scientific reports

Check for updates

OPEN

Human-organ-scale x-ray fluorescence ghost imaging for radioisotope-free diagnostics

Eli Levinson[1]✉, Rachel H. Shukrun[1], Nicola Viganò[2] & Sharon Shwartz[1]✉

A wide range of diagnostic information in medicine is currently obtained using radioactive tracers. While central to nuclear medicine, these methods face inherent limitations: radiation dose restricts repeat examinations, short tracer half-lives and complex logistics limit accessibility and increase costs, and relatively low spatial resolution often necessitates complementary CT or MRI. Here we present a first proof-of-concept demonstration of an alternative imaging modality based on x-ray fluorescence (XRF) computational ghost imaging (CGI) at the human-organ scale. Using a thyroid phantom filled with iodine solution as a model system, we show that structured illuminations combined with fluorescence detection reconstruct the iodine distribution with high fidelity. This approach avoids the use of radioactive tracers while preserving key aspects of the image structure and contrast. We report a spatial resolution of ~ 4.5 mm in the current implementation, with potential for further improvement through system optimization. Beyond this demonstration, XRF-CGI establishes a generalizable framework for tracer imaging without radionuclides, opening a route toward safer, repeatable, and more accessible diagnostics.

In nuclear medicine, methods such as scintigraphy, positron emission tomography (PET), and single-photon emission computed tomography (SPECT) enable visualization of metabolic processes and tissue-specific uptake using radioactive tracers. In scintigraphy, a detector records gamma emissions from radionuclides to produce two-dimensional functional images. SPECT extends this concept by rotating the detector to reconstruct 3D tracer distributions. PET detects coincident annihilation photons from positron-emitting isotopes, mainly ^{18}F[1], allowing highly sensitive and quantitatively accurate tomographic imaging of metabolic activity. Together, these methods have become indispensable in modern diagnostics, particularly in oncology and cardiology, where they provide quantitative insights that guide both diagnosis and therapy[2–6].

Despite their wide utility, all nuclear imaging methods fundamentally rely on the radioactive decay of unstable isotopes, whether through direct gamma emission or positron annihilation, imposing unavoidable physical constraints. Ionizing radiation increases the risk of cancer and other diseases, limiting the permissible dose and consequently the frequency of repeated examinations[7,8]. Typically, short-lived radiotracers with half-lives ranging from a few hours to several days are employed, creating major logistical challenges in isotope production, transport, and storage[9,10]. Moreover, the relatively low spatial resolution of SPECT (~ 8–10 mm) and PET (~ 4–5 mm under clinical conditions) often necessitates the use of hybrid imaging with CT or MRI, increasing complexity and cost[10,11].

Recent advances in SPECT and PET have improved both hardware and reconstruction capabilities[12–15]. These include the development of solid-state detector technologies[12], time-of-flight (TOF) PET[13], optimized system geometries and iterative as well as machine-learning-based reconstruction methods[14], which together enhance sensitivity, quantitative accuracy, and spatial resolution. In addition, hybrid imaging platforms such as SPECT-CT[12] and PET-CT[15] provide improved anatomical localization, attenuation and scatter correction, further strengthening their clinical utility. Nevertheless, these advancements cannot eliminate the intrinsic drawbacks of radioactive imaging itself, including dependence on a limited selection of short-lived radionuclides and the associated logistical constraints.

These remaining limitations motivate the search for alternative approaches that avoid radioactive tracers, while maintaining clinically relevant spatial resolution and clinically acceptable radiation doses. One intriguing direction lies in x-ray fluorescence (XRF)-based modalities. In XRF, incident x-rays excite inner-shell electrons in the sample, and the subsequent emission of element-specific fluorescent photons provides a direct fingerprint

[1]Physics Department, Institute of Nanotechnology and Advanced Materials, Bar-Ilan University, Ramat Gan 52900, Israel. [2]Université Grenoble Alpes, CEA, IRIG-MEM, Grenoble 38000, France. ✉email: eliyahu.levinson@mail.huji.ac.il; Sharon.shwartz@biu.ac.il

of elemental composition. XRF therefore enables elemental concentration mapping, in close analogy to how SPECT maps radionuclide distributions, and can in principle deliver comparable functional information without reliance on radioactive decays[16].

A key advantage of XRF is that, unlike radioactive tracers, it does not rely on ongoing decay: ionizing radiation is present only during x-ray excitation and vanishes immediately once the illumination is turned off. This property has established XRF as a powerful analytical tool and imaging modality across a broad range of disciplines, including biomedical research and preclinical imaging[17], archaeological studies of pigments and artifacts[18], and environmental monitoring of trace elements in soils and sediment[19]. Early efforts to translate XRF into medical imaging date back several decades, including fluorescent thyroid scanning in humans, which demonstrated that it could provide sufficient elemental contrast at low dose without radioactive tracers[20]. However, these studies also highlighted a major limitation of conventional XRF: its reliance on raster scanning using a narrow pencil beam. Although pencil beams can produce high-resolution elemental maps, acquisition times scale prohibitively with field of view, making human-organ imaging impractical[21]. Consequently, XRF imaging has remained largely confined to small-animal studies and small-scale applications[22–24], while extending it to human-organ scales poses formidable challenges[16]. Importantly, this limitation is fundamentally architectural rather than algorithmic, arising from the inherently serial nature of point-by-point XRF acquisition. This motivates the development of approaches that convert XRF from a serial raster-scanning modality into a parallel and information-efficient imaging framework.

To overcome this challenge, here we introduce a computational approach that replaces sequential pencil-beam scanning with parallel structured illumination and algorithmic reconstruction. The transition from conventional raster-scanned XRF to the CGI-based approach employed in this work is illustrated schematically in Fig. 1a. In this x-ray fluorescence computational ghost imaging (XRF-CGI) scheme, the element map is recovered by correlating the known illumination patterns with the total fluorescence measured for each pattern[25].

In our proof-of-concept experiment, we image a human thyroid phantom and reconstruct the iodine distribution with a spatial resolution of ~4 mm. Although demonstrated here for structural mapping, the same principle could naturally extend to functional imaging using element-targeted contrast agents. Together, these results establish a pathway toward XRF-based medical imaging that combines high spatial precision, rapid acquisition, and the safety of non-radioactive tracers.

