incorporating central panel stagger in the case where an additional restriction of $p_{em} = 2.3$ cm is applied. In this case, we again find that imposing a central panel offset results in an improvement in performance with respect to non-offset designs with the same SID and p_{em} constraints. Offset PC and DC arrangements result in Q_{2D}^{mid} values of 44.7 and 42.3, compared with non-offset P and D values of 42.0 and 40.4.

In figure 7, we display the irradiation maps of the optimised DC panel arrangements with SID and p_{em} constraints, with and without a central panel offset. The figure includes top-down and side-on renderings of the panel positions for both configurations. The non-offset configuration results in a relatively poor

performance, with $Q_{2D}^{det} = 58.0$, $Q_{2D}^{mid} = 36.9$, and $Q_{3D} = 37.1$. There are some small locations in the patient volume with no incident radiation, albeit at the far edges. Introducing the central panel offset results in a modest improvement in performance, with $Q_{2D}^{det} = 42.3$, $Q_{2D}^{mid} = 63.6$, and $Q_{3D} = 42.3$. While there are still regions with no irradiation in the volume, the detector is fully irradiated. The central panel offset, z_{stag} , is found to be optimal at 29.5 cm, effectively its maximum allowed value. Combined with SID = 70 cm, this in effect recreates a preferred device height of 1 m that has been found throughout this work. As might be expected given their different heights above the detector, the optimal cone angles for the lower and central panels are different, with the central panel preferring $\theta_{cone}^{cent} = 29.3^\circ$ and the lower $\theta_{cone} = 36.1^\circ$.

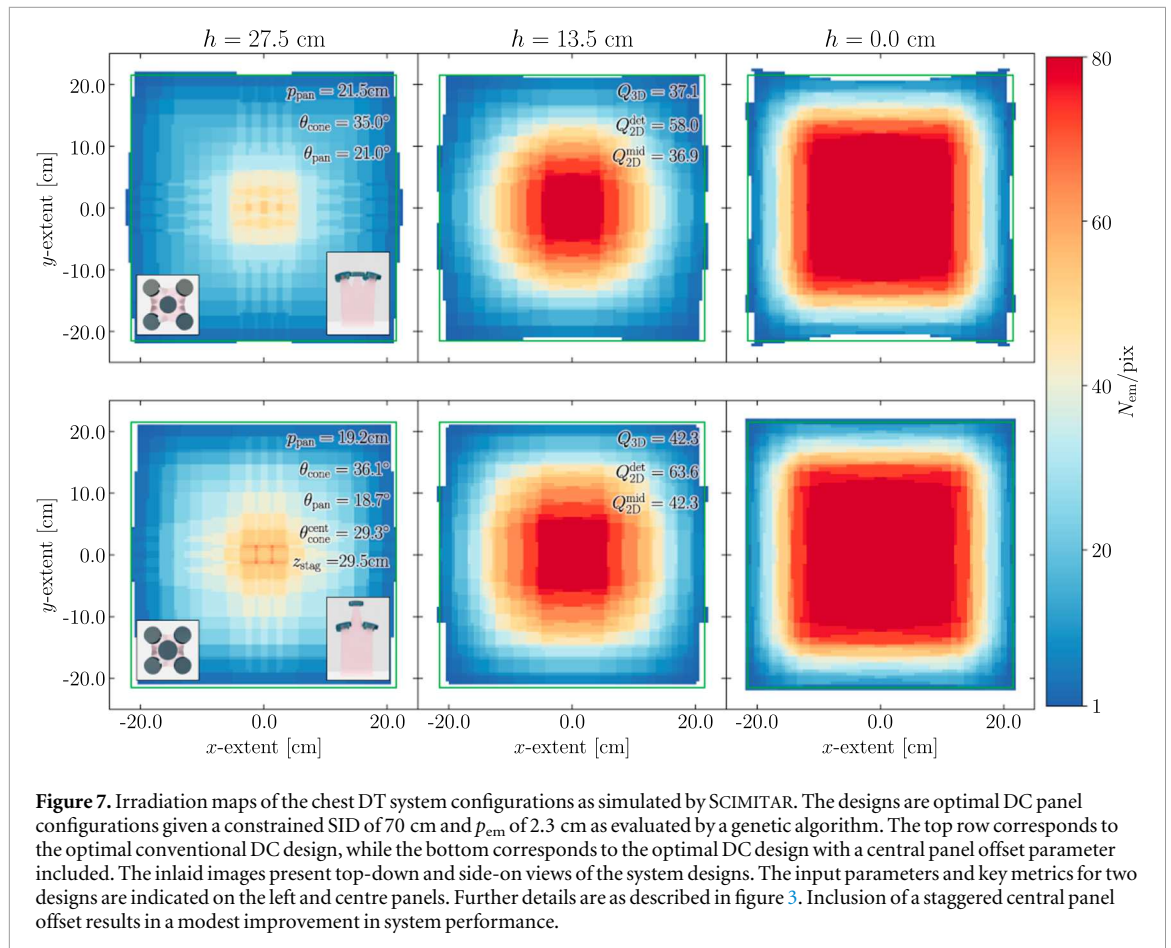
A further finding is that p_{pan} has indeed decreased, allowing the lower panels to be effectively as close together as they can be without being irradiated by beams from the raised central panel. For lower values of z_{stag} , the outer panels would be able to be brought even further in due to the reduced 2D extent of the central panel's radiation. At the extreme of z_{stag} being equal to the panel thickness (5 cm) the lower panels would only have to be further apart than the extent of the central panel's square wafer, upon which the emitters are etched. This, however, does not seem to be the optimal strategy in maximising Q_{2D}^{mid} , as at all times the GA converges on maximising the height of the sources above the detector.

