



OPEN ACCESS

EDITED BY

Omneya Attallah,
Arab Academy for Science, Technology and
Maritime Transport (AASTMT), Egypt

REVIEWED BY

Oliver Zettinig,
ImFusion GmbH, Germany
Bishesh Khanal,
NepAI Applied Mathematics and Informatics
Institute for Research, Nepal
RabinAdhikari,
NepAI Applied Mathematics and Informatics
Institute for Research, Nepal, in collaboration
with reviewer BK

*CORRESPONDENCE

Abdoul Aziz Amadou
✉ abdoul.a.amadou@kcl.ac.uk

[†]Work done while at Siemens Healthineers,
Digital Technology and Innovation, Princeton,
NJ, United States

RECEIVED 15 February 2024

ACCEPTED 19 July 2024

PUBLISHED 13 August 2024

CITATION

Amadou AA, Peralta L, Dryburgh P, Klein P,
Petkov K, Housden RJ, Singh V, Liao R,
Kim Y-H, Ghesu FC, Mansi T, Rajani R, Young A
and Rhode K (2024) Cardiac ultrasound
simulation for autonomous ultrasound
navigation.
Front. Cardiovasc. Med. 11:1384421.
doi: 10.3389/fcvm.2024.1384421

COPYRIGHT

© 2024 Amadou, Peralta, Dryburgh, Klein,
Petkov, Housden, Singh, Liao, Kim, Ghesu,
Mansi, Rajani, Young and Rhode. This is an
open-access article distributed under the
terms of the [Creative Commons Attribution
License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or
reproduction in other forums is permitted,
provided the original author(s) and the
copyright owner(s) are credited and that the
original publication in this journal is cited, in
accordance with accepted academic practice.
No use, distribution or reproduction is
permitted which does not comply with
these terms.

Cardiac ultrasound simulation for autonomous ultrasound navigation

Abdoul Aziz Amadou^{1,2*}, Laura Peralta¹, Paul Dryburgh¹,
Paul Klein³, Kaloian Petkov³, R. James Housden¹, Vivek Singh³,
Rui Liao^{3†}, Young-Ho Kim³, Florin C. Ghesu⁴, Tommaso Mansi^{3†},
Ronak Rajani¹, Alistair Young¹ and Kawal Rhode¹

¹Department of Surgical & Interventional Engineering, King's College London, School of Biomedical Engineering & Imaging Sciences, London, United Kingdom, ²Digital Technology and Innovation, Siemens Healthcare Limited, Camberley, United Kingdom, ³Siemens Healthineers, Digital Technology and Innovation, Princeton, NJ, United States, ⁴Siemens Healthineers AG, Digital Technology and Innovation, Erlangen, Germany

Introduction: Ultrasound is well-established as an imaging modality for diagnostic and interventional purposes. However, the image quality varies with operator skills as acquiring and interpreting ultrasound images requires extensive training due to the imaging artefacts, the range of acquisition parameters and the variability of patient anatomies. Automating the image acquisition task could improve acquisition reproducibility and quality but training such an algorithm requires large amounts of navigation data, not saved in routine examinations.

Methods: We propose a method to generate large amounts of ultrasound images from other modalities and from arbitrary positions, such that this pipeline can later be used by learning algorithms for navigation. We present a novel simulation pipeline which uses segmentations from other modalities, an optimized volumetric data representation and GPU-accelerated Monte Carlo path tracing to generate view-dependent and patient-specific ultrasound images.

Results: We extensively validate the correctness of our pipeline with a phantom experiment, where structures' sizes, contrast and speckle noise properties are assessed. Furthermore, we demonstrate its usability to train neural networks for navigation in an echocardiography view classification experiment by generating synthetic images from more than 1,000 patients. Networks pre-trained with our simulations achieve significantly superior performance in settings where large real datasets are not available, especially for under-represented classes.

Discussion: The proposed approach allows for fast and accurate patient-specific ultrasound image generation, and its usability for training networks for navigation-related tasks is demonstrated.

KEYWORDS

ultrasound, Monte-Carlo integration, path tracing, simulation, echocardiography

1 Introduction

Ultrasound (US) is pivotal in the diagnosis, treatment and follow-up of patients in several medical specialities such as cardiology, obstetrics, gynaecology and hepatology. However, the quality of acquired images varies greatly depending on operators' skills, which can impact diagnostic and interventional outcomes (1).

Providing guidance or automation for the image acquisition process would allow for reproducible imaging, increase both the workflow efficiency and throughput of echo

departments and improve access to ultrasound examinations. This requires an intelligent system, capable of acquiring images by taking into consideration the high variability of patient anatomies.

Several works are investigating US acquisition automation but commercially available systems do not go beyond teleoperated ultrasound (2). Recent research towards autonomous navigation has used imitation learning (3) and deep reinforcement learning (4, 5). While these methods achieve varying degrees of success, they struggle to adapt to unseen anatomies, can only manage simple scanning patterns or are tested on small datasets.

The main advantage of a simulation environment is the ability to generate views that occur when operators navigate to a given standard view or anatomical landmark but are not saved in clinical routine. These datasets, which we call navigation data, can also contain imaging artefacts (e.g. shadowing caused by ribs). Hence, recent ultrasound image synthesis methods using neural networks (6, 7) face significant challenges in generating these views due to the necessity of comprehending ultrasound physics and the unavailability of large-scale datasets of complete ultrasound acquisitions. Besides, learning-based approaches for navigation (5) require a large number of images for training, including non-standard views, which are not available in classical ultrasound training datasets.

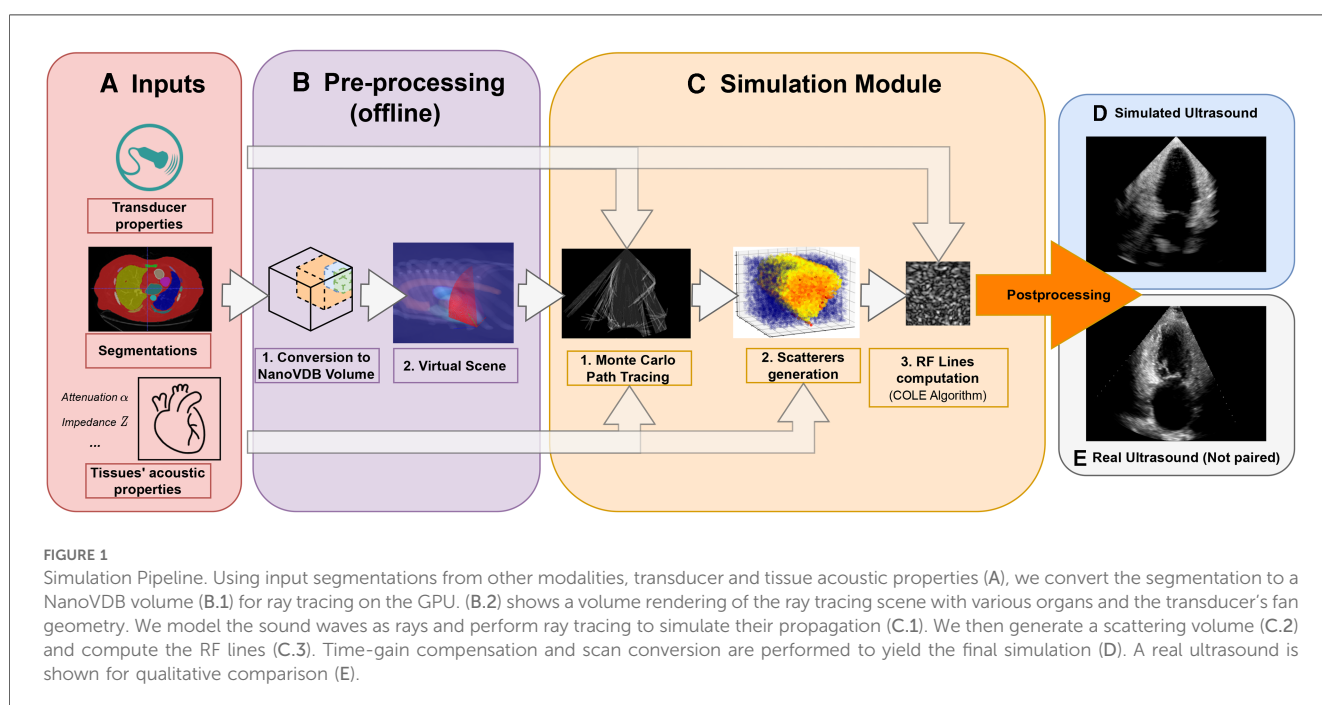
Using a simulation environment to train such a system would have several benefits. The trained model could learn while being exposed to a varying range of anatomies and image qualities, hence improving its robustness, and the training could be done safely, preventing the wear of mechanical components and potential injuries. This simulation environment should be: (1) Fast, to enable the use of state-of-the-art reinforcement learning algorithms. (2) Reproduce patients' anatomies with high fidelity. (3) Recreate attenuation artefacts such as shadowing. Moreover, exposing the system to a wide range of anatomies

requires large-scale data generation capabilities, meaning the pre-processing of data must be streamlined.

This paper presents an ultrasound simulation pipeline using Graphical Processing Unit (GPU) based ray tracing on NVIDIA OptiX (8), capable of generating US images in less than a second. By combining networks capable of segmenting a wide range of tissues and a volumetric data representation, we overcome the scene modelling limitations of previous mesh-based simulation methods, enabling efficient processing of numerous datasets from different modalities. Our pipeline, described in Figure 1 takes as input segmentations of the organs of interest and, coupled with user-defined transducer and tissue properties, generates a simulated US by combining Monte Carlo path tracing (MCPT) and convolutional approaches.

Our contributions are the following:

- Our pipeline is able to generate images from a large number of datasets from other modalities. Using an efficient GPU volumetric representation that allows for the modelling of arbitrary patient anatomies, and a Monte Carlo path tracing algorithm, we are able to synthesize more than 10,000 images per hour using a NVIDIA Quadro K5000 GPU. Furthermore, we demonstrate scalability by generating images from 1,000 CT patient datasets in our experiments. In contrast, existing ray tracing methods limit their experiments to datasets one or two orders of magnitude smaller.
- We extensively validate the ability of our pipeline to preserve anatomical features through a phantom experiment by looking at distances and contrast between structures. Ultrasound image properties are further assessed by looking at first-order speckle statistics.
- We demonstrate the usability of our pipeline in training neural networks for transthoracic echocardiography (TTE) standard



view classification, a task critical in ultrasound navigation guidance. The neural networks were initially pre-trained on synthetic images and subsequently fine-tuned using varying amounts of real data. With around half of the real samples, fine-tuned networks reach a performance level comparable to those trained with all the real data. We also report an improved classification performance when using pre-trained networks, particularly for under-represented classes.

This paper is organised in the following way: In [Section 2.1](#), we provide an overview of relevant ultrasound simulation methods and highlight their limitations in terms of suitability as simulation environments. The next subsections in [Section 2](#) detail our simulation implementation. Experimental results using a virtual phantom and a view classification network are shown in [Section 3](#). This is followed by a discussion and a conclusion.

2 Methods

2.1 Related work

Early methods were attempting to simulate the US image formation process by solving the wave equation using various strategies ([9–13](#)). While being accurate, these methods take a substantial amount of time to generate images [in the order of several minutes to hours ([9–12](#))], which is not scalable for large-scale training.