Results

Our approach is based on computational ghost imaging (CGI)[26], in which spatial information is recovered from correlations between a set of known structured illumination patterns and a single integrated (bucket) signal recorded for each pattern. In a discrete representation, the measurement process can be described as[27]:

$$y = Wx + \epsilon \tag{1}$$

where y is the vector of measured bucket fluorescence signals, W contains the applied patterns, x is the unknown elemental distribution, and ϵ is random uncorrelated noise. In practice, W is a matrix and each of its rows is one of the illumination beam structures. We implement CGI here in an XRF setting by projecting a sequence of structured x-ray patterns onto the thyroid phantom and recording, for each pattern, the integrated iodine fluorescence counts in the selected energy window. The elemental map x is reconstructed by solving the inverse problem using compressive sensing or machine-learning methods, both of which substantially reduce the number of required measurements[28,29]. This approach enables large-area elemental imaging without pencil-beam raster scanning.

A schematic of the experimental setup is shown in Fig. 1b. A binary amplitude mask mounted on a translation stage is positioned upstream of the object. For each realization, a different region of the mask is illuminated, generating a sequence of structured illumination patterns at the sample plane. Fluorescence photons are detected by a silicon-drift-detector (SDD).

The XRF-CGI reconstruction of the image was performed using the Noise2Ghost (N2G) algorithm[27], as detailed in the Methods. The reconstructions presented in this work are two-dimensional and correspond to planar projections of the iodine distribution within the phantom.

To demonstrate human-scale XRF-CGI, Fig. 2a shows the reconstruction of the thyroid insert filled with an iodine solution, embedded within the phantom. For comparison and validation, Fig. 2b and c show the corresponding flat-panel detector (FPD) transmission images of the isolated thyroid insert and of the entire phantom, respectively.

The XRF-CGI reconstruction accurately recovers the thyroid geometry and internal structure and is consistent with the transmission reference. Notably, the thyroid is more clearly delineated in the fluorescence reconstruction than in the transmission image of the full phantom. This difference reflects the intrinsic selectivity of XRF: the fluorescence signal originates primarily from iodine-containing regions, whereas transmission imaging integrates attenuation from all tissues along the beam path. As a result, XRF-CGI naturally suppresses anatomical background and isolates the elemental distribution of interest.

In the upper left region of the thyroid (indicated by arrows in all panels of Fig. 2), a small missing section is visible, especially in the isolated insert and in our XRF-CGI reconstruction. This feature originates from an air bubble trapped during filling of the cavity with the iodine solution. Resolving this localized void highlights the ability of XRF-CGI to capture internal inhomogeneities within a human-scale organ phantom.

Having demonstrated the capability of our method to image the thyroid insert, we next assess its accuracy. Figure 3a shows one-dimensional profiles through the center of the thyroid insert, extracted from the FPD reference image (Fig. 2b) and from the N2G reconstruction. The two profiles exhibit nearly identical shapes of the thyroid lobes, confirming that our approach accurately recovers the iodine distribution. Agreement

