

Research article

Fast aberration correction in 3D transcranial photoacoustic computed tomography via a learning-based image reconstruction method

Hsuan-Kai Huang^a, Joseph Kuo^a, Yang Zhang^b, Yousuf Aborahama^b, Manxiu Cui^b,
Karteekeya Sastry^b, Seonyeong Park^d, Umberto Villa^c, Lihong V. Wang^b,
Mark A. Anastasio^{d,a} *

^a Department of Electrical and Computer Engineering, University of Illinois Urbana-Champaign, Urbana, 61801, IL, United States

^b Andrew and Peggy Cherng Department of Medical Engineering, California Institute of Technology, Pasadena, 91125, CA, United States

^c Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, 78712, TX, United States

^d Department of Bioengineering, University of Illinois Urbana-Champaign, Urbana, 61801, IL, United States

ARTICLE INFO

Keywords:

Photoacoustic computed tomography

Transcranial imaging

Aberration compensation

Deep learning

ABSTRACT

Transcranial photoacoustic computed tomography (PACT) holds significant potential as a neuroimaging modality. However, compensating for skull-induced aberrations in reconstructed images remains a challenge. Although optimization-based image reconstruction methods (OBRMs) can account for the relevant wave physics, they are computationally demanding and generally require accurate estimates of the skull's viscoelastic parameters. To circumvent these issues, a learning-based image reconstruction method was investigated for three-dimensional (3D) transcranial PACT. The method was systematically assessed in virtual imaging studies that involved stochastic 3D numerical head phantoms and applied to experimental data acquired by use of a physical head phantom that involved a human skull. The results demonstrated that the learning-based method yielded accurate images and exhibited robustness to errors in the assumed skull properties, while substantially reducing computational times compared to an OBRM. To the best of our knowledge, this is the first demonstration of a learned image reconstruction method for 3D transcranial PACT.

1. Introduction

Photoacoustic computed tomography (PACT) is an emerging imaging modality that combines the rich contrast of optical imaging with