The COLE Algorithm developed by Gao et al. ([14](#)) is at the core of Convolutional Ray Tracing (CRT) methods. This approach allows for a fast simulation of ultrasound images with speckle noise by convolving a separable Point-Spread Function (PSF) with a scatterer distribution. Methods in ([15–17](#)) replace the ray casting by ray tracing and combine it with the COLE algorithm to simulate images on the GPU. These methods follow a similar methodology where the input volumes are segmented and acoustic properties from the literature are assigned to each tissue. Scatterers amplitude are hyperparameters chosen such that the generated ultrasounds look plausible. Ray tracing is used to model large-scale effects at boundaries (reflection and refraction) and attenuation within tissue. Finally, the COLE algorithm is applied to yield the final image. The method developed in Mattausch et al. ([17](#)) distinguishes itself by employing Monte-Carlo Path Tracing (MCPT) to approximate the ray intensity at given points by taking into account contributions from multiple directions.

CRT methods enable fast simulations and the recreation of imaging artefacts. Methods in ([15](#), [17](#)) both make use of meshes to represent the boundaries between organs. However, using meshes comes with a set of issues as specific pre-processing and algorithms are needed to manage overlapping boundaries. This can lead to the erroneous rendering of tissues, hence limiting the type of scene that can be modelled, as reported in Mattausch et al. ([17](#)). A further limitation of CRT methods lies in tissue parameterization, where scatterers belonging to the same tissue have similar properties, preventing the modelling of fine-tissue variations, and thus limiting the realism of the images.

Another line of work generates synthetic ultrasound images by directly sampling scatterers' intensities from template ultrasound images and using electromechanical models to apply cardiac motion ([18](#), [19](#)). These are different from our line of work as they require pre-existing ultrasound recordings for a given patient, while we generate synthetic images from other modalities, which also enables us to simulate different types of organs other than the heart.

Finally, as deep learning has become increasingly popular, the field shifted towards the use of generative adversarial networks (GAN) or diffusion models for image synthesis. These generative models have been used in several ways for image simulation: Either for generating images directly from segmentations ([6](#), [7](#), [20](#)), calibrated coordinates ([21](#)), or for improving the quality of images generated from CRT simulators ([22–24](#)). However, using GANs comes with several challenges: For instance, authors in Hu et al. ([21](#)) report mode collapse when generating images for poses where training data was not available and authors in Gilbert et al. ([6](#)) report hallucination of structures if anatomical structures are not equally represented in datasets. This suggests generative neural networks would struggle in generating out-of-distribution views or with image artefacts such as shadowing. This would be problematic for ultrasound navigation guidance as out-of-distribution views are frequently encountered before reaching a desired standard view.

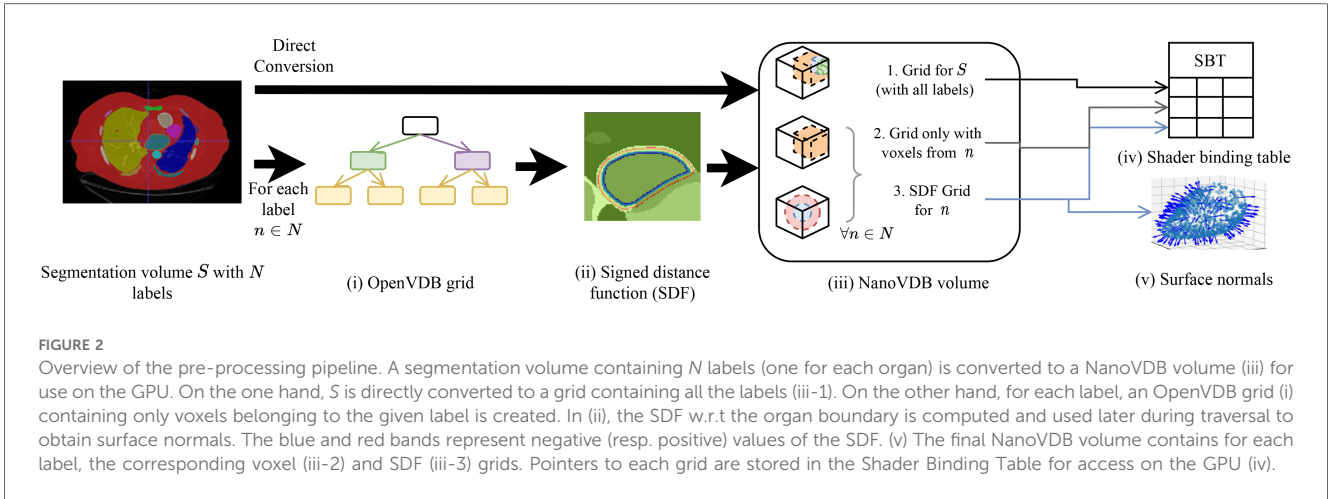
Methods taking as input low-quality images from CRT simulators seem the most promising, but several works report issues in preventing the GANs from distorting the anatomy ([24](#)) or introducing unrealistic image artefacts ([22](#)). While CRT methods are limited in realism, they match our requirements (speed, artefacts recreation, anatomical fidelity through accurate geometry) to train navigation/guidance algorithms.

2.2 Pre-processing pipeline

This section presents our novel pre-processing pipeline, shown in [Figure 1](#), which enables large-scale data generation by avoiding technical pitfalls caused by the use of meshes ([17](#)), thus allowing us to model any anatomy. Besides, the use of segmentations is essential to implement constraints on the environment for navigation tasks.

Input volumes ([Figure 1A](#)) are segmentations obtained from either CT or Magnetic Resonance Imaging (MRI) datasets, which are processed by a four chamber ([25](#)) and multi-organ segmentation algorithm ([26](#)). The segmentation output contains all the structures relevant for echocardiography, e.g. individual ribs, sternum, heart chambers, aorta, and lungs.

During ray tracing, voxels need to be accessed at random. The access speed is highly dependent on the memory layout of the data. This problem has been addressed by OpenVDB ([27](#)) with its optimized B+ tree data structure and by its compacted, read-only and GPU-compatible version, NanoVDB ([28](#)). Data in Open/NanoVDB are stored in grids. These grids can be written together into a single file, which we call an Open/NanoVDB



volume. We convert the segmentation volumes into NanoVDB volumes (Figure 1B) as described below.

A detailed overview of the pre-processing pipeline is shown in Figure 2. Firstly, the segmentation volume with all labels is converted to a NanoVDB grid (Figure 2iii-1). This grid is used during ray tracing to access a label associated with a given voxel. Then, for each label in the segmentation volume, a narrow-band signed distance function (SDF) is computed such that the distance from voxels in the neighbourhood of the organ to its boundary is known (Figure 2ii). Blue (resp. red) bands represent the voxels with negative (resp. positive) distance to that boundary, i.e. inside (resp. outside) it. The SDF grids are written to the output volume (Figure 2iii-3) and are later used during traversal to compute smooth surface normals by looking at the SDF's gradient (Figure 2v).

A separate grid containing only the voxels associated with the current organ is also saved (Figure 2iii-2) in the output volume. Hence, the final NanoVDB volume (Figure 2iii) contains the original voxel grid and, for each label, two grids: the SDF grid as well as the voxel grid. In practice, the pre-processing takes less than five minutes per volume and we use several worker processes to perform this task on multiple volumes in parallel.

2.3 Scene setup

Similarly to previous work (15–17), the sound wave is modelled as a ray. The simulation is done using OptiX (8), which is a CUDA/C++ general-purpose ray tracing library providing its users with fast intersection primitives on the GPU. The previously generated NanoVDB volume is loaded and the voxel grids corresponding to each label (Figure 2iii-2) are represented as Axis-Aligned Bounding Boxes (AABB) which are grouped together to create the Acceleration Structure (AS) used by OptiX to compute intersections. We assign acoustic properties from the literature (29) to each organ. A summary of all the assigned properties is listed in Table 1. Values for μ_0 , μ_1 , σ_0 are detailed in Table A1 in the Appendix. To retrieve data during traversal, OptiX uses a Shader Binding Table (SBT). We populate it with tissue properties, pointers to the organs' SDFs and a pointer to the original voxel

TABLE 1 List of properties assigned to tissues.

Property	Description	Domain
Impedance (Z)	Tissue-specific acoustic impedance in $\text{kg}/(\text{m}^2 \cdot \text{s})$	–
Attenuation coefficient (α)	Tissue-specific attenuation in $\text{dB}/(\text{cm} \cdot \text{Hz})$	–
Sound speed (c)	Sound speed in a given tissue, in $\text{m} \cdot \text{s}^{-1}$	–
μ_0, σ_0, μ_1	Scatterer distribution parameter, from (15). μ_0, σ_0 control the scatterer amplitude while μ_1 controls the probability of a scatterer being generated	$\mu_0, \sigma_0 \in [0, 1]$
τ	Coefficient used to specify whether a reflection is more diffuse or specular, as in (15)	$\tau \in [0, 3]$
γ	Coefficient used to amplify small reflections, as in (15)	$\gamma \in [-2, 2]$

Domain values are indicated for hyperparameters. Please refer to Tables A1, A2 in the Appendix for the impedance and attenuation coefficient values per tissue.

TABLE 2 List of parameters used to configure the transducer.

Property	Description
Center frequency	Transducer center frequency (in Hz)
Sampling frequency	Signal sampling frequency (in Hz)
Element width	Width (in mm) of an element
Element height	Height (in mm) of an element
Kerf	Spacing between two elements (in mm)
Number of elements	Number of elements making up the matrix array
Scan geometry	Type of scan geometry (e.g. linear, phased)

grid (Figure 2iv). Finally, a virtual transducer is positioned in the scene. Transducer parameters are listed in Table 2.

2.4 Simulation module

The goal of the simulation module (Figure 1C) is to generate view-dependent US images. This module is made of two parts.

The first part performs the ray tracing using OptiX. The goal of this module is to model large-scale effects (reflections, refractions and attenuation). This is done by computing, for each point along a scanline, the intensity I sent back to the transducer. The

second part generates the US image by convolving the point spread function (PSF) with the scatterer distribution while taking into account the corresponding intensity $I(l)$ along the scanline.

2.4.1 Background

2.4.1.1 Ultrasound physics

Here we first describe the phenomena happening during ray propagation: The wave loses energy due to attenuation following $I(l) = I_0 e^{-lf\alpha}$, with I_0 the initial wave intensity and l the distance travelled in a given medium with attenuation α at frequency f . When it reaches a boundary, it is partially reflected and transmitted depending on the difference in impedance between the two media. The reflection and transmission coefficients R and T are written following Equations 1, 2:

$$R(Z_1, Z_2, \theta_1, \theta_2) = \left(\frac{Z_2 \cos(\theta_2) - Z_1 \cos(\theta_1)}{Z_2 \cos(\theta_2) + Z_1 \cos(\theta_1)} \right)^2 \quad (1)$$

$$T(Z_1, Z_2, \theta_1, \theta_2) = 1 - R(Z_1, Z_2, \theta_1, \theta_2) \quad (2)$$

$$\cos(\theta_1) = \vec{n} \cdot \vec{v} \quad (3)$$

$$\cos(\theta_2) = \sqrt{1 - \left(\frac{Z_1}{Z_2} \right)^2 (1 - \cos^2(\theta_1))} \quad (4)$$

Where Z_1 and Z_2 are the impedances of the media at the boundary. θ_1 computed following Equation 3, is the angle between the incident ray \vec{v} and the surface normal \vec{n} . Finally, θ_2 is the refracted angle and computed following Equation 4.

2.4.1.2 Rendering equation

When the wave propagates in tissue, it can encounter several boundaries and bounce multiple times, depending on the scene geometry. Hence, retrieving the total intensity at a given point P requires taking into account contributions coming from multiple directions. The field of computer graphics has faced similar challenges to compute global illumination.