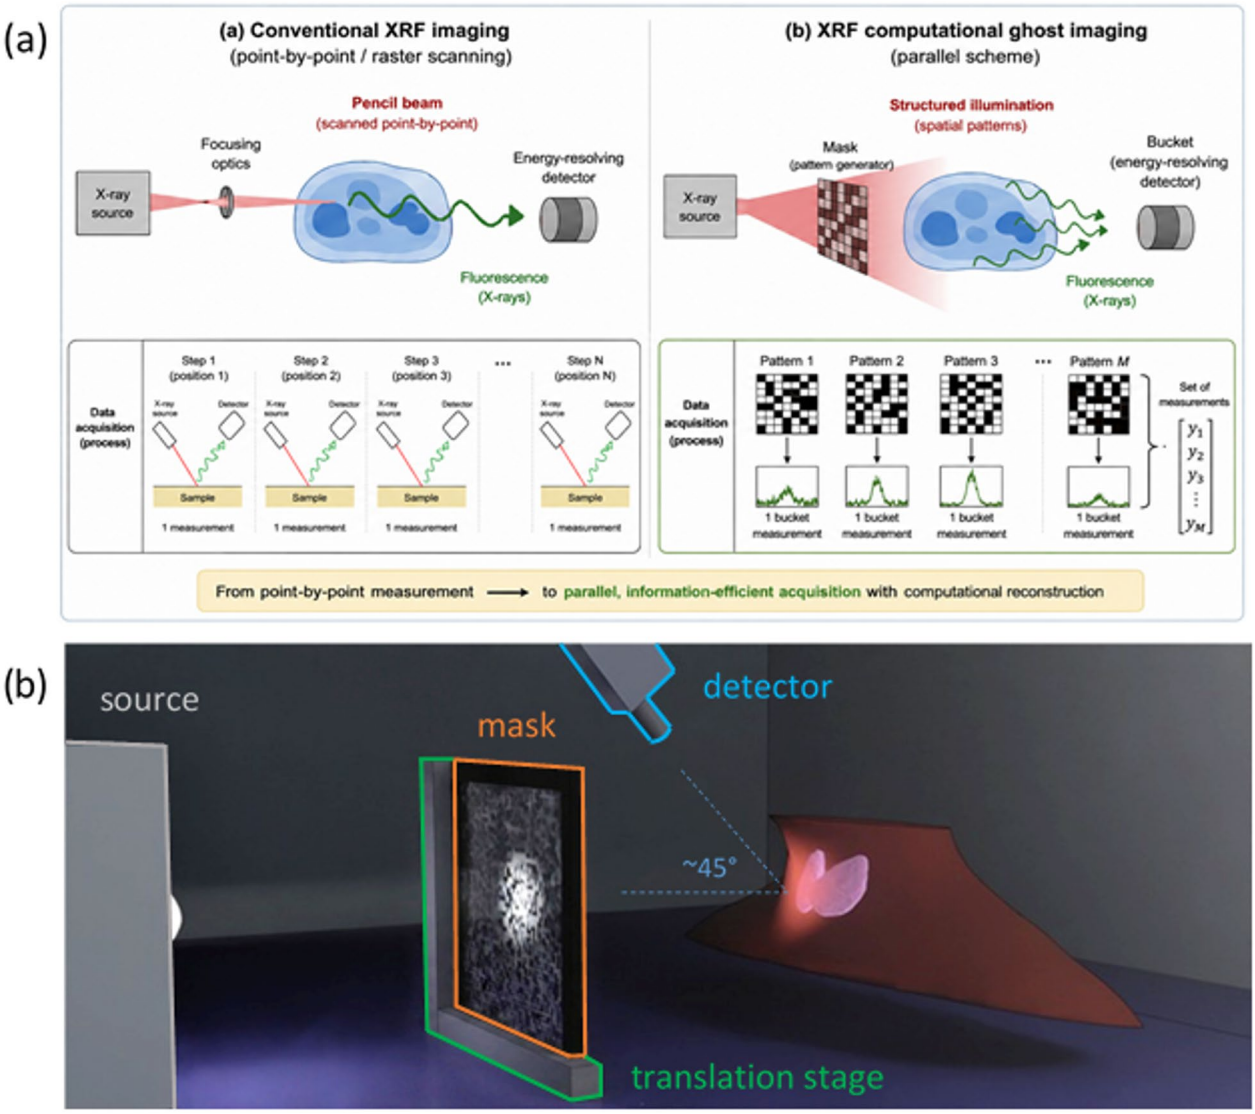


Fig. 1. (**a**) A schematic illustration of the transition from conventional point-by-point XRF raster scanning to parallel acquisitions using the computational ghost imaging method. (**b**) A schematic of the experimental system. Incident x-ray photons pass through the mask and excite the iodine solution within the thyroid phantom, generating characteristic fluorescence emission. The emitted fluorescence photons are collected by an energy-resolving detector, which discriminates the iodine fluorescence lines from the scattered background.

is obtained both on the organ scale and for internal features: the separation between the two thyroid lobes is 45.2 ± 0.2 mm in the reference image and 44.7 ± 4.5 mm in the reconstruction. The air bubble void has a horizontal width of 9.0 ± 0.2 mm in the reference and 8.2 ± 4.5 mm in the reconstruction. Small differences are expected because XRF-CGI is selective to iodine-bearing regions, whereas transmission integrates attenuation over the entire thyroid volume, including surrounding tissue. To provide an independent and quantitative estimate of spatial resolution, Fig. 3b presents the Fourier ring correlation (FRC) curve computed between two N2G reconstructions, each obtained from 3000 independent illumination realizations. Using the standard FRC threshold criterion (see Methods) we obtain an effective spatial resolution of approximately 4.5 mm. This value is consistent with the smallest resolved features in Fig. 2 and with the profile comparison in Fig. 3a, confirming that the reconstruction retains millimeter-scale structural information across the human-scale field of view.

Having established that high-quality images can be demonstrated under these conditions, we next quantify how image quality scales with the number of mask realizations, which directly determines the ionizing-radiation dose in XRF-CGI. In Fig. 4a, b we show the dependence of the Contrast-to-Noise Ratio (CNR) and Mean-squared-error (MSE) on the number of realizations, for data without binning and with a binning ×16 factor. Binning was introduced to accelerate post-acquisition processing and to assess how much spatial information could be discarded before the reconstruction quality noticeably degraded. As seen in Fig. 4a, both cases follow the same overall trend, although the binned data yields higher CNR at low realization counts. In Fig. 4b, the binned data exhibits a somewhat faster decrease in the initial stages. Beyond this initial drop, however, the

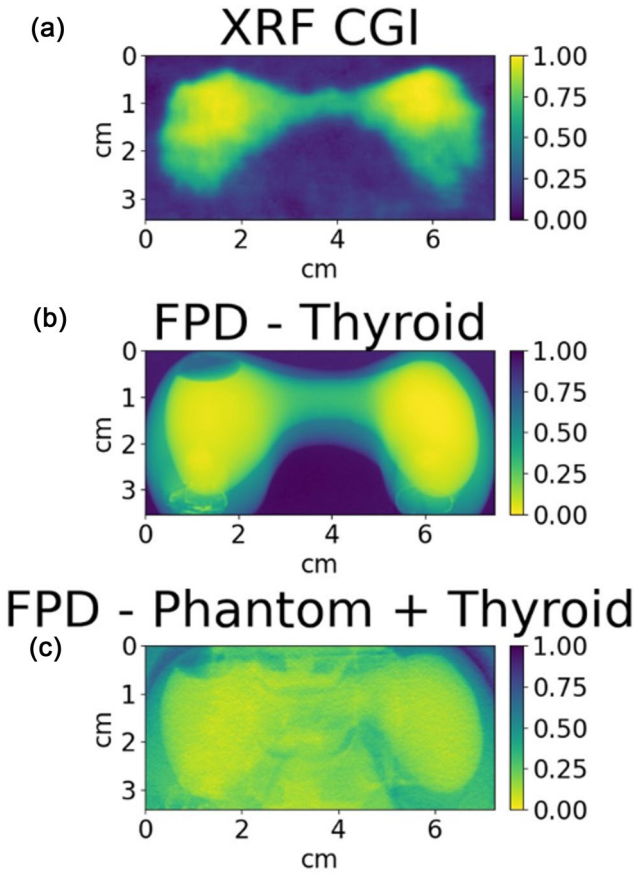


Fig. 2. Images of the thyroid phantom filled with the iodine solution. (**a**) XRF-CGI reconstruction obtained using the self-supervised N2G approach with 5000 illumination realizations, demonstrating faithful recovery of the thyroid insert structure (**b**) Transmission Flat-panel detector (FPD) image of the thyroid insert measured separately from the phantom. (**c**) FPD transmission image of the complete phantom. All images are displayed using normalized intensity scales. For the XRF images, the color bar corresponds to the fluorescence yield, while for the transmission images it corresponds to the inverse of the transmission.

curves converge to similar values, indicating that even with substantial binning the reconstruction quality was not compromised in this experiment.