We note that while central panel offset results in an improved performance with respect to arrangements with SID constraints to 70 cm, this is not the case for configurations with no SID restrictions. As seen in figure 4, the optimised designs of each panel arrangement result in $Q_{2D}^{mid} > 56$, greater than the highest value of our staggered designs.

3.3. Device adaptation to different patient sizes

The volume on which our previous analysis has been based is $43 \times 43 \times 28\text{cm}^3$. While this may be appropriate for certain patient sizes, this will be excessive or potentially dangerous for smaller patient sizes. While optimising the design for separate chest DT devices for different patients would be possible, a more economical solution would be to adapt the device in real time to suit different patient needs. We therefore explore whether variation of individual parameters would allow for the adaptation of the device to different desired volumes. We consider two additional target volumes, $32 \times 32 \times 23\text{cm}^3$ and $23 \times 23 \times 18\text{cm}^3$ as described in section 2.1.1.

As examples of the practical consequence of varying SCIMITAR's parameters, changing the SID would involve raising or lowering the device relative to the patient, while varying θ_{pan} and p_{pan} would require the panels to be secured on adjustable mounts. To



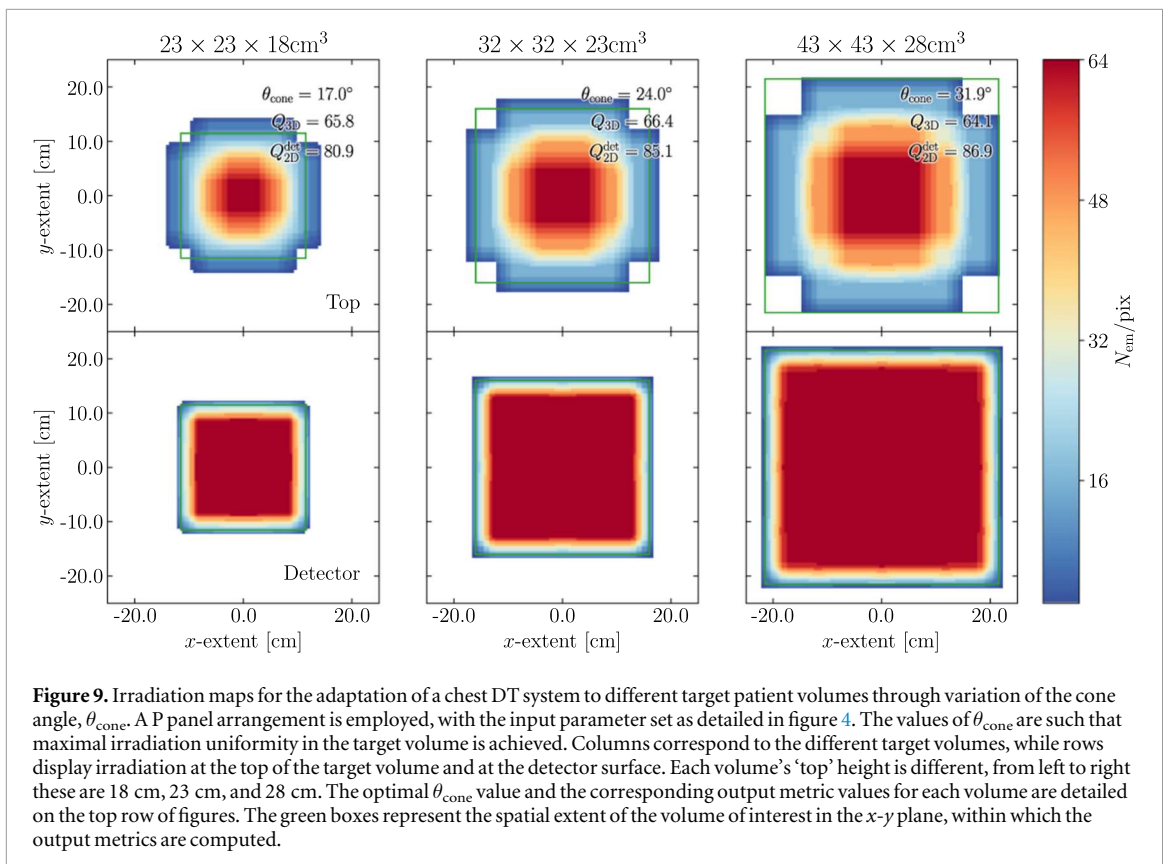
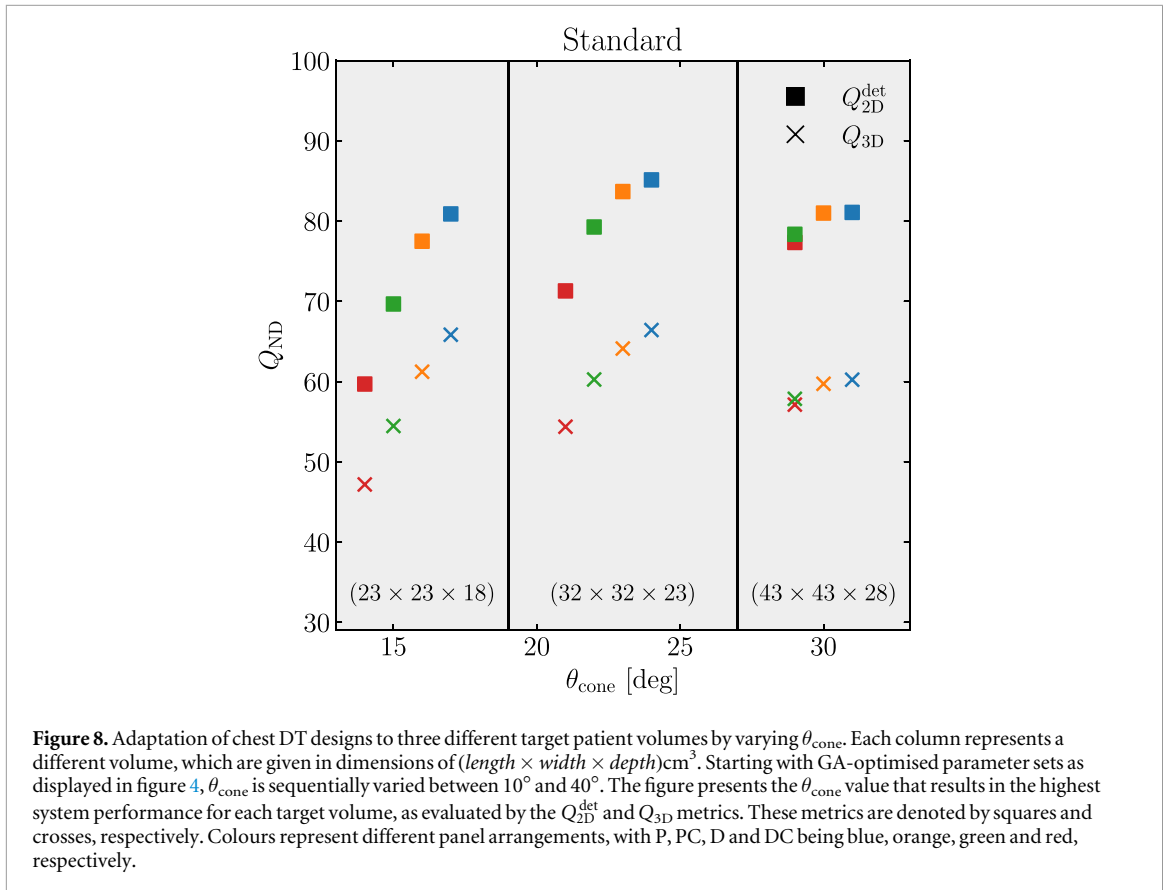
determine which parameter is the most promising in supporting patient size adaptation, we start with the parameter sets of our GA-optimised solutions, such as those in presented in figure 4, and sequentially vary the value of one parameter only while keeping all others fixed. As we have not explicitly computed the strength of correlation between Q_{2D}^{mid} and Q_{3D} for these different volumes, we revert here to measuring performance via the Q_{2D}^{det} and Q_{3D} metrics. For each input parameter and each target volume, we identify the parameter value that maximises Q_{3D} . In our analyses, we find that the most promising parameter to change while fixing all others in order to adapt the chest DT system to different patient volumes is θ_{cone} .