We take inspiration from the rendering equation (30):

$$L_{P \rightarrow v} = O_{P \rightarrow v} + \int_{\Omega} f_{P, \omega \rightarrow v} L_{P \rightarrow \omega} \cos(\theta) d\omega \quad (5)$$

where:

- Ω is the surface hemisphere around the surface normal at point P .
- $L_{P \rightarrow v}$ is the amount of light leaving point P in direction v .
- $O_{P \rightarrow v}$ is the light emitted at P in direction v .
- $f_{P, \omega \rightarrow v}$ is a Bidirectional Scattering Distribution Function (BSDF) giving the amount of light sent back by a given material in direction v when it receives light from direction ω at point P .
- $L_{P \rightarrow \omega}$ is the amount of light received by P in direction ω .
- Finally, θ is the angle between the surface normal at P , \vec{n}_P and the incoming light direction ω .

2.4.2 Model derivation

Several modifications are made to adapt Equation 5 to US physics. Firstly, the term $O_{P \rightarrow v}$ is zero in our case as scatterers do not emit echoes.

We can then refer to the intensity sent back to the transducer from P as I_{Tr} . This term depends on the intensity $I(P)$ arriving at P , expressed following Equation 6:

$$I(P) = \int_{\Omega} I_{P' \rightarrow \omega} A_{P' \rightarrow P} d\omega \quad (6)$$

This represents the accumulation of echoes reaching P along directions ω from several points P' located on other boundaries in the scene. This is illustrated in Figure 3 where contributions from P_3 and P_2 are gathered at P . $I_{P' \rightarrow \omega}$ is the intensity leaving P' in direction ω and $A_{P' \rightarrow P}$ is the attenuation affecting the wave from P' to P along ω (denoted as $\sim \alpha^-$ in Figure 3). $I_{P' \rightarrow \omega}$ depends in turn on the intensity accumulated at P' (illustrated by incident rays at $P_1 \dots P_3$ in Figure 3) following Equation 7:

$$I_{P' \rightarrow \omega} = I(P') f_{P', \omega' \rightarrow \omega} \cos(\theta') \quad (7)$$

With θ' the angle between the incident ray ω' and $\vec{n}_{P'}$, and $f_{P', \omega' \rightarrow \omega} = R(Z_1, Z_2, \omega', \omega)^\beta T(Z_1, Z_2, \omega', \omega)^{1-\beta}$ where β is a binary variable equal to one when the ray is reflected, and zero otherwise. We randomly choose whether to reflect or refract a ray and $\beta = 1$ when $u < R(Z_1, Z_2, \theta_1, \theta_2)$, with $u \sim U(0, 1)$, otherwise $\beta = 0$. Here f is analogous to the BSDF in rendering and the corresponding loss of energy is represented at boundaries by $|| - ||$ in Figure 3.

As we now have an expression for $I(P)$, we can compute I_{Tr} . This term depends on whether or not P lies on an organ's surface. The two cases are described below:

- Similarly to Burger et al. (15), on a boundary, the intensity reflected to the transducer $I_{Tr}(P) = I_R(P)$ is written following Equation 8:

$$I_R(P) = \left(\frac{Z_2 - Z_1}{Z_2 + Z_1} \right)^2 I(P)^\tau \cos(\theta)^\gamma \quad (8)$$

- Otherwise, it is simply equal to $I(P)$ as shown in Equation 9:

$$I_{Tr}(P) = I(P) \quad (9)$$

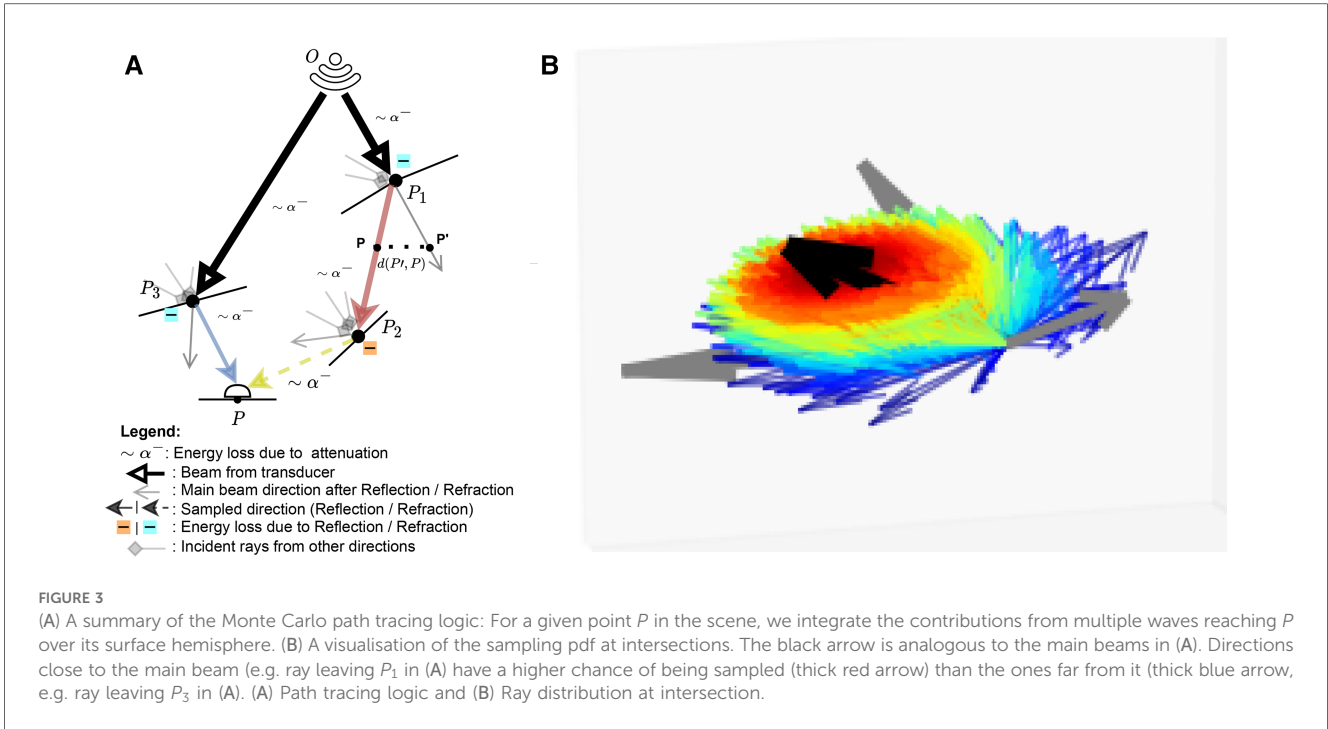
For a given point along a scanline with radial, lateral and elevation coordinates (r, l, e) , the expression for the received echo is described by Equation 10:

$$E(r, l, e) = I_{Tr}(r, l, e) \rho(r, l, e) \otimes T(r, l, e) \quad (10)$$

where $\rho(r, l, e)$ is a cosine modulated PSF and $T(r, l, e)$ the scatterer distribution. Their expressions are given in Equations 11, 12, respectively.

$$\rho(x, y, z) = \exp\left(-\frac{1}{2} \left(\frac{r^2}{\sigma_r^2} + \frac{l^2}{\sigma_l^2} + \frac{e^2}{\sigma_e^2} \right)\right) \cos(2\pi fr) \quad (11)$$

$$T(r, l, e) = \sum_{q=1}^N w_q a_q \delta(r - r_q) \quad (12)$$



N is the number of scatterers, a_q is the tissue-dependent scatterer amplitude, computed similarly to (15–17). Each scatterer is projected onto the scanline and associated with the closest radial sample r_q . Finally, w_q is used to weight the contribution of a scatterer depending on its distance to the scanline. Let's write Δ_L and $\Delta_E = e - e_S$ as the lateral and elevational distances of a scatterer to a scanline. Then w_q can be computed in two ways:

- Using an analytical beam profile, defined by a gaussian PSF with lateral and elevational variance σ_L, σ_E and expressed in Equation 13:

$$w_q = \exp\left(-\frac{1}{2}\left(\frac{\Delta_L^2}{\sigma_L^2} + \frac{\Delta_E^2}{\sigma_E^2}\right)\right) \quad (13)$$

- Using a pulse echo field generated from Field II (offline) with the desired transducer configuration. The field is sampled based on Δ_L and Δ_E and the scatterer's radial depth.

The computation of E is done using the fast implementation of the COLE algorithm from Storve et al. (31).

2.4.3 Monte Carlo path tracing

By substituting $I(P')$ in Equation 7 by its expression in Equation 6, it is easy to see the recursive nature of the integral, which makes the problem hard to solve. Hence, we resort to Monte-Carlo integration, which is a useful tool to approximate high-dimensional integrals.

This allows us to express $I(P)$ following Equation 14:

$$I(P) = \frac{1}{N} \sum_{i=1}^N \frac{I(P_i) f_{P_i, \omega' \rightarrow \omega} A_{P_i \rightarrow P} \cos(\theta_i)}{p(\omega_i)} \quad (14)$$

Unlike in Mattausch et al. (17), we explicitly weight the pdf's contribution, $p(\omega_i)$, which is the probability of generating a sample in direction ω_i . Indeed, at boundaries, rather than randomly varying the surface normal to choose a direction to trace reflected/refracted rays, we choose a random direction by sampling in a cone around the reflection/refraction directions, represented by the black arrow in Figure 3B. Indeed, when the wave hits large spherical scatterers, the reflected wavefront is a replica of the shape of the intersected area, which would take a conic shape in the case of spherical scatterers (29).

We generate random directions by sampling in spherical coordinates. More precisely, we have $\theta \sim U(0, 2\pi)$ and $\phi \sim \psi(\sigma, \mu, a, b)$ where $\psi(\sigma, \mu, a, b)$ is a truncated normal distribution. ϕ is sampled using inverse transform sampling. The joint distribution is $p(\theta, \phi) = \frac{\psi(\sigma, \mu, a, b)}{2\pi}$ and is illustrated in Figure 3B, where directions close to the reflection/refraction direction have a higher chance of being sampled (red colour) than the ones far from it (blue colour).

Finally, since we are working with solid angles, the distribution needs to be converted accordingly, with:

$$p(\omega) = \frac{p(\theta, \phi)}{\sin(\theta)} = \frac{\psi(\mu, \sigma, a, b)}{2\pi \sin(\theta)} \quad (15)$$

When propagating, the sampled ray deviates from its main beam (blue, red and yellow rays in [Figure 3A](#), yielding a reduced echo intensity. Thus, in addition to the attenuation due to propagation through tissue, the sampled rays' intensities are further reduced by weighting them with a factor w_R corresponding to the beam coherence (BC) as done in [Mattausch et al. \(17\)](#). For each point P' along the sampled ray, the amplitude is weighted by $w_R = \frac{C_0}{C_0 + d(P, P')}$, where C_0 is a user-defined constant and $d(P, P')$ is the distance between P' and its projection on the main beam P , as illustrated in [Figure 3A](#). We typically use C_0 values in the range [0, 1].

2.4.4 Traversal

Rays are sent from the virtual transducer depending on its scan geometry. The intersection with the volume is computed and from that point, we march stepwise along the ray using a hierarchical digital differential analyser (HDDA) ([32](#)). At each step, the ray is attenuated and once a boundary is reached, we randomly reflect or refract the ray. We repeat the process until a maximum number of collisions is reached.

Once the RF scanlines are computed, we apply time-gain compensation, log compression, dynamic range adjustment and scan conversion to obtain the final simulated US.

3 Experiments

In the following sections, we begin by presenting qualitative results, where we examine the impact of different parameterizations and evaluate the pipeline's ability to replicate image artefacts and patient anatomies ([Section 3.1](#)).

Subsequently, we detail our phantom experiments, which serve as a validation of essential aspects of our simulation pipeline for its role as a learning environment. We assess its capability to reproduce anatomical structures by measuring physical distances and assessing contrast, using a calibration phantom as a reference. We further investigate its aptitude in generating a fully-formed speckle pattern, as speckle is an inherent property of ultrasound images ([Section 3.2](#)).