To evaluate the agreement between the XRF-CGI reconstructions and the corresponding transmission reference images, we show the structural similarity index measure (SSIM) as a function of the increasing number of realizations in Fig. 5. SSIM was selected because it captures similarities in local intensity, contrast, and structural information. The SSIM exhibits an initial rapid increase with the number of realizations, followed by gradual saturation for both ×1 and ×16 binning factors, with overall higher SSIM values for the latter.

It is important to note that a unity SSIM value is not expected in this comparison. The XRF-CGI reconstruction represents the spatial distribution of iodine fluorescence, whereas the FPD image reflects x-ray attenuation through the phantom. The two modalities therefore probe different physical contrast mechanisms: fluorescence selectively maps the iodine-containing region, while transmission integrates attenuation from all materials along the beam path, including the surrounding enclosure of the iodine solution. Nevertheless, because both measurements are two-dimensional projections of the same thyroid phantom geometry, SSIM provides a useful indicator for structural agreement. The observed overall increase in SSIM with realization number therefore supports the conclusion that additional realizations improve not only contrast and pixel-wise convergence, but also the recovery of the thyroid structure.

To further illustrate the dependence of the image quality on the number of realizations, Fig. 6 shows reconstructions obtained with increasing numbers of illumination realizations. Image quality improves monotonically as additional realizations are acquired, with thyroid boundaries and internal features becoming progressively clearer. These examples visually corroborate the CNR, MSE and SSIM trends and highlight the acquisition-quality tradeoff that sets the dose requirements of XRF-CGI.

As seen in Figs. 4 and 5, the maximum number of illumination realizations was 5000, as the image-quality metrics (CNR, MSE, and SSIM) exhibit a clear convergence behavior, with additional measurements above ~4500 realizations providing only marginal improvements in reconstruction quality.

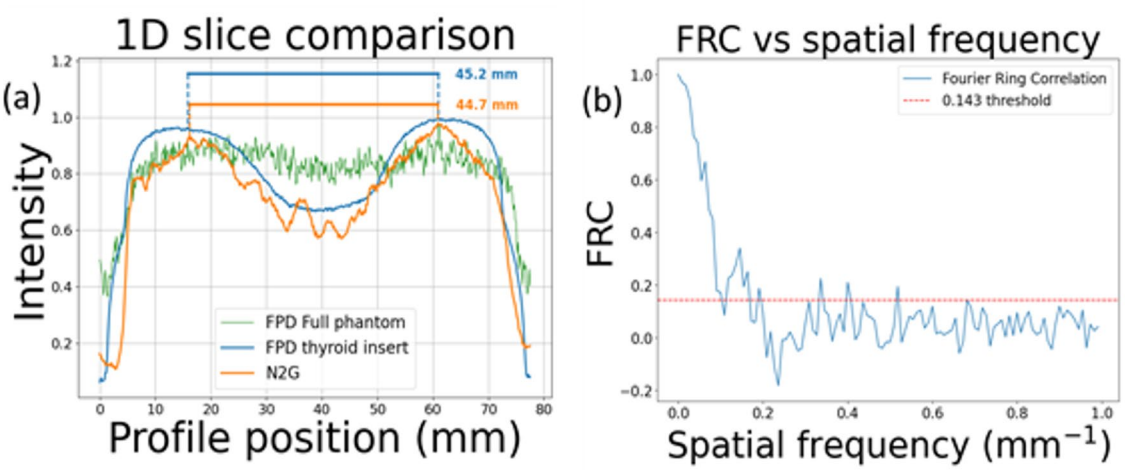


Fig. 3. (**a**) One-dimensional profiles through the center of the thyroid. Shown are lines extracted from the FPD image of the complete phantom, the FPD image of the thyroid insert separated from the phantom, and from a N2G reconstruction. The profiles of the isolated thyroid and the N2G reconstruction exhibit closely matching shapes and similar widths, with lobe separations of 45.2 mm and 44.7 mm, respectively, demonstrating strong agreement between the fluorescence-based reconstruction and the reference transmission measurement. (**b**) Fourier ring correlation (FRC) between two independent ghost-imaging reconstructions (solid line). The dashed red line indicates the 0.143 threshold. The first crossing occurs at 0.109 mm^{-1}, giving a spatial resolution of approximately 4.5 mm.

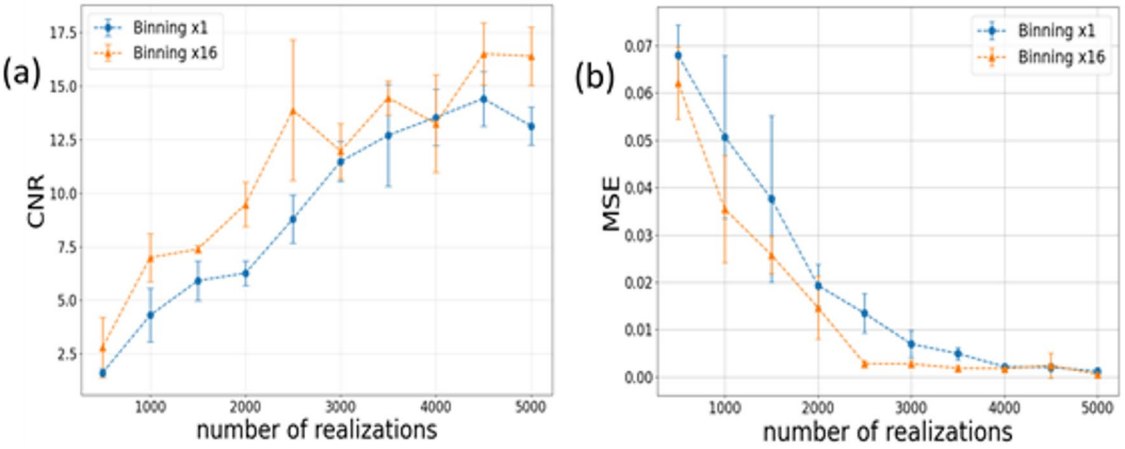


Fig. 4. Dependence of (**a**) Contrast-to-Noise Ratio (CNR) and (**b**) Mean-squared-error (MSE) on the number of realizations. Each data point represents the mean of five independent reconstructions for a given number of realizations, and different colors denote distinct binning factors. Lines are guides to the eye. Error bars indicate the standard deviation.