Figure 8 displays the results varying θ_{cone} to adapt to different volumes, starting with parameter sets as detailed in figure 4, named here the ‘Standard’ set. We vary θ_{cone} between 10° and 40° while fixing all remaining parameters, and plot the θ_{cone} value which maximises Q_{2D}^{det} and Q_{3D} for each panel arrangement within a given volume. We find that one can adapt to smaller patient volumes by decreasing θ_{cone} . For the ‘Standard’ case, P is the optimal panel arrangement for each volume. Further, the rank ordering of the panel arrangements remain consistent for the different volumes. The performance of P does not deteriorate for smaller volumes as measured by Q_{2D}^{det} and Q_{3D} , indeed there is evidence that it improves. For P, Q_{2D}^{det} is found to be 80.8, 85.1 and 81.1 for the three

volumes in increasing order of size, while Q_{3D} equals 65.8, 66.4 and 60.2 in turn. As patient volume decreases, there is an increasing spread in the performance of the panel arrangements. While there is only a difference of 3.1 in Q_{3D} between P and DC for the largest volume, this increases to 18.7 for the smallest volume.

In figure 9 we display the irradiation maps one obtains when altering θ_{cone} for the optimal P panel arrangement, i.e. the blue points in figure 8. Changing θ_{cone} to the optimal values for each volume does not result in a significant change in the bulk appearance of the irradiation profile. That is, the majority of the detector surfaces are irradiated by all the emitters, while the maps at the tops of the volumes are cross shaped, with only the central regions fully illuminated. We note that smaller volumes experience potentially undesirable stray radiation outside the volume of interest above the detector. As the stray radiation is acceptable at the detector surface itself the arrangements do not fail the safety requirement outlined in section 2 of the ESM. Further work is needed to consider the safety and device efficiency implications of radiation above the detector falling outside of a $43 \times 43 \text{ cm}^2$ area.

In section 6 of the ESM, we repeat this experiment in varying θ_{cone} for the GA-optimised parameter sets found when constraining SID (70 cm) and p_{em} (2.3 cm). We find that contrary to the ‘Standard’ case, there is a marked deterioration in performance when



adapting to smaller patient sizes by decreasing θ_{cone} . While the rank ordering of the various panel arrangement is broadly the same as it has been throughout

this work, for the P arrangement the value of $Q_{2\text{D}}^{\text{det}}$ decreases by 24.9 between the largest and smallest volumes. There does, however, seem to be a tighter

clustering in the performance of the various panel arrangements across the volumes than is seen in figure 8. Observing the irradiation maps of the P arrangement for this SID and p_{em} -fixed case, it is apparent that the irradiation uniformity at the top of the volumes is poor. Additionally, as the target volume size decreases, the fraction of the detector area that is well-illuminated also decreases. The viability of the adaptation of the chest DT device via variation of θ_{cone} will therefore depend on other design choices and engineering restrictions.

4. Discussion

As evaluated by our performance metrics, our optimisation of a chest DT system based on multiple flat panel sources (FPSs) reveals several consistent and actionable insights. Across all panel arrangements and engineering constraints, maximising the distance between the x-ray sources and the target volume is always desirable. Fixing the source-to-image distance (SID) at smaller values results in a predictable quasi-linear deterioration in irradiation uniformity. Additionally, reducing the separation between emitters as much as possible improves irradiation uniformity, while increasing this separation leads to degradation, albeit to a lesser extent than decreasing SID. The optimal x-ray cone opening angle is primarily determined by the maximum allowed SID and the dimensions of the detector.