Lastly, we showcase the utility of these simulations in training a neural network for cardiac standard view classification, a critical task for ultrasound navigation guidance ([Section 3.3](#)).

3.1 Qualitative results

[Figure 4](#) shows examples of simulated echocardiograms with various parameterizations: Firstly, the number of rays traced is critical in allowing the Monte Carlo process to converge and reveal the anatomy in the scene. Indeed, the left atrium is hardly visible in [Figure 4B](#) without MCPT, as rays reflect in deterministic directions, thus not propagating in the whole scene. Using MCPT with a greater number of rays improves the visibility of the anatomical structures as demonstrated in [Figures 4C,D](#). The beam coherence value C_0 impacts the

intensity of the rays deviating from the main beams. This is illustrated in [Figure 4E](#) where a higher C_0 reveals the aorta as deviating rays are less attenuated. Additionally, we report the Kernel Inception Distance (KID) ([33](#)) for each synthetic image, where the features are extracted from the penultimate layer of a neural network trained on real images for view classification (see [Section 3.3](#) for network and training details) and compared to features from unseen real apical two chamber (A2C) images. The reported values for each image are obtained after averaging KID values from 100 random resamplings on the real features, following the methodology in ([33](#)). For our experiments and for future use as a training environment, the preferred simulation outcome would be similar to [Figure 4E](#), as critical structures for identifying the view are clearly visible and the KID is the lowest.

[Figures 4F,G](#) show an amplification of myocardium reflections in two ways using γ and τ . The reflection intensities in [Figure 4F](#) are angle-dependent while in [Figure 4G](#) all reflections are amplified. When using an analytical profile in [Figure 4H](#), the axial distance of the scatterers along the scanline is not taken into account in w_q , meaning their amplitude is not attenuated with depth, yielding a brighter image in the far field.

[Figure 5](#) shows real acquisitions (left column) apical 5, 4, 3 chamber views (top to bottom) alongside simulations (right column). The chambers appear clearly in the images but the simulations lack fine tissue detail, as this information is lost when segmenting the input data. This is highlighted by the orange box in the four-chamber view, where the papillary muscles and valve leaflets in the real left ventricle acquisition make the ventricle's border fuzzier than in our simulation. Nevertheless, this shows the potential of the pipeline in generating any type of view.

[Figure 6](#) demonstrates post-acoustic enhancement and shadowing artefacts using a virtual sphere placed in a propagating medium. Post-acoustic enhancement is demonstrated in [Figure 6A](#), similar to artefacts caused by fluid-filled cystic structures in clinical settings. When using a highly reflective and attenuating sphere, a shadow is cast as in [Figure 6B](#). [Figures 6C,D](#) illustrate acoustic shadowing in a more complex scene, where a rib is in front of the transducer. The advantage of our pipeline lies in its ability to produce such views, which are neither routinely saved nor available in open-source ultrasound datasets.

3.2 Phantom experiments

We use a commercial calibration phantom (Multi-Purpose Multi-Tissue Ultrasound Phantom, model 040GSE, Sun Nuclear, USA) to perform the validation. Real acquisitions with a SiemensTM Healthineers ACUSON P500TM system (P4-2 phased transducer) are taken for lesion detectability comparison with the simulated images. To generate our simulations, a virtual phantom is built following the technical sheet describing the arrangement of structures in the phantom. Each type of structure is assigned a label and a segmentation volume is built. We simulate three different views, with each containing a different set of targets and perform various measurements on

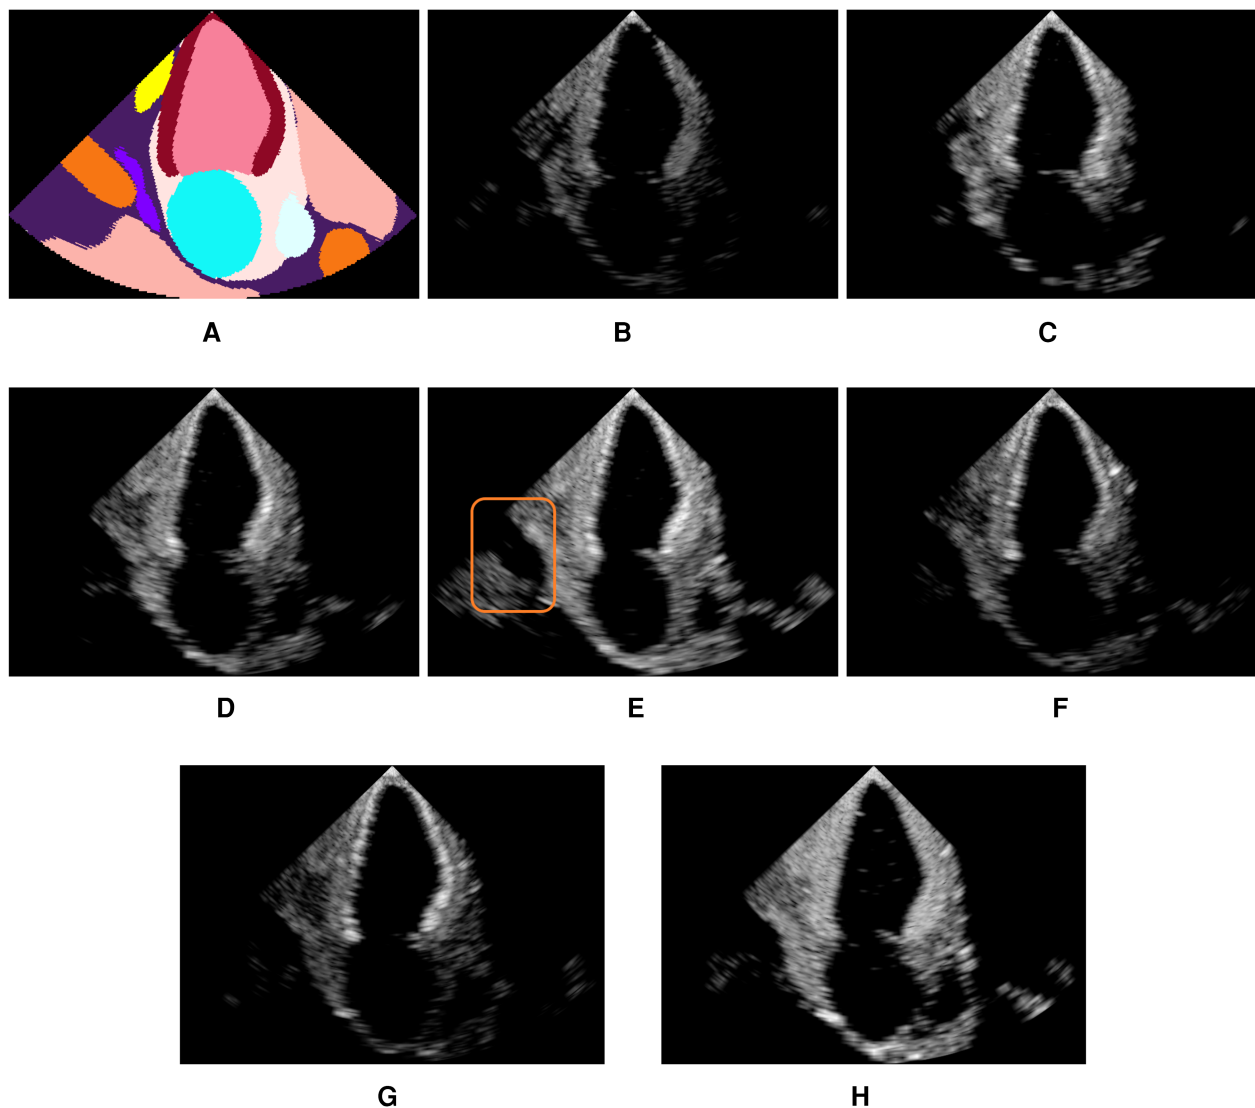


FIGURE 4

Illustration of the influence of the MCPT, beam coherence C_0 value, scatterer weighting strategy, τ and γ terms. All simulations use MCPT, 2,500 rays, a pulse-echo field from Field II with a focus at 60mm, $C_0 = 0.1$ and the myocardium properties are $\tau = 2.0$ and $\gamma = 0.1$ unless stated otherwise. The values in parentheses indicate the Kernel Inception Distance (KID) for each image, computed w.r.t real A2C images from the test dataset. Features for KID computation were extracted from a network trained on real images. (A) is an input segmentation map for an A2C view, where the orange label is associated with the aorta. In (E), the orange box denotes the aorta, showing the simulations reproduce patient-specific anatomy with fidelity. (A) Segmentation. (B) No MCPT (1.494). (C) 500 rays (1.490). (D) 2,500 rays (1.493). (E) $C_0 = 0.2$ (1.488). (F) $\gamma = -1.8$ (1.493). (G) $\tau = 2.8$ (1.492) and (H) $\sigma_L = 1.5^{-3}$ (1.492).

each synthesized view. As we perform a comparison of lesion detectability in simulated and real images, we set the image pixel spacing of our simulations to the same value as the real acquisitions, i.e. at 0.23 mm. All simulations are done using a Desktop computer equipped with an NVIDIA Quadro K5000 GPU.

3.2.1 Experiment parameters

The transducer, simulation and post-processing parameters for phantom and view classification experiments are listed in Table 3. For the phantom experiment, the transducer is parameterized similarly to the real one following the parameters

listed in Table 2. The parameters for the truncated normal distribution ϕ are $\mu = 0$, $\sigma = \frac{\pi}{4}$, $a = 0$, $b = \frac{\pi}{2}$.

3.2.2 Distance measurements

We sampled pixels along a 1-D line going through nylon targets. The coordinates of the line were automatically computed given the technical phantom sheet. A 1-D signal was extracted from this line and peaks (corresponding to the centre of nylon wires) were identified. Knowing the virtual transducer's position as well as the peaks' location along the line allowed us to compute a Target Registration Error (TRE) between the expected and simulated nylon wire positions.