Discussion

Our human-organ-scale demonstration of XRF-CGI with a thyroid phantom provides a proof-of-concept for retrieving the two-dimensional spatial distribution of non-radioactive markers in clinically sized objects. A particularly clear indicator of reconstruction fidelity is the visibility of the air bubble within the thyroid insert: despite the non-ideal, early-stage implementation, the reconstruction resolves a local inhomogeneity in the marker distribution. This result demonstrates sensitivity to spatial variations in marker concentration at the human-organ scale and supports the potential of XRF-CGI for future functional imaging applications, following substantial improvements in acquisition efficiency and dose performance.

Because non-radioactive elemental markers can, in many cases, be chemically analogous to their radioactive counterparts, XRF-CGI has the potential to provide functional information analogous to that obtained in PET or SPECT while reducing the logistical burden associated with radioisotopes and avoiding the administration of radionuclides. The method, however, still involves exposure to x-ray radiation. The practical utility of such an approach therefore depends critically on achieving clinically relevant acquisition times and radiation doses.

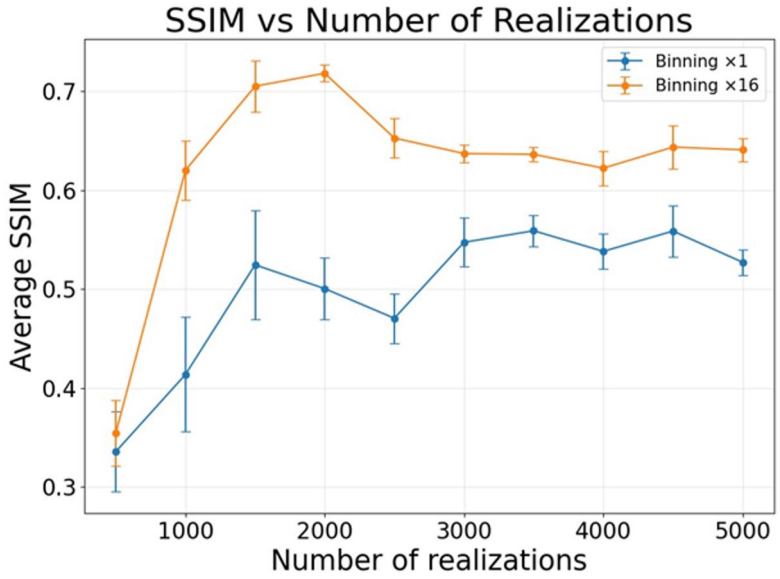


Fig. 5. Structural similarity index measure (SSIM) between the N2G reconstructions and the corresponding FPD transmission reference images as a function of the number of realizations. Each data point represents the mean SSIM of five independent reconstructions. Different colors denote distinct binning factors. Lines are guides to the eye. Error bars indicate the standard deviation.

Translating the present proof-of-concept into a clinically practical modality requires substantial reductions in both acquisition time and radiation dose. We therefore discuss possible improvements based on the current system limitations and on realistic future developments in detector technology and reconstruction methods. In the present implementation we acquired 5000 realizations with 20 s per realization, corresponding to a total acquisition time of nearly 30 h and an estimated entrance dose of approximately 3000 Gy, clearly beyond any clinically practical workflow. The prolonged duration and high dose are primarily dictated by limited fluorescence collection efficiency: the SDD used here has an efficiency of ~ 13% at the iodine Kα fluorescence line (~ 28 keV)[30] and an active area of 17 mm^2. Replacing it with modern CdTe detector arrays, which can approach near unity efficiency at 28 keV[31] together with substantially larger effective collection areas, could significantly increase the collected fluorescence yield, thereby proportionally reducing the required acquisition time and dose. Additional reductions could be obtained by decreasing the detector-phantom distance, which increases the detected rate approximately according to the inverse-square dependence. However, operating at shorter standoff requires careful optimization to control scattered-radiation backgrounds (especially in anatomically heterogeneous environments) and to mitigate dead time through appropriate choices of shaping time and front-end electronics.

A second bottleneck is spectral mismatch between the excitation spectrum and the iodine absorption edge. Our current tube spectrum (80 kVp) is not optimized for iodine excitation, and we estimate that only a small fraction of the emitted photons contributes effectively to fluorescence generation. In addition, geometric constraints introduced to suppress background further reduce the flux reaching the thyroid region. Future implementations could improve spectral efficiency by using excitation spectrum closer to 40 keV together with filtration schemes that suppress low energy non-contributing photons while transmitting those above the iodine K-edge[16]. The achievable gain is expected to depend strongly on the specific system geometry and filtration design. Finally, continued advances in reconstruction methods, particularly those based on machine learning and compressive strategies, are expected to further reduce the number of realizations required for high-fidelity reconstruction, with the achievable reduction depending on the specific reconstruction framework employed. Looking further ahead, quasi-monochromatic sources such as inverse-Compton systems could further increase the spectral efficiency and reduce the administered patient dose[32].

Taken together, these considerations suggest a plausible pathway toward substantially shorter acquisition times and reduced doses while relying on tracers that do not contain radionuclides.

The extension of the present approach to 3D imaging can in principle be achieved tomographically, similarly to other imaging modalities. In the context of the present framework, this would most naturally be achieved using a rotating source-detector gantry. While tomographic acquisition would necessarily increase the total acquisition time and radiation dose relative to a single-projection measurement, the additional spatial redundancy and sparsity constraints available in 3D datasets are known to substantially improve the performance of compressive and sparse-view reconstruction methods, as demonstrated in related GI tomography approaches[33]. These effects may help partially offset the increased acquisition burden associated with angular sampling.

Finally, we emphasize that XRF-CGI is not specific to the thyroid or to iodine. In principle, it can be generalized to other organs and to a broad class of non-radioactive markers, including elements for which no suitable clinical radioisotopes exist today.