Among the panel arrangements investigated, the P configuration (depicted in figure 1) demonstrates superior performance in all scenarios. This advantage is particularly pronounced when considering engineering constraints on the input parameters and when adapting the device for different patient sizes through modifications to the x-ray cone opening angle. Furthermore, while PC and DC designs exhibit modest improvement when incorporating a vertically offset central panel, they fall short of the standard P arrangement.

Modification of the x-ray cone opening angle (θ_{cone}) offers a promising means to tailor the chest DT system to different patient sizes. The predictable change in irradiated volume with adjustments to θ_{cone} suggests that mechanical collimation of individual emitters could allow for patient-specific adaptation via software rather than manual repositioning, a significant design advantage. Fortuitously, θ_{cone} is the only input parameter in SCIMITAR whose modification does not result in changes in the positioning of source emitters. This eliminates any need for recalibration between uses, providing a major practical benefit and motivating further investigation into the FPSs with adaptive collimation.

There are further avenues of study with SCIMITAR that could provide new insights. Considering the relationship between the input parameters when

implementing small changes in the fixed values of emitter separation (p_{em}) and SID, we find that panel angle (θ_{pan}) and θ_{cone} exhibit predictable trends in their optimal values to compensate for changes in the fixed parameters. However, panel separation (p_{pan}) displays less predictable behaviour, occasionally exhibiting large, irregular jumps. This raises the possibility that the genetic algorithm (GA) does not always reliably identify the optimal value for p_{pan} . Modifications to the GA structure or extended run times may alleviate this, however the observed variability also suggests a degree of flexibility in p_{pan} when designing an effective chest DT system. Additional constraints on p_{pan} , such as those aimed at minimising image artifacts caused by widely distributed sources, may therefore be introduced without significantly compromising radiation uniformity.

A further point to note regarding p_{pan} is that panel arrangements without a central panel can accommodate values lower than 20 cm, the lower bound used in this study, without physical collision between panels. Should density and consistency in x-ray source location be critical for image reconstruction, the absence of a central panel, which allows for tighter clustering of sources, may become more advantageous than is already indicated. However, further study is needed to quantify the trade-offs between source positioning, image quality, and practical system design constraints.

A major limitation of SCIMITAR in its current form is its inability to evaluate image quality for a given chest DT design or to estimate patient radiation dose, both crucial considerations in radiography (e.g. Uffmann and Schaefer-Prokop 2009). For instance, while GA-optimisation of SCIMITAR suggests that central panels directly above the patient volume do not yield the best performance in terms of irradiation uniformity, practical image reconstruction may require source positions directly above the patient. Although unrestricted optimisation of different panel arrangements in SCIMITAR produces only slightly poorer irradiation uniformity in configurations with central panels, adapting these designs to different patient sizes appears to result in poorer performance relative to the P arrangement. Moreover, the influence of panel distribution and orientation on radiation flux incident on the patient and detector, particularly in relation to the inverse square law, must be further investigated. While our analysis identifies promising configurations and quantifies the decline in performance metrics when certain parameters are constrained, it remains unclear how these reductions affect image quality and whether they render specific designs unworkable. Understanding these trade-offs will be essential for balancing power consumption, radiation dose, and image resolution.

Several key questions must be addressed to improve SCIMITAR's utility. A critical area of study is the relationship between image quality and overlapping x-ray source properties, such as the required

number of overlapping beams in a given voxel, their angular range, and their spatial clustering. This has been studied extensively in earlier work relating to linear tomosynthesis (e.g. Zhao and Zhao 2008). Ongoing experimental and simulation work aims to clarify these relationships in the context of devices based on a multiple FPS design, with emitters spatially distributed over numerous dimensions. SCIMITAR and its associated genetic algorithm are highly flexible, allowing for the incorporation of new performance metrics as research progresses. Future optimisations could involve multi-objective criteria, which for example may include rewarding configurations that ensure a minimum number of overlapping emitters while avoiding the continued rewarding of a increased number of overlapping beams if the improvement in image quality is minimal. The primary function of SCIMITAR is to rapidly identify optimal designs within a large parameter space based on an evolving understanding of chest DT system requirements. Its adaptability ensures that it can readily integrate new insights as they emerge, ultimately contributing to the development of more effective imaging systems.