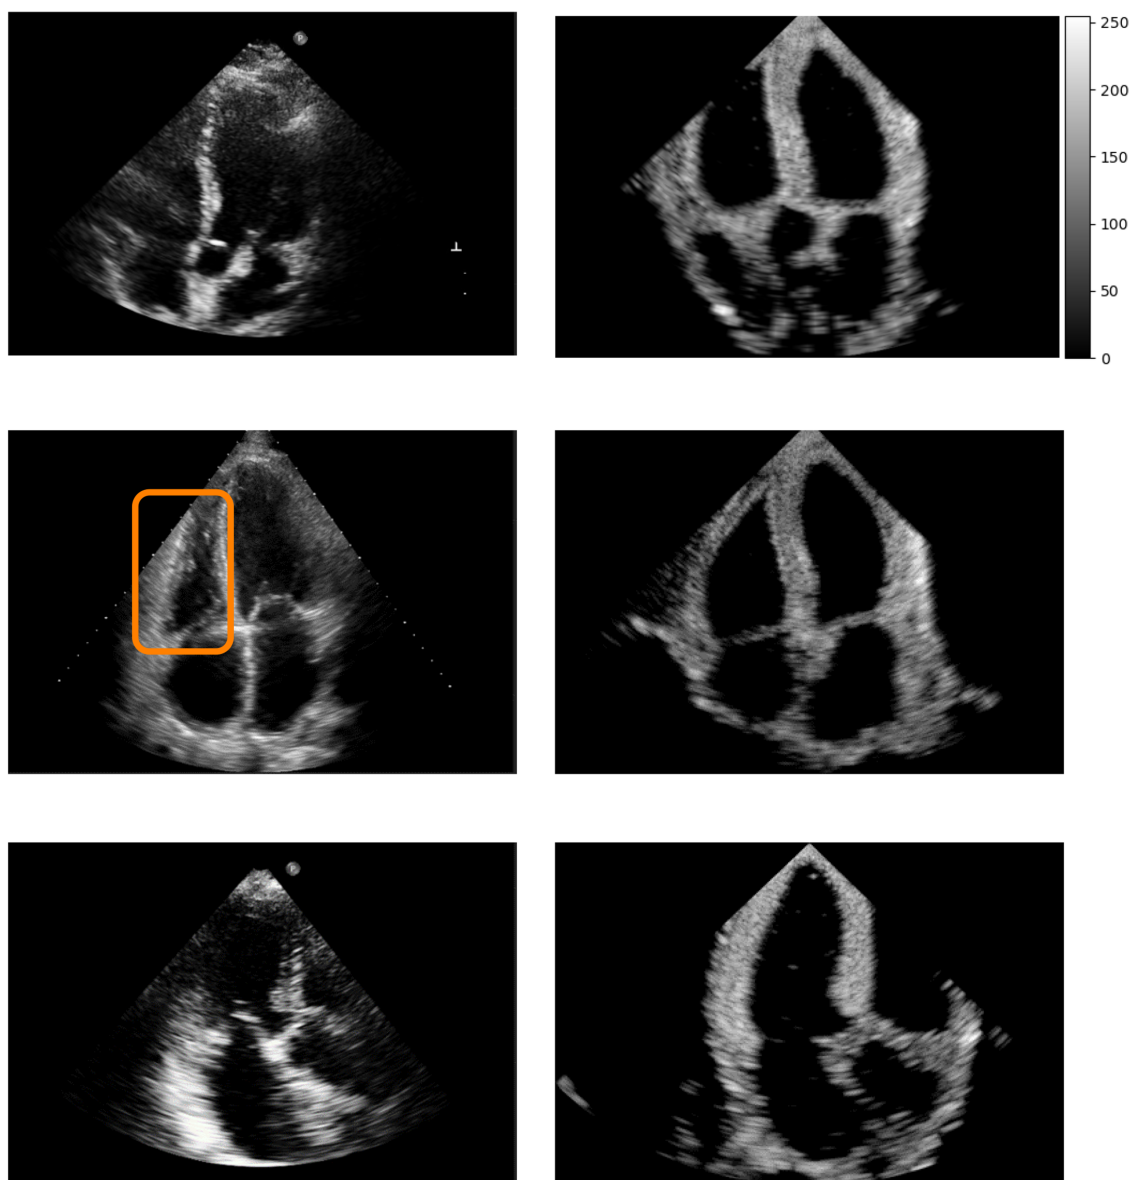


FIGURE 5
Real (left column) and simulated (right column) Apical 5, 4, 3 chambers views (top to bottom, not paired). The orange box denotes papillary muscles and fine cardiac structures which are not captured by the simulations, making the ventricles' borders sharper in the synthetic images.

Examples of targets used in this experiment are shown by the orange box in [Figure 7D](#). A detailed summary of the error per view and per target group is given in [Table 4](#). An error of 0.20 ± 0.32 mm was reported when measuring the TRE from 60 targets.

A pattern emerges from [Table 4](#), where the error increases with depth (Horizontal Distance Groups 1 and 2). This is due to beam divergence in the far field, which decreases the lateral resolution. This agrees with experimental measurements.

3.2.3 Lesion detectability and contrast

Having an accurate contrast between background and surrounding structures is critical in ultrasound as it allows users

to discriminate between tissues. Thus, we investigate the ability of our pipeline to simulate structures of various contrast. To this end, we compare anechoic and hyperechoic lesions from our virtual phantom to the same lesions from real acquisitions.

In addition to classical metrics such as Contrast to Noise Ratio (CNR) and contrast, we reported the generalized Contrast-to-Noise Ratio (gCNR) (34), a metric robust to dynamic range alterations and with a simple interpretation. Since our post-processing pipeline differs from the P500's as it is a commercial system, this metric would provide a way to compare the lesion detectability independently of post-processing differences.

We computed gCNR, CNR and contrast between lesions and background using two views. The background patch size

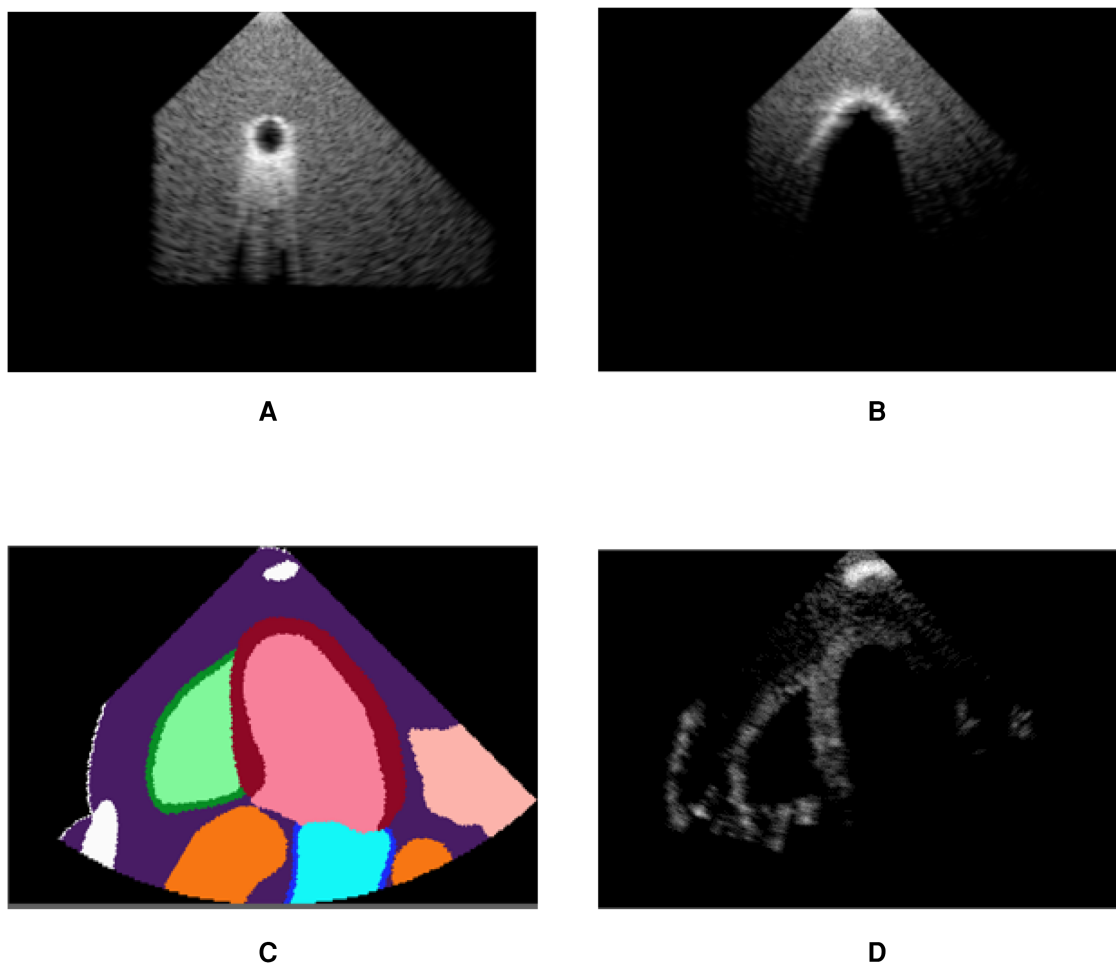


FIGURE 6

Our pipeline is able to recreate some artefacts such as (A) post-acoustic enhancement and (B) shadowing. Spheres filled with fluid (A) and with high attenuation (B) were used to recreate the artefacts. (C) shows segmentation labels of a scene with a rib in front of the transducer (white label) and (D) is the corresponding simulated image, demonstrating acoustic shadowing. (A) Post-acoustic enhancement. (B) Acoustic shadowing. (C) Segmentation map and (D) Rib shadowing.

was calculated to have a sample size similar to the lesion patch. Real and simulated acquisitions, as well as histograms of the lesions and background distributions, are illustrated in Figure 7.

A summary of the scores between real and simulated images is indicated in Table 5. Overall, gCNR, CNR and contrast values between real and simulated values are close, suggesting our pipeline reproduces lesions with fidelity. Contrast values for the second and third anechoic lesions differ as in the real acquisition, the far field is more hypoechoic compared to the focus area in the centre of the image.

3.2.4 Speckle pattern analysis

In this section, we analyze the capability of our simulator to generate a fully-developed speckle pattern. To this end, we conduct a comparative analysis similar to Gao et al. (35), where random scatterers at a density of 600 mm^{-2} and fixed amplitude of 1 are distributed in a $40 \times 40 \text{ mm}^2$ volume. It is known for

such an experiment that the envelope detected signal follows a Rayleigh distribution and its signal-to-noise ratio (SNR) reaches a value of 1.91 (36). The experiment is repeated 10 times to take into account its stochastic nature. Here, we use an analytical beam profile with $\Delta_E, \Delta_L = 2.0 \text{ mm}$. For each run, the SNR is computed and the sum-of-squared errors (SSE) w.r.t a fitted Rayleigh distribution is calculated. An example histogram and fitted distribution from a run is shown in Figure 8.

We obtain a mean SSE of $1.81e^{-5}$ and SNR of 1.89 ± 0.01 , which is in the ranges reported in the literature (14, 16, 35). This suggests that our pipeline is able to create a fully developed speckle pattern.

3.3 View classification

Our last experiment assesses the usability of simulated images to train neural networks for view classification. This task is intrinsic

TABLE 3 Parameters used in the experiments.

Property	Value in experiment	
	Phantom	View classification
Transducer parameters		
Sampling frequency (MHz)	50	
Center frequency (MHz)	3.6	3.6
Field II pulse echo field focus (cm)	5	6.5
Analytical profile std (Δ_L , Δ_R) (mm)	N/A	1.0
Simulation parameters		
Max num. collisions	7	10
Beam coherence C_0	0.1	[0.01, 0.05, 0.075, 0.1]
Num. rays per element	1,000	[1,000, 2,000, 3,000, 5,000]
Post-processing parameters		
Dynamic Range (dB)	75	[65, 75, 85, 95]
Time-Gain compensation (dB/cm)	1.5	
Reject threshold (dB)	40	[35, 45, 50, 60]
τ	1.0	2.0
γ	0.7	0.1

For the view classification experiment, brackets indicate the range of values sampled.

to navigation as a network must be capable of identifying when a target view has been reached. We train networks to classify real apical views (A2C, A3C, A4C, A5C) and investigate the impact of fine-tuning with real data on the networks' performance, as networks trained in the simulation environment would likely be fine-tuned to adapt to real scenarios. Synthetic and real dataset generation are described in Sections 3.3.1 and 3.3.2. Furthermore, we conduct an ablation study where we evaluate the impact of parameters we believe impact the image quality the most, namely: the use of MCPT and the weighting method when projecting scatterers, i.e. with an analytical function or using a

TABLE 4 Target Registration Error (TRE) between expected and simulated wire positions (mean \pm std).

Structure	TRE per view in mm		
	View 1	View 2	View 3
Target groups			
Vertical distance	0.03 \pm 0.03	0.09 \pm 0.07	0.12 \pm 0.09
Horizontal distance 1 (near-field)	0.05 \pm 0.04	0.16 \pm 0.14	0.28 \pm 0.05
Horizontal distance 2 (far-field)	0.34 \pm 0.33	0.25 \pm 0.59	0.33 \pm 0.37

pulse echo field from Field II. The experimental setup is detailed in Section 3.3.3, followed by the results in Section 3.3.4.

3.3.1 Simulated TTE dataset

Chest CTs and Cardiac CTs from 1,019 patients from the FUMPE (37) and The Cancer Imaging Archive (38) [LIDC-IDRI (39)] datasets were used to generate simulated images. The volumes were automatically segmented using (40) and pre-processed according to the pipeline described in Figure 2 and several landmarks were automatically obtained (apex, the centre of the heart chambers...) and used to find the appropriate transducer orientations and positions to acquire the standard views.