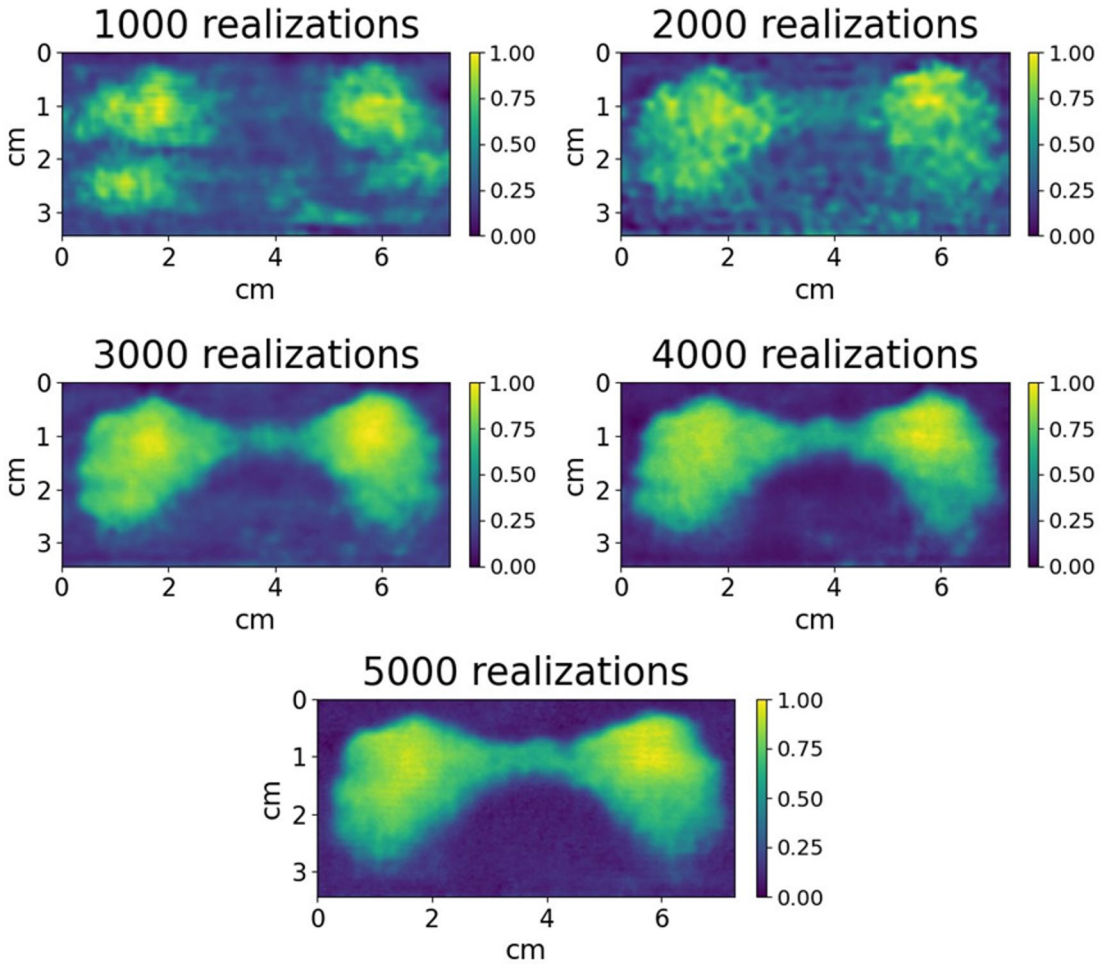


Fig. 6. Dependence of the N2G reconstructions on the number of realizations. Reconstructions of the thyroid phantom are shown for (**a**) 1,000, (**b**) 2,000, (**c**) 3,000, (**d**) 4,000, and (**e**) 5,000 realizations. The two thyroid lobes are discernible even at low realization counts, and image quality improves progressively with additional realizations. All images are displayed using normalized intensity scales.

Methods

Experimental setup

The x-ray source used was a VJ X-ray Ltd. Tungsten tube that operated at 80 kVp and 2 mA. A lead collimator with a 2×4 cm^2 rectangular aperture was placed 25 cm from to the source, limiting irradiation of the downstream phantom to radiation transmitted only through the aperture. The x-ray spot size was approximately matched to the aperture dimensions at a distance of ~ 25 cm with a divergence of ± 4°, and a field of view (FOV) of roughly 4×8 cm^2 at the sample position. No additional beam filtration was applied. The source-to-sample and sample-to-detector distances were ~ 50 cm and ~ 14 cm, respectively. The detector used for fluorescence measurements was an Amptek XR-100SDD FAST silicon drift detector (SDD), with an efficiency of roughly 13% at the iodine Kα fluorescence line of 28 keV, and an active area of 17 mm^2. The detector was positioned at an angle of approximately 45° relative to the incident beam to optimize fluorescence collection and reduce direct transmission to the detector. The exposure time per illumination pattern was 20 s. The geometry and position of the detector were chosen to balance primary fluorescence signal and scattered background. The structured illuminations were generated using a 1.5 mm-thick Ag mask patterned with 0.5×0.5 mm^2 openings over a 16×16 cm^2 area, fabricated by Nano-dimension Ltd, using high-precision multi-layer 3D printing[33,34](open fraction ~ 50%). The object was an RS-542 phantom (RSD Inc.) comprising tissue-equivalent materials and anatomical structures, including a neck section with vertebrae, clavicle, and sternum elements, with overall dimensions of 36 W × 36D × 36 H cm^3. The phantom contained a thyroid insert filled with a 200 mg/mL solution of sodium diatrizoate hydrate, corresponding to approximately 120 mg I/ml. The insert dimensions were 7.5×6 cm^2, representing a human-scale organ.