Section 7 of the ESM presents a preliminary demonstration of SCIMITAR role within in a wider chest DT design pipeline. We use the CTACardio phantom from 3DSLICER (Fedorov *et al* 2012) in combination with the Insight Toolkit (ITK) (Yoo *et al* 2002, McCormick *et al* 2014) to generate digitally reconstructed radiographs (DRRs) from emitter positions determined via SCIMITAR optimisation. The resulting DT reconstruction reasonably well resolves anatomical structures in the central region of the phantom, however the upper and lower sections are poorly reconstructed, with projections from the five emitter clusters appearing either overly magnified or spatially disjointed across slices. These findings highlight that while initial optimisation based on beam uniformity is a reasonable starting point for geometry selection, additional evaluation criteria are required to ensure adequate sampling and image quality across the full volume. Refining these criteria will be a key focus in future development of the multi-simulation pipeline.

5. Summary

This work introduces SCIMITAR, a geometry-based simulation designed to model the radiation fields produced by chest digital tomosynthesis (DT) devices. This is motivated by developments from Adaptix Ltd., who have created innovative radiographic systems based on flat panel sources (FPS), units comprising many individual x-ray emitters based on cold cathode technology. These FPS systems have enabled low-dose, mobile 3D imaging in the veterinary and orthopaedic sectors. Building on this, efforts are now

underway to develop a multiple-FPS device for low-dose, mobile 3D imaging of the chest.

SCIMITAR allows the rapid evaluation of possible design configurations by simulating radiation distributions and constructing irradiation maps throughout a specified target volume (figure 1). From these maps key performance metrics are computed, which design optimisation via a single-objective genetic algorithm. This work's target metric is Q_{2D}^{mid} , which acts as a proxy for the quantity and uniformity of the radiation passing through the volume (figure S2 of the ESM), and is therefore presumed to be indicative of image quality. In this study, each panel hosts 16 emitters arranged in a square grid, and four main panel arrangements are trialed: P, PC, D, and DC PC and DC correspond to five-panel systems arranged in horizontal-vertical and diagonal configurations, respectively, while P and D represent equivalent four-panel systems with the central panel omitted.

A summary of our results is as follows:

(i) Devices with square x-ray collimation generally outperform those with circular collimation (Figure 2), likely due to better alignment with the square detector area and cuboid target volume (figure 3). For both collimator types, the uniformity of overlapping emitter coverage increases closer to the detector. Across all panel arrangements, the most significant input parameters affecting Q_{2D}^{mid} are the x-ray cone angle (θ_{cone}) and the source-to-image distance (SID), while the emitter pitch (p_{em}) is the least influential (table 1). Random sampling input parameters proves inefficient for identifying promising design configurations, highlighting the need for systematic optimisation methods such as genetic algorithms.

(ii) Genetic algorithm optimisation produced consistent trends across the four panel arrangements (figure 4). The optimal SID is consistently near the maximum allowed ($\sim 1m$), while conversely the optimal p_{em} is at the minimum allowed (1.0 cm). The optimal θ_{cone} is consistently close to 30° , while more variability is seen for the panel angle (θ_{pan}) and panel pitch (p_{pan}). Horizontal-vertical cross arrangements yield higher optimised Q_{2D}^{mid} values than diagonal arrangements, and the four-panel arrangements outperform their five-panel counterparts.

(iii) Clinical and engineering considerations, such as power requirements, may necessitate constraints on certain design aspects. Imposing such constraints prior optimisation systematically reduces the final Q_{2D}^{mid} . Systematically lowering the value of the SID leads to a predictable reduction in Q_{2D}^{mid} , with corresponding increases in the optimal θ_{cone} and θ_{pan} values (figure 5). Similarly, systematically increasing the value of p_{em} reduces the optimised Q_{2D}^{mid} (figure 6). Only modest changes in the other parameters are found in response, with a slight decrease in the optimal θ_{cone} . Joint constraints on both SID and p_{em} are found to have a compounded deleterious effect on

Q_{2D}^{mid} . Across all scenarios, the relationship between constraints and Q_{2D}^{mid} remains approximately linear and predictable.

(iv) Allowing the vertical offset of the central panel in PC and DC arrangements improves the optimised Q_{2D}^{mid} , as this enables closer spacing in the lower panels, thus enhancing overall radiation overlap (figure 7). However the improvement is somewhat modest, and the performance remains inferior to that of the equivalent P and D arrangements, which do not include a central panel.

(v) Dynamic adjustment of θ_{cone} enables the chest DT device to adapt to different patient target volumes, with smaller cone angles better suited to smaller volumes (figures 8 and 9). Across all target volumes, the P panel arrangement consistently outperforms other configurations, particularly at smaller volumes where the other arrangements perform significantly worse.