For each view, we generate multiple synthetic samples by varying simulation parameters as described in Table 3. We generated more synthetic samples for the A5C view to compensate for the low number of datasets where we were able to automatically obtain a suitable view. The final dataset distribution is 30%, 30%, 30% and 10% resp. for the A2C, A3C, A4C and A5C classes.

All the samples from the simulated dataset are used for training. The average simulation time per image was 300

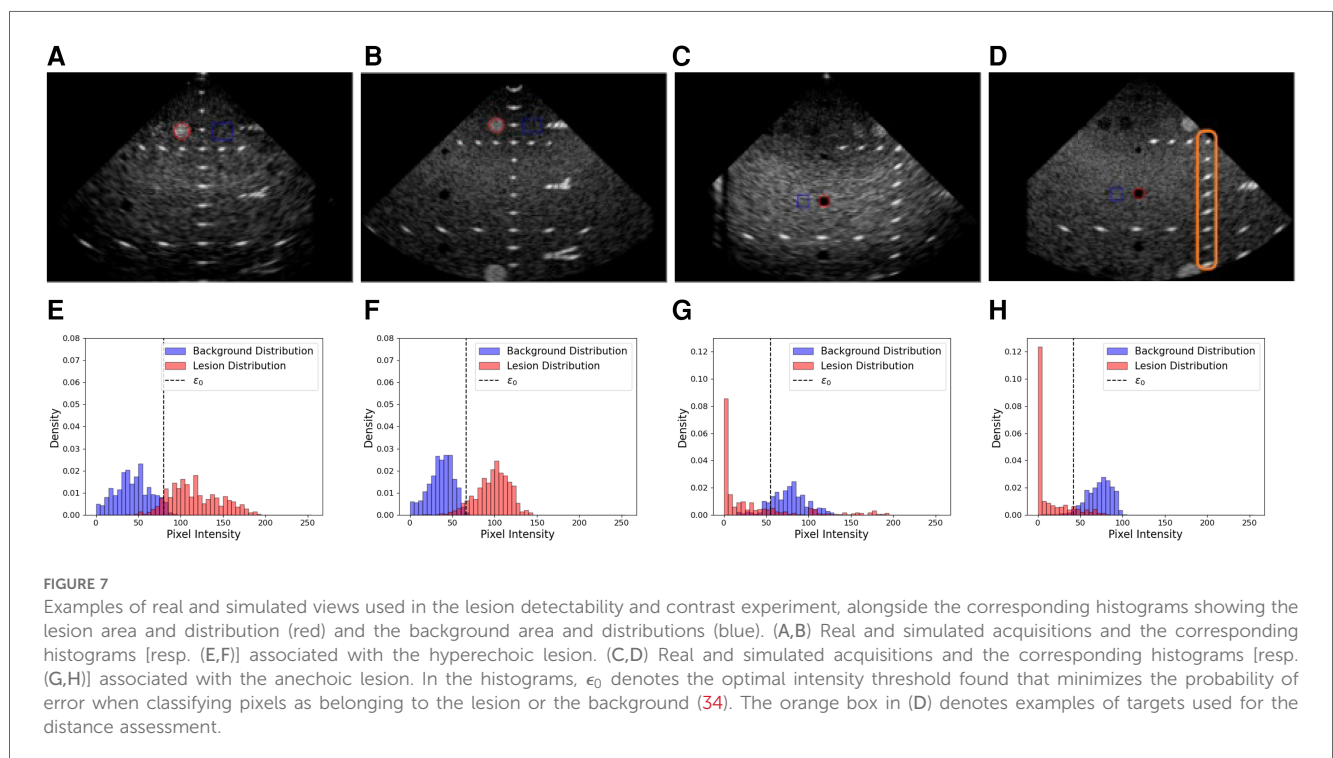
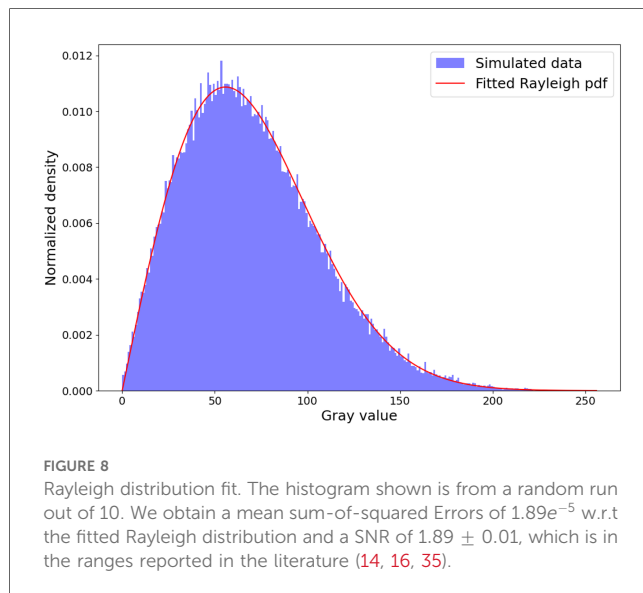


TABLE 5 gCNR, CNR and contrast (in dB), values from lesions in real and simulated US acquisitions (mean \pm std).

Lesion	Metric	Real	Sim
Hyperechoic (+6 dB)	gCNR	0.19	0.22 \pm 0.08
	CNR	0.27	0.29 \pm 0.1
	Contrast	4.72	4.88 \pm 1.82
Hyperechoic (+15 dB)	gCNR	0.88	0.89 \pm 0.02
	CNR	0.74	0.77 \pm 0.03
	Contrast	16.07	16.97 \pm 2.30
Anechoic 1	gCNR	0.80	0.82 \pm 0.03
	CNR	0.64	0.70 \pm 0.05
	Contrast	-14.17	-13.38 \pm 1.89
Anechoic 2	gCNR	0.87	0.78 \pm 0.04
	CNR	0.72	0.70 \pm 0.04
	Contrast	-16.61	-13.38 \pm 1.34
Anechoic 3	gCNR	0.71	0.71 \pm 0.06
	CNR	0.66	0.69 \pm 0.06
	Contrast	-19.41	-13.69 \pm 2.43

Simulations were generated 10 times to take in account the stochasticity of the MCPT and the scatterers' generation in $T(x)$.



milliseconds. This number includes only the simulation step (i.e. Figure 1C).

Finally, to conduct the ablation study, 3 different simulated datasets are created.

- sim NO MCPT, where MCPT was disabled. Thus all samples are generated with deterministic raytracing.
- sim + MCPT, where MCPT was enabled and an analytical beam profile used.
- sim + MCPT + FIELD, where MCPT was enabled and a pulse echo field from Field II was used to weight the scatterers' contributions.

3.3.2 Real TTE dataset

We used real US acquisitions to train and test the view classification network. The video sequences came from Siemens

and Philips systems. During training, we sample randomly one frame from a given sequence and add it to the training batch. The real training dataset is also imbalanced, where the sample distribution in each fold for A2C, A3C, A4C and A5C classes is around 21%, 18%, 51% and 10%.

3.3.3 Evaluation methodology

For this experiment, we used a Convolutional Neural Network (CNN) with a DenseNet architecture (41) to classify views. The network architecture is kept fixed for all experiments. Random weighted sampling is used to fight class imbalance. We divide the real dataset into 5 folds for cross-validation but always use the same synthetic dataset for pre-training.

In each fold, we create subsets d_r of the real training dataset D_{real} with varying amounts of real data. For each d_r , we train four networks: One network on d_r only, to establish a baseline and then we pre-train 3 other networks on each one of the simulated datasets and fine-tune them on d_r . Validation and testing are always done on the same real datasets, independently of d_r 's size.

When pre-training, we employ the following data augmentations on the synthetic samples to match the variations observed in the real dataset: Cropping/zooming (e.g. to mimic real sequences where there's a zoom on a chamber or a valve), Gaussian smoothing, brightness and contrast jittering (to replicate varying texture qualities), fan angle variation (for real sequences where the fan angle is changed by the operator). No augmentations are applied to the real data.

When evaluating, for each video sequence, we perform a majority vote on the network's predictions on each frame to determine which label to assign to the sequence.

3.3.4 Results

We report averaged F1-score and accuracy for all the classes in Figures 9A,B and F1-score for the A5C and A4C classes in Figures 9C,D.

Figure 9 suggests pre-trained networks achieve a performance level comparable to networks trained on all real datasets when fine-tuned with at least half of the real data.

Fine-tuned networks show significant improvements over their counterparts trained on real data (when $d_r < 800$). This trend is accentuated for the A5C class, which is the most under-represented in the dataset. Using simulated data for pre-training still benefits the dominant A4C class, as shown in Figure 9D. Results for networks trained on simulated data only are not reported as they overfitted easily and performed poorly on the real test dataset.

Confusion matrices for $d_r = 450$ are reported in Figure 10 for the baseline trained on real data only (Figure 10A) and the network pre-trained on sim + MCPT (Figure 10B). There is a noticeable improvement in the results, highlighted by a reduction in confusion between the A5C and A4C classes.

Finally, no statistically significant differences were found when comparing the results of the networks pre-trained on sim + MCPT and sim + MCPT + FIELD ($p > 0.05$). This suggests the choice of the weighting method for scatterers has little influence on neural

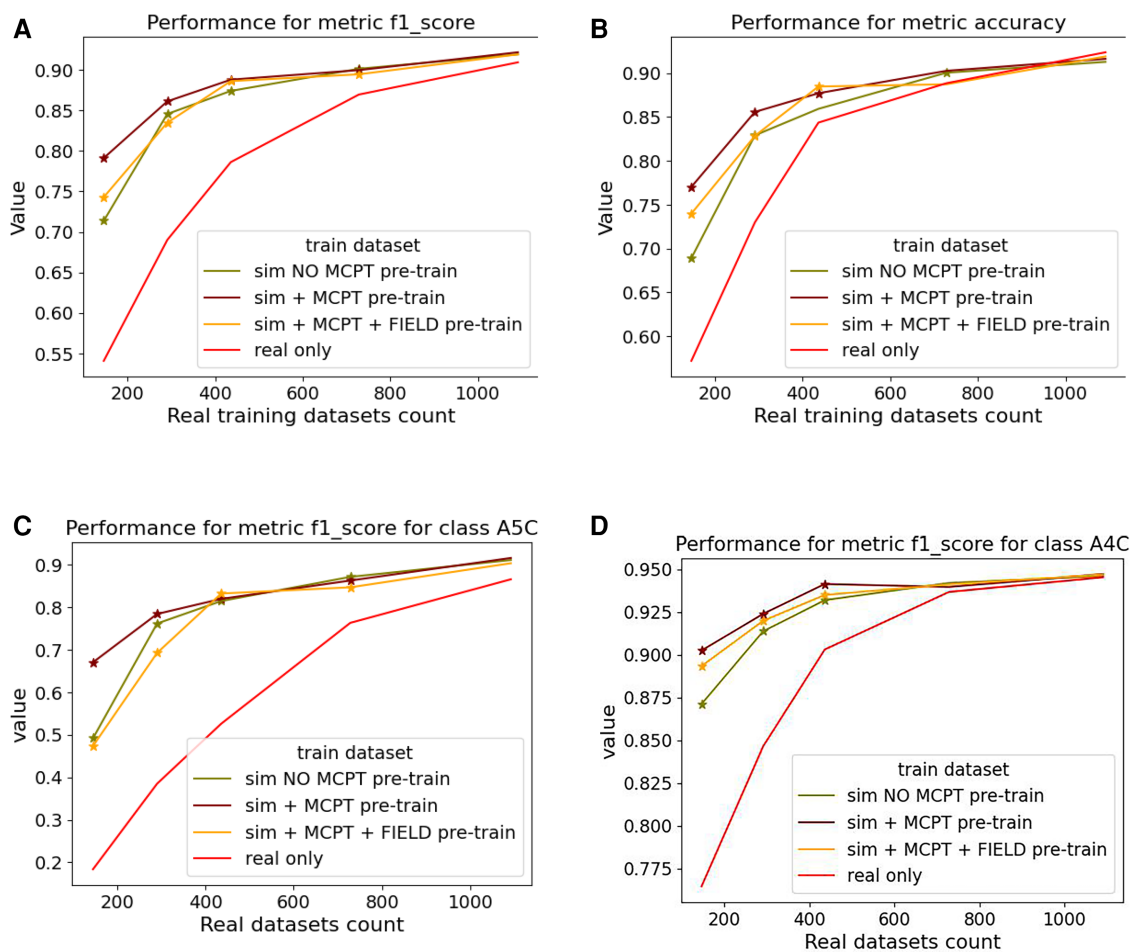


FIGURE 9

Results of the view classification ablation study averaged over 5 folds. Networks pre-trained with simulations and then fine-tuned on real samples were compared to networks trained on real data only. The x-axis indicates the size of the subset of real data d_r . (A,B) report the F1-score and accuracy over the 4 classes while (C,D) report the metrics for the (most-represented) A4C and (under-represented) A5C classes. For a given d_r , a star is displayed on a graph if the p -value from a right-tailed Wilcoxon signed rank-test is < 0.05 .

network training on this task. Results were statistically different between sim NO MCPT and sim + MCPT when $d_r < 450$ and only different between sim NO MCPT and sim + MCPT + FIELD when $d_r < 150$.