Reconstruction method

The reconstructions presented were obtained using Noise2Ghost (N2G), a self-supervised convolutional neural network (CNN) method for CGI[27]. In the present work, we used the permutation-augmented N2G formulation. The acquired realizations are permuted, and each permutation is partitioned into independent subsets containing

different groups of illumination patterns and corresponding bucket values. Lower quality sub-reconstructions are first generated independently from each subset using the forward model. Although these sub-reconstructions represent the same object, they contain different noise contributions and missing-realization artifacts. During training, the CNN is optimized such that the reconstruction generated from one subset is consistent with bucket measurements of the other independent subsets. Consequently, subset-specific noise is not reinforced through the training loss, whereas image features that are consistent across the independent sets are preserved. We consider the forward model we introduced in Eq. (1). N2G seeks a neural network parametrization $\widehat{\theta}$ that minimizes the following expression:

$$\widehat{\theta} = \underset{\theta}{\operatorname{argmin}} \frac{1}{2} \sum_{p,k} ||WN_{\theta}(x_{p,k}) - y||_2^2 + \lambda R(N_{\theta}(x_{p,k})), \tag{2}$$

where p and k index the number of permutations and splits, respectively. For a fixed permutation and target subset, the expected prediction error will then be:

$$E_{b,i\neq k} ||W_i N_{\theta}(x_k) - b_i||_2^2 + E_{\epsilon,i\neq k} ||\epsilon_i||_2^2 , \tag{3}$$

where the first term is the supervised reconstruction error and does not depend on the noise. The reconstructed image is obtained by averaging the network outputs and splits:

$$\widehat{x} = \frac{1}{PK} \sum_{k} N_{\widehat{\theta}}(x_{p,k}), \tag{4}$$

where P and K denote the corresponding permutations and splits, respectively. This averaging further reduces subset-specific fluctuations and yields the final N2G reconstruction used in the present work. The theoretical derivation, network schematic, and benchmarking of the N2G method are reported in Ref[27]. Training hyperparameters specific to the present work are reported in the Supplementary Information.

Spatial binning was applied only to the structured illumination patterns, while the measured bucket fluorescence signals were kept unchanged. This was done by averaging non-overlapping N×N pixel blocks over each pattern.

Evaluation metrics

To quantify reconstruction quality, we used MSE, CNR, FRC and SSIM.

For the MSE, the deviations between corresponding pixels in the reconstructed and reference images are computed as:

$$MSE = \frac{1}{N} \sum_{j=1}^{N} (I_j - R_j)^2, \tag{5}$$

where I_j and R_j denote the intensity values of the reconstructed and reference image pixels, respectively, and **N** is the number of pixels. Because an external ground-truth image is unavailable, the reference image is defined as the average of five independent N2G reconstructions obtained using 5000 realizations. The MSE therefore measures convergence toward this stable reference.

To quantify contrast, we computed the CNR, which accounts for the signal and background levels as well as signal variability:

$$CNR = \frac{|S_A - S_B|}{\sigma_A} \tag{6}$$

where S_A and S_B are the mean intensities in the signal and background regions, respectively, and σ_A is the standard deviation within the signal region. We normalize by σ_A rather than the background standard deviation because the background level decreases strongly with the number of realizations, which would otherwise artificially inflate the CNR at large realization numbers.

FRC is used as an independent and objective estimate of spatial resolution. Two reconstructions obtained from statistically independent sets of realizations are Fourier-transformed, and their normalized cross-correlation is computed over concentric rings in Fourier space as a function of spatial frequency r:

$$FRC(r) = \frac{\sum_{i\in r} \boldsymbol{F}_1(\boldsymbol{r_i}) \boldsymbol{F}_2(\boldsymbol{r_i})^*}{\sqrt{\sum_{i\in r} |\boldsymbol{F}_1(\boldsymbol{r_i})|^2 \sum_{i\in r} |\boldsymbol{F}_2(\boldsymbol{r_i})|^2}} \tag{7}$$

where F_1 and F_2 are the Fourier transforms of the two reconstructions, and the summation is performed over all Fourier pixels r_i on a ring of radius r. The resolution is defined as the spatial frequency at which the FRC curve crosses the 0.143 threshold, a commonly used empirical criterion[35]. Spatial resolution was reported using the half-period convention, calculated as $d = \frac{1}{2f_c}$, where f_c is the threshold-crossing spatial frequency. The uncertainty in the reconstructions presented here was estimated based on the system's lateral resolution (~ 4.5 mm), as determined by the FRC analysis, and used as a conservative proxy for the measurement precision.

Structural similarity index measure (SSIM) was used as an additional image quality metric to quantify the agreement between the XRF-CGI reconstructions and the corresponding FPD transmission images. Prior to the

calculation, both the reconstructed fluorescence images and the FPD reference images were normalized. SSIM compares two images by evaluating local similarity in luminance, contrast, and structure. For two images x and y, SSIM is defined as:

$$SSIM\,(x,y) = \frac{\left(2\mu_x\mu_y + C_1\right)\left(2\sigma_{xy} + C_2\right)}{\left(\mu_x^2 + \mu_y^2 + C_1\right)\left(\sigma_x^2 + \sigma_y^2 + C_2\right)} \tag{8}$$

where μ_x, μ_y are the corresponding mean intensities, σ_x^2, σ_y^2 are the corresponding variances, σ_{xy} is the local covariance, and C_1, C_2 are stabilization constants. SSIM is dimensionless, with larger values indicating higher structural similarity.

Data availability

The representative raw datasets supporting the findings of this study have been deposited in a Zenodo repository and will be made publicly available upon publication.

Code availability

The Noise2Ghost code implementation is publicly available at: https://github.com/CEA-MetroCarac/noise2ghost.