This work has used SCIMITAR to identify promising designs of a low-dose, mobile chest DT device, providing a foundation for ongoing experimental and simulation work. We find that possible engineering constraints on key aspects of the design, namely the SID and emitter spacing, influence radiation uniformity throughout the patient volume in a predictable manner. This predictability offers engineers a framework within which to balance competing considerations, such as image quality, power consumption, and overall cost. An important result is that dynamic adjustment of the x-ray cone opening angle facilitates adaptation of the device to different patient sizes, suggesting that a single device could be used for both paediatric and adult patients. Further experimental and simulation-based work is needed to validate our findings in terms of image quality and patient radiation dose.

The main utility of SCIMITAR is its ability to rapidly identify optimal designs within a large parameter space given a series of imposed constraints and assumptions. This adaptability enables it to evolve alongside an improving understanding of system requirements. By enabling rapid and adaptable optimisation, SCIMITAR is well-positioned to play a central role in the future development of clinically effective and practically viable chest DT systems.

Acknowledgments

This project has benefitted from funding from the STFC-funded Liverpool Centre for Doctoral Training for Innovation in Data Intensive Science (LIV.INNO) under training grant ST/W006766/1. This work has also received funding from UKRI's late-stage commercialisation scheme via the OptiX project, grant number ST/Z510324/1. The authors thank Thomas Primidis for his foundational contribution to this work and our subsequent conversations. We are

grateful to the referees for their thorough and constructive feedback, which has improved the quality of this manuscript.

Data availability statement

All data that support the findings of this study are included within the article (and any supplementary files).

References

- Acciavatti R J, Barufaldi B, Vent T L, Wileyto E P and Maidment A D 2019 Personalization of x-ray tube motion in digital breast tomosynthesis using virtual Defrise phantoms *Medical Imaging 2019: Physics of Medical Imaging* **10948** 65–73
- Alabousi M et al 2021 *JNCI: Journal of the National Cancer Institute* **113** 680
- Alsheik N H et al 2019 *Academic Radiology* **26** 597
- Andresen E N, Frydland M and Usinger L 2016 *Journal of Medical Case Reports* **10** 1
- Armao D and Smith J K 2014 *North Carolina Medical Journal* **75** 126
- Billingsley A, Inscoc C, Lu J, Zhou O and Lee Y Z 2025 *Med. Phys.* **52** 542
- Boroumand G, Teberian I, Parker L, Rao V M and Levin D C 2018 *American Journal of Roentgenology* **210** 1092
- Bowen D K et al 2023 2023 IEEE 36th International Vacuum Nanoelectronics Conference (IVNC) **192194**
- Brenner D J and Hall E J 2007 *New Engl. J. Med.* **357** 2277
- Brenner D J, Elliston C, Hall E J and Berdon W E 2001 *AJR. American Journal of Roentgenology* **176** 289
- (Cancer Research UK) 2025b *Early Diagnosis Hub* <https://crukancerintelligence.shinyapps.io/EarlyDiagnosis/>
- Cancer Research UK 2025a *Lung Cancer Statistics* (<https://cancerresearchuk.org/health-professional/cancer-statistics/statistics-by-cancer-type/lung-cancer>)
- Choo J Y, Lee K Y, Yu A, Kim J-H, Lee S H, Choi J W, Kang E-Y and Oh Y W 2016 *European Radiology* **26** 3147
- Cutti A, Santi M, Hansen A and Fatone S 2024 *Sensors* **24** 2350
- Deb K, Pratap A, Agarwal S and Meyarivan T 2002 *IEEE Trans. Evol. Comput.* **6** 182
- Fedorov A et al 2012 *Magn. Reson. Imaging* **30** 1323
- Ferrari A, Bertolaccini L, Solli P, Di Salvia P O and Scaraduzzi D 2018 *Annals of Translational Medicine* **6** 91
- Gange C, Ku J, Gosangi B, Liu J and Maolinbay M 2024 *Journal of Clinical Imaging Science* **14** 22
- Gay S B, Siström C L, Holder C A and Suratt P M 1994 *Investigative Radiology* **29** 848
- Gunnell E T et al 2019 *European Radiology* **29** 1665
- Healthcare Safety Investigation Branch 2020 *Placement of Nasogastric Tubes* (<https://hssib.org.uk/patient-safety-investigations/placement-of-nasogastric-tubes/investigation-report/>)
- Hill A, Holloway F and Aflyatunova D (2025) *SCIMITAR: Publicly Available Version* (https://github.com/Alex-Hill94/SCIMITAR_PUBLIC)
- Inscoc C et al 2024 X-ray source array optimization for mobile chest tomosynthesis *Medical Imaging 2024: Physics of Medical Imaging* **12925** 146–53
- Johnsson A A, Vikgren J and Bath M 2014 Chest tomosynthesis: technical and clinical perspectives *Seminars in Respiratory and Critical Care Medicine* **35** 017–26
- Katoch S, Chauhan S S and Kumar V 2021 *Multimedia Tools Appl.* **80** 8091
- Kleinman P L, Strauss K J, Zurakowski D, Buckley K S and Taylor G A 2010 *American Journal of Roentgenology* **194** 1611
- Machida H, Yuhara T, Mori T, Ueno E, Moribe Y and Sabol J M 2010 *Radiographics* **30** 549