4 Discussion

In this section, we first discuss experimental results from the view classification experiment in Section 4.1. We then address the limitations of our proposed simulation pipeline in Section 4.2 and finish by expanding on potential applications of the pipeline and future work in Section 4.3.

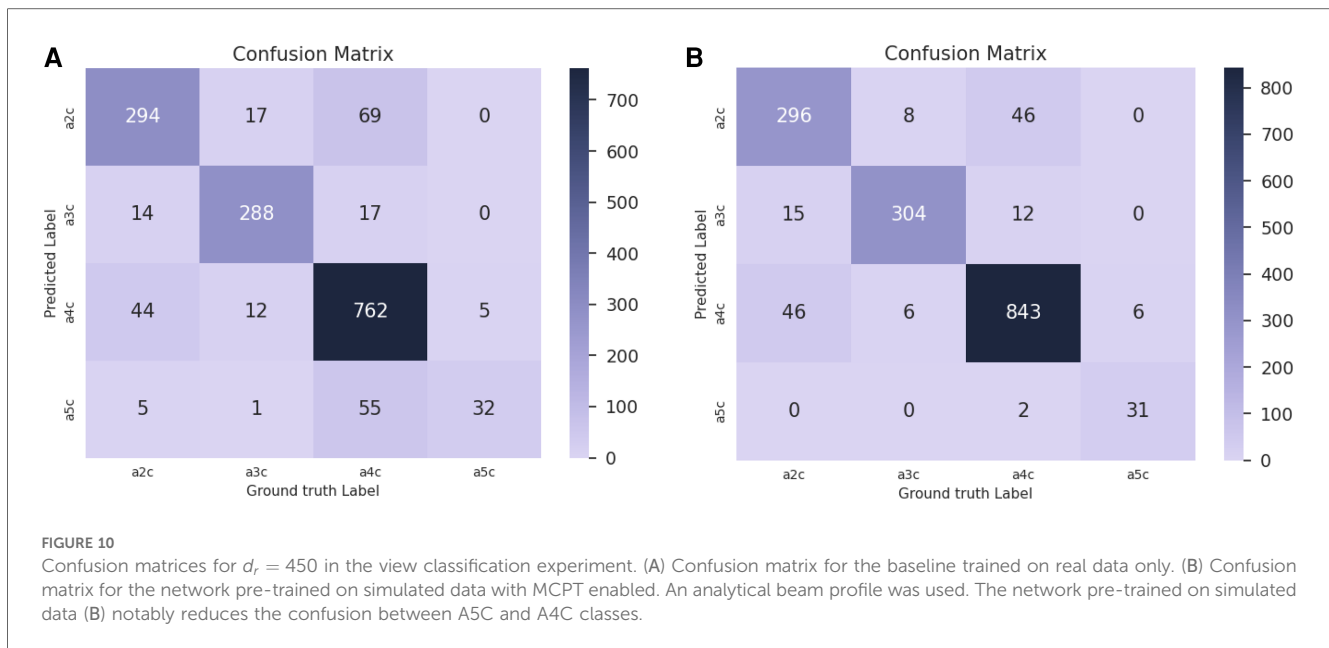
4.1 View classification

In Figure 9, pre-trained networks show improved performance compared to the ones trained on real data only, meaning the simulations can be used to generate data when large datasets are

not readily available or to target a sub-population which is less prevalent. This suggests our pipeline could be used to generate data for other tasks, given some improvements are made to circumvent the limitations caused by using segmentations. We expand on this in Section 4.2.

Moreover, networks pre-trained without MCPT achieved in some cases performances similar to their counterparts trained with MCPT. While MCPT allows for a better visibility of the anatomical structures as demonstrated in Figure 4, the discriminating features between views (i.e. heart chambers) are still present in the images without using MCPT. This would explain why the networks can still learn from such images. However, we believe using MCPT might be more critical in applications where all structures need to be clearly observable, such as image segmentation.

We limited the view classification experiment to four views as apical views were the only ones we could obtain robustly in an automatic way. Even then, we were not always successful in obtaining correct transducer orientations for each apical view in every patient dataset, especially for the A5C view. Indeed, view planes for each patient are obtained by finding landmarks using



segmentations and morphological operations and then fitting a plane. Our automated method failed to consistently find a plane where the aorta and the four chambers were visible in the simulations. This is related to the fact that we obtain our segmentations from CT data, where the patients are lying supine, and it is known that finding A5C views when patients are in the supine position is complicated in clinical settings as the imaging plane is suboptimal (42). This explains the synthetic training data distribution in the view classification experiment. However, using an algorithm capable of navigating between views (which is what we intend to develop using the simulator), we could potentially generate datasets with a greater number of standard views.

Finally, in Figure 10, there is a confusion between A2C and A4C classes. Our data is annotated such that all frames in a video sequence have the same label. However, there are multiple A4C sequences in which some frames resemble A2C views (due to suboptimal probe positioning or cardiac phase) but are labelled as A4C, which introduces confusion for the network during training.

4.2 Proposed simulation pipeline

While this pipeline allows for the fast simulation of arbitrary anatomies from a large number of patient datasets, it presents limitations:

(1) Similarly to other raytracing methods, we cannot simulate non-linear propagation. This prevents us from using techniques such as tissue harmonic imaging. Furthermore, we cannot reproduce reverberations. These could be simulated by summing the ray contributions temporally (i.e. by keeping track of the distance travelled by a ray) rather than spatially. However, this requires a careful weighting of the contribution of the randomly sampled rays with the beam coherence, so as to not yield incorrect results.

(2) As seen in Figure 5, the border with the blood pool is sharp and the inhomogeneities of tissues are not represented in the simulations. This is due, respectively, to smaller cardiac structures (e.g. papillary muscles, trabeculae ...) which are difficult to annotate and segment, and to the assumption of homogeneity within the tissue (i.e. all scatterers' intensities in a given medium follow the same distribution) since the intensity variation between pixels is lost with the segmentations. Additionally, the segmentation algorithm can also produce inaccurate labels. The impact of these factors depends on the downstream tasks.

A quantitative evaluation of segmentation errors could be made using a CT with accurate segmentation labels and then altering the labels with geometric and morphological operations to assess the impact of inaccuracies on the outcomes. Furthermore, given pairs of registered CT and real ultrasound images, one could assess the ability of the pipeline in simulating pathological cases, as in Figures 6A,B.

The impact of the tissue homogeneity assumption was illustrated in our attempt to train networks solely on simulated data for the view classification experiment, but the performance was poor. We noticed the network quickly overfitted the data. While the range of anatomies simulated is wide (+1,000 patients), the lack of fine-tissue detail seems to limit the diversity of generated samples. We believe a potential solution to this challenge would be a combination of our pipeline with generative models, to improve the realism and quality of simulations. This could enable the generation of large and realistic ultrasound datasets, with readily available anatomical labels.

(3) While we do not address the topic of cardiac motion in this manuscript, it is possible to generate such sequences with our pipeline, given input volumes for each timestep of the cardiac cycle.

(4) We recognize that the pressure applied by sonographers on the patient's chest during TTE examinations can impact the

image quality. We plan to address this in future work by incorporating a volumetric deformation model over the anatomical volume. Nevertheless, we note that the proposed framework would still be sufficient for training navigation algorithms for transesophageal imaging, where the impact on images of such anatomical shape deformations due to the ultrasound probe would be significantly smaller.

4.3 Applications and future work

We aim to use our pipeline as a simulator to train navigation algorithms, similar to Li et al. (5). While the motivation behind the development of our pipeline is autonomous navigation, its capabilities could enable numerous downstream applications. Large dataset generation from segmentations could allow for the training of neural networks for tasks such as view classification, image segmentation or automated anatomical measurements.

In addition to the proposed use for automated acquisition, the method could be used for training or guidance of a semi-trained or novice ultrasound operator. Typically, guidance methods use 2D images from a pre-acquired 3D dataset. However, a simulation method would enable larger adjustments to the probe position.

While we focus on cardiac TTE imaging in this paper, other organs or modalities such as Transoesophageal Echocardiography (TEE) or Intracardiac Echocardiography (ICE), in 2D or 3D, could be simulated as a result of the built-in flexibility of our pipeline. Our future work will investigate both the use of the simulation pipeline as an environment to train deep reinforcement learning agents for autonomous navigation and the use of generative networks to improve the realism and train networks for several downstream tasks.

5 Conclusion

We have presented an ultrasound simulation pipeline capable of processing numerous patient datasets and generating patient-specific images in under half a second. In the first experiment, we assessed several properties of the simulated images (distances, contrast, speckle statistics) using a virtual calibration phantom. The geometry of our simulations is accurate, the contrast of different tissues is reproduced with fidelity and we are able to generate a fully developed speckle pattern.

We then synthesized cardiac views from more than 1,000 real patient CT datasets and pre-trained networks using simulated datasets. The pre-trained networks required around half the real data for fine-tuning to reach a performance level comparable to networks trained with all the real samples, demonstrating the usefulness of simulations when large real datasets are not available.

The main limitation lies in the use of segmentations, unable to capture smaller cardiac structures or intensity variations between neighbouring pixels. Using a generative neural network to augment the simulations is a potential workaround. Such a pipeline enables a large number of downstream applications,

ranging from data generation for neural network training (segmentation, classification, navigation) to sonographer training.

Data availability statement

Publicly available datasets were analyzed in this study. This data can be found here: https://imaging.cancer.gov/informatics/cancer_imaging_archive.htm, <https://www.kaggle.com/datasets/andrewmvd/pulmonary-embolism-in-ct-images>.

Author contributions

AAA: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Methodology, Investigation, Formal Analysis, Data curation, Conceptualization; LP: Writing – review & editing, Writing – original draft, Validation, Supervision, Methodology; PD: Writing – review & editing, Supervision, Methodology, Formal Analysis; PK: Writing – review & editing, Data curation; KP: Writing – review & editing, Methodology; RJH: Writing – review & editing, Supervision, Methodology; VS: Writing – review & editing, Writing – original draft, Supervision, Methodology, Investigation, Formal Analysis; RL: Writing – original draft, Supervision, Methodology; Y-HK: Writing – review & editing, Writing – original draft, Methodology; FCG: Writing – review & editing, Writing – original draft, Methodology, Investigation; TM: Writing – original draft, Supervision, Methodology; RR: Writing – review & editing, Supervision; AY: Writing – review & editing, Writing – original draft, Validation, Methodology; KR: Writing – review & editing, Writing – original draft, Validation, Supervision, Resources, Methodology.