Received: 23 March 2026; Accepted: 13 July 2026
Published online: 17 July 2026

References

1. Cole, E., Stewart, M., Littich, R., Hoareau, R. & Scott, P. Radiosyntheses using fluorine-18: the art and science of late stage fluorination. *Curr. Top. Med. Chem.* **14**, 875 (2014).
2. Manabe, O. et al. Radiopharmaceutical tracers for cardiac imaging. *J. Nucl. Cardiol.* **25**, 1204 (2018).
3. Werner, R. A. et al. The changing face of nuclear cardiology: guiding cardiovascular care toward molecular medicine. *J. Nucl. Med.* **61**, 951 (2020).
4. McAteer, M. A. et al. Translation of PET radiotracers for cancer imaging: recommendations from the National Cancer Imaging Translational Accelerator (NCITA) consensus meeting. *BMC Med* **23** (2025).
5. Watabe, T. et al. Recent advances in theranostics and oncology PET: emerging radionuclides and targets. *Ann. Nucl. Med.* **39**, 909 (2025).
6. Oldan, J. D., Pomper, M. G., Werner, R. A., Higuchi, T. & Rowe, S. P. The cutting edge: Promising oncology radiotracers in clinical development. *Diagn. Interv. Imaging* **105**, 400 (2024).
7. Frush, D. P. et al. Recurrent medical imaging exposures for the care of patients: one way forward. *Eur. Radiol.* **34**, 6475 (2024).
8. Hosono, M. et al. Cumulative radiation doses from recurrent PET–CT examinations. *Br. J. Radiol.* https://doi.org/10.1259/bjr.20210388 (2021).
9. Al-Ibraheem, A. et al. Implementation of radiotheranostics: challenges, barriers, and IAEA-driven strategies for sustainable access. *Semin Nucl. Med.* **55**, 1032 (2025).
10. Hansen, S. B. & Bender, D. Advancement in production of radiotracers. *Semin Nucl. Med.* **52**, 266 (2022).
11. Voss, S. D. SPECT/CT, PET/CT and PET/MRI: Oncologic and infectious applications and protocol considerations. *Pediatr. Radiol.* **53**, 1443 (2023).
12. Willowson, K. P. & Bailey, D. L. Evolving SPECT-CT technology. *Br J. Radiol* **97** (2024).
13. Surti, S. & Karp, J. S. Update on latest advances in time-of-flight PET. *Phys. Med.* **80**, 251 (2020).
14. Cheng, Z., Chen, P. & Yan, J. A review of state-of-the-art resolution improvement techniques in SPECT imaging. *EJNMMI Phys.* https://doi.org/10.1186/s40658-025-00724-9 (2025).
15. Singh, M. K. A review of digital PET-CT technology: comparing performance parameters in SiPM integrated digital PET-CT systems. *Radiography* **30**, 13 (2024).
16. Staufer, T. & Grüner, F. Review of development and recent advances in biomedical x-ray fluorescence imaging. *Int. J. Mol. Sci.* **24**, 10990 (2023).
17. Larsson, J. C. et al. High-spatial-resolution x-ray fluorescence tomography with spectrally matched nanoparticles. *Phys. Med. Biol.* **63**, 164001 (2018).
18. Shugar, A. N. & Mass, J. L. Handheld x-ray fluorescence spectrometry for art and archaeology. **1**, 17–36 (2012).
19. Marguí, E., Queralt, I. & Hidalgo, M. Application of x-ray fluorescence spectrometry to determination and quantitation of metals in vegetal material. *TrAC Trends Anal. Chem.* **28**, 362 (2009).
20. Hoffer, P. B., Bernstein, J. & Gottschalk, A. Fluorescent techniques in thyroid imaging. *Semin. Nucl. Med.* **1**, 379 (1971).
21. Romero, I. O., Fang, Y., Lun, M. & Li, C. X-ray fluorescence computed tomography (XFCT) imaging with a superfine pencil beam x-ray source. *Photonics* **8**, 236 (2021).
22. Manohar, N., Reynoso, F. J., Diagaradjane, P., Krishnan, S. & Cho, S. H. Quantitative imaging of gold nanoparticle distribution in a tumor-bearing mouse using benchtop X-ray fluorescence computed tomography. *Sci. Rep.* **6**, 22079 (2016).
23. Grüner, F. et al. Localising functionalised gold-nanoparticles in murine spinal cords by X-ray fluorescence imaging and background-reduction through spatial filtering for human-sized objects. *Sci. Rep.* **8**, 16561 (2018).
24. Pushie, M. J. et al. X-ray fluorescence microscopy methods for biological tissues. *Metallomics* **14**, mfac032 (2022).
25. Klein, Y., Sefi, O., Schwartz, H. & Shwartz, S. Chemical element mapping by x-ray computational ghost fluorescence. *Optica* **9**, 63 (2022).
26. Shapiro, J. H. Computational ghost imaging. *Phys. Rev. A* **78**, 061802 (2008).
27. Manni, M., Karpov, D., Batenburg, K. J., Shwartz, S. & Viganò, N. Noise2Ghost: Self-supervised deep convolutional reconstruction for ghost imaging. Opt. Express 34, 24787–24804 (2026) (2025).
28. Zhang, H., Xia, Y. & Duan, D. Computational ghost imaging with deep compressed sensing. *Chin. Phys. B.* **30**, 124209 (2021).
29. Duarte, M. F. et al. Single-pixel imaging via compressive sampling. *IEEE Signal Process. Mag.* **25**, 83 (2008).
30. https://www.amptek.com/internal-products/obsolete-products/sdd-x-ray-detectors-for-xrf/xr-100sdd-silicon-drift-detector
31. https://www.amptek.com/products/x-ray-detectors/cdte-x-ray-and-gamma-ray-detectors/cdte-x-ray-and-gamma-ray-detector
32. Kulpe, S. et al. Spectroscopic imaging at compact inverse Compton x-ray sources. *Phys. Med.* **79**, 137 (2020).
33. Ben-Yehuda, A. et al. High-resolution computed tomography with scattered x-ray radiation and a single pixel detector. *Commun. Eng.* **3**, 184 (2024).
34. Nano Dimension. Industrial 3D printing of electronics for manufacturing and AME. accessed (2026). http://www.nano-di.com

35. van Heel, M. & Schatz, M. Fourier shell correlation threshold criteria. *J. Struct. Biol.* **151**, 250 (2005).

Acknowledgements
This study was funded by the Israel Science Foundation (ISF) (847/21). The funder played no role in study design, data collection, analysis and interpretation of data, or the writing of this manuscript.

Author contributions
E.L. performed all experiments and led the data analysis. R.H.S. conceived the project and provided feedback on the manuscript. N.V. developed the Noise2Ghost algorithm used for the image reconstructions. S. Shwartz supervised the research. E.L. and S. Shwartz wrote the manuscript. All authors contributed to the work presented here and to the final paper.

Declarations

Competing interests
The authors declare no competing interests.

Additional information
Supplementary Information The online version contains supplementary material available at https://doi.org/10.1038/s41598-026-62593-1.

Correspondence and requests for materials should be addressed to E.L. or S.S.

Reprints and permissions information is available at www.nature.com/reprints.

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

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

© The Author(s) 2026