- Marinovich M L, Hunter K E, Macaskill P and Houssami N 2018 *JNCI: Journal of the National Cancer Institute* **110** 942
- McCormick M, Liu X, Jomier J, Marion C and Ibanez L 2014 *Frontiers in Neuroinformatics* **8** 13
- Metheny N A, Krieger M M, Healey F and Meert K L 2019 *Heart & Lung* **48** 226
- Phillips S et al 2019 Digital Tomosynthesis: The Next Generation *European Congress of Radiology ECR 2019 C-2547*
- Primidis T 2022a *Design and Optimisation of Ultra-compact, High-resolution, 3D X-ray Imaging Systems* (The University of Liverpool) (<https://doi.org/10.17638/03161341>)
- Primidis T 2022b *Design and optimisation of ultra-compact, high-resolution, 3D X-ray imaging systems* University of Liverpool 10.17638/03161341
- Primidis T G, Wells S G, Soloviev V Y and Welsch C P 2021 *Biomed. Phys. Eng. Express* **8** 015006
- Primidis T, Soloviev V, Wells S and Welsch C 2022 *Proc. 13th International Particle Accelerator Conference (IPAC'22). No. 13 in International Particle Accelerator Conference (JACoW Publishing)* (<https://doi.org/10.18429/JACoW-IPAC2022-THPOMS033>)
- Quaia E, Baratella E, Poillucci G, Kus S, Cioffi V and Cova M A 2013 *Academic Radiology* **20** 546
- Sardanelli F, Fallenberg E M, Clauser P, Trimboli R M, Camps-Herrero J, Helbich T H and Forrai G 2017 *European Society of Breast Imaging (EUSOBI) w. l. r. b. E. D.-T. E. B. C. C 8 11* Insights into imaging
- Schroeder W, Martin K and Lorensen B 2006 *The Visualization Toolkit* 4th edn (Kitware)
- Semelka R C, Armao D M, Elias J and Huda W 2007 *J. Magn. Reson. Imaging* **25**
- Sodickson A, Baeyens P, Andriole K, Prevedello L, Nawfel R, Hanson R and Khorasani R 2009 *Radiology* **251** 175
- Uffmann M and Schaefer-Prokop C 2009 *Eur. J. Radiol.* **72** 202
- Vikgren J, Zachrisson S, Svallkvist A, Johnsson A A, Boijesen M, Flinck A, Kheddache S and Bath M 2008 *Radiology* **249** 1034
- Vonder M, Dorrius M D and Vliegthart R 2021 *Translational Lung Cancer Research* **10** 1154
- Wells S G, Elangovan P, Dance D R, Wells K, Soloviev V Y, Renforth K L and Young K C 2020 *Medical Imaging 2020: Physics of Medical Imaging* (<https://doi.org/10.1117/12.2548438>)
- Wheeler G, Deng S, Toussaint N, Pushparajah K, Schnabel J, Simpson J and Gomez A 2018 *Healthcare Technology Letters* **5**
- Yoo T S, Ackerman M J, Lorensen W E, Schroeder W, Chalana V, Aylward S, Metaxas D and Whitaker R 2002 *Medicine Meets Virtual Reality 02/10* (IOS Press) 586–92
- Yun S J, Ryu C-W, Rhee S J, Ryu J K and Oh J Y 2017 *Breast Cancer Research and Treatment* **164** 557
- Zhao B and Zhao W 2008 *Med. Phys.* **35** 5219
- Zheng J, Fessler J A and Chan H-P 2019 *Med. Phys.* **46** 5572
- Goldberg D 1989 Genetic algorithms in search, optimization, and machine learning *Addion wesley* (**102**) 36