Funding

The author(s) declare financial support was received for the research, authorship, and/or publication of this article.

This research was funded in part, by the Wellcome Trust under Grant WT203148/Z/16/Z and in part by Siemens Healthineers.

Acknowledgments

The authors acknowledge the National Cancer Institute and the Foundation for the National Institutes of Health, and their critical role in the creation of the free publicly available LIDC/IDRI Database used in this study.

Conflict of interest

PK, KP, VS, Y-HK and FCG are employed by Siemens Healthineers. RL and TM were employed by Siemens Healthineers.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated

organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Pinto AV, Pinto F, Faggian A, Rubini G, Caranci F, Macarini L, et al. Sources of error in emergency ultrasonography. *Crit Ultrasound J.* (2013) 5:S1. doi: 10.1186/2036-7902-5-S1-S1
- Haxthausen FV, Böttger S, Wulff D, Hagenah J, García-Vázquez V, Ipsen S. Medical robotics for ultrasound imaging: Current systems and future trends. *Curr Robot Rep.* (2021) 2:55–71. doi: 10.1007/s43154-020-00037-y
- Huang Y, Xiao W, Wang C, Liu H, Huang RP, Sun Z. Towards fully autonomous ultrasound scanning robot with imitation learning based on clinical protocols. *IEEE Robot Autom Lett.* (2021) 6:3671–8. doi: 10.1109/LRA.2021.3064283
- Hase H, Azampour MF, Tirindelli M, Paschali M, Simson W, Fatemizadeh E, et al. Ultrasound-guided robotic navigation with deep reinforcement learning. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2020). p. 5534–41.
- Li K, Li A, Xu Y, Xiong H, Meng MQH. RL-tee: autonomous probe guidance for transesophageal echocardiography based on attention-augmented deep reinforcement learning. *IEEE Trans Autom Sci Eng.* (2023) 21(2):1526–38. doi: 10.1109/TASE.2023.3246089
- Gilbert A, Marciniak M, Rodero C, Lamata P, Samset E, Mcleod K. Generating synthetic labeled data from existing anatomical models: An example with echocardiography segmentation. *IEEE Trans Med Imaging.* (2021) 40:2783–94. doi: 10.1109/TMI.2021.3051806
- Tiago C, Snare SR, Sprem J, Mcleod K. A domain translation framework with an adversarial denoising diffusion model to generate synthetic datasets of echocardiography images. *IEEE Access.* (2023) 11:17594–602. doi: 10.1109/ACCESS.2023.3246762
- Parker SG, Bigler J, Dietrich A, Friedrich H, Hoberock J, Luebke DP, et al. Optix: a general purpose ray tracing engine. *ACM SIGGRAPH 2010 papers* (2010).
- Jensen JA. Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers. *IEEE Trans Ultrason Ferroelectr Freq Control.* (1992) 39:262–7. doi: 10.1109/18.139123
- Arendt J. *Paper Presented at the 10th Nordic-Baltic Conference on Biomedical Imaging: Field: A Program for Simulating Ultrasound Systems* (1996).
- Pinton GF, Dahl J, Rosenzweig SJ, Trahey GE. A heterogeneous nonlinear attenuating full-wave model of ultrasound. *IEEE Trans Ultrason Ferroelectr Freq Control.* (2009) 56:474–88. doi: 10.1109/TUFFC.2009.1066
- Karamalis A, Wein W, Navab N. Fast ultrasound image simulation using the westervelt equation. *Med Image Comput Assist Interv.* (2010) 13 Pt 1:243–50. doi: 10.1007/978-3-642-15705-9_30
- Treeby BE, Cox BT. k-wave: Matlab toolbox for the simulation and reconstruction of photoacoustic wave fields. *J Biomed Opt.* (2010) 15:021314. doi: 10.1117/1.3360308
- Gao H, Choi HF, Claus P, Boonen S, Jaecques SVN, van Lenthe GH, et al. A fast convolution-based methodology to simulate 2-Dd/3-D cardiac ultrasound images. *IEEE Trans Ultrason Ferroelectr Freq Control.* (2009) 56:404–9. doi: 10.1109/TUFFC.2009.1051
- Bürger B, Bettinghausen S, Rädle M, Hesser JW. Real-time GPU-based ultrasound simulation using deformable mesh models. *IEEE Trans Med Imaging.* (2013) 32:609–18. doi: 10.1109/TMI.2012.2234474
- Salehi M, Ahmadi SA, Prevost R, Navab N, Wein W. Patient-specific 3D ultrasound simulation based on convolutional ray-tracing and appearance optimization. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention* (2015).
- Mattausch O, Makhinya M, Goksel O. Realistic ultrasound simulation of complex surface models using interactive monte-carlo path tracing. *Comput Graph Forum.* (2018) 37:202–13. doi: 10.1111/cgf.13260
- Alessandrini M, Craene MD, Bernard O, Giffard-Roisin S, Allain P, Waechter-Stehle I, et al. A pipeline for the generation of realistic 3D synthetic echocardiographic sequences: methodology and open-access database. *IEEE Trans Med Imaging.* (2015) 34:1436–51. doi: 10.1109/TMI.2015.2396632
- Alessandrini M, Chakraborty B, Heyde B, Bernard O, Craene MD, Sermesant M, et al. Realistic vendor-specific synthetic ultrasound data for quality assurance of 2-D speckle tracking echocardiography: simulation pipeline and open access database. *IEEE Trans Ultrason Ferroelectr Freq Control.* (2018) 65:411–22. doi: 10.1109/TUFFC.2017.2786300
- Stojanovski D, Hermida U, Lamata P, Beqiri A, Gómez A. Echo from noise: synthetic ultrasound image generation using diffusion models for real image segmentation. *ASMUS@MICCAI* (2023).
- Hu Y, Gibson E, Lee LL, Xie W, Barratt DC, Vercauteren TKM, et al. Freehand ultrasound image simulation with spatially-conditioned generative adversarial networks. *arXiv arXiv:abs/1707.05392* (2017).
- Vitale S, Orlando JJ, Iarusi E, Larrabide I. Improving realism in patient-specific abdominal ultrasound simulation using cyclegans. *Int J Comput Assist Radiol Surg.* (2020) 15:183–92. doi: 10.1007/s11548-019-02046-5
- Zhang L, Portenier T, Goksel O. Unpaired translation from semantic label maps to images by leveraging domain-specific simulations. *arXiv arXiv:abs/2302.10698* (2023).
- Tomar D, Zhang L, Portenier T, Goksel O. Content-preserving unpaired translation from simulated to realistic ultrasound images. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention* (2021)
- Jacob AJ, Abdelkarim O, Zook S, Cocker MS, Gupta P, Giraldo JCR, et al. AI-based cardiac chamber analysis from non-contrast, gated cardiac CT. *J Am Coll Cardiol.* (2022) 79(9):1250. doi: 10.1016/j.jacc.2022.02.024-0
- Hu Y, Nguyen H, Smith C, Chen T, Byrne M, Archibald-Heeren B, et al. Clinical assessment of a novel machine-learning automated contouring tool for radiotherapy planning. *J Appl Clin Med Phys.* (2023) 24. doi: 10.1002/acm2.13949
- Museth K. VDB: High-resolution sparse volumes with dynamic topology. *ACM Trans Graph.* (2013) 32:27:1–27:22. doi: 10.1145/2487228.2487235
- Museth K. NanoVDB: A GPU-friendly and portable VDB data structure for real-time rendering and simulation. *ACM SIGGRAPH 2021 Talks* (2021).
- Szabo TL. *Diagnostic ultrasound imaging: inside out* (second edition). *Ultras Med Biol.* (2015) 41:622. doi: 10.1016/C2011-0-07261-7
- Kajiya JT. The rendering equation. In: *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques* (1986).
- Storve S, Torp H. Fast simulation of dynamic ultrasound images using the GPU. *IEEE Trans Ultrason Ferroelectr Freq Control.* (2017) 64:1465–77. doi: 10.1109/TUFFC.2017.2731944
- Museth K. Hierarchical digital differential analyzer for efficient ray-marching in openVDB. *ACM SIGGRAPH 2014 Talks* (2014).
- Binkowski M, Sutherland DJ, Arbel M, Gretton A. Demystifying mmd gans. In: *International Conference on Learning Representations* (2018).
- Rodriguez-Molares A, Rindal OMH, D'hooge J, Måsøy SE, Austeng A, Bell MAL, et al. The generalized contrast-to-noise ratio: a formal definition for lesion detectability. *IEEE Trans Ultrason Ferroelectr Freq Control.* (2019) 67:745–59. doi: 10.1109/TUFFC.2019.2956855
- Gao H, D'hooge J, Hergum T, Torp H. Comparison of the performance of different tools for fast simulation of ultrasound data. In: *2008 IEEE Ultrasonics Symposium* (2012). p. 1318–21.
- Tuthill TA, Sperry RH, Parker KJ. Deviations from rayleigh statistics in ultrasonic speckle. *Ultrason Imaging.* (1988) 10:81–9. doi: 10.1177/016173468801000201
- Masoudi M, Pourreza HR, Saadatmand-Tarzan M, Eftekhari N, Zargar FS, Rad MP. A new dataset of computed-tomography angiography images for computer-aided detection of pulmonary embolism. *Sci Data.* (2018) 5. doi: 10.1038/sdata.2018.180
- Clark KW, Vendt BA, Smith KE, Freymann JB, Kirby JS, Koppel P, et al. The cancer imaging archive (TCIA): maintaining and operating a public information repository. *J Digit Imaging.* (2013) 26:1045–57. doi: 10.1007/s10278-013-9622-7
- Armato SG, McNitt-Gray MF. The lung image database consortium (LIDC) and image database resource initiative (IDRI): a completed reference database of lung nodules on ct scans. *Med Phys.* (2011) 38:2915–31. doi: 10.1118/1.3528204
- Yang D, Xu D, Zhou SK, Georgescu B, Chen M, Grbic S, et al. Automatic liver segmentation using an adversarial image-to-image network. *arXiv arXiv:abs/1707.08037* (2017).
- Huang G, Liu Z, Weinberger KQ. Densely connected convolutional networks. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016). p. 2261–69.
- Ugalde D, Medel JN, Romero C, Cornejo RA. Transthoracic cardiac ultrasound in prone position: a technique variation description. *Intensive Care Med.* (2018) 44:986–7. doi: 10.1007/s00134-018-5049-4

Appendix

TABLE A1 Acoustic properties of different media.

Medium	Impedance [kg/(m ² s)] × 10 ⁶	Attenuation [dB/MHz]	μ_0	μ_1	σ_0	γ	τ
Skull bone	7.8	10.0	0.78	0.56	0.1	0.5	2.0
Soft tissue	1.63	0.75	0.25	0.25	0.25	0.5	1.0
Cardiac muscle	1.71	0.52	0.45	0.51	0.05	0.1	2.25
Blood	1.62	0.18	0.02	0.005	0.015	0.7	1.0
Lung	0.26	41.0	0.10	0.2	0.1	0.9	1.0

TABLE A2 Structures used in the simulation and associated media.

Structure	Associated media
Aorta, pulmonary veins, ventricles and atrias	Blood
Lungs	Lungs
Pericardium	Cardiac muscle
Body	Soft tissue
Ribs	Skull bone

“Body” is the label associated to the tissues between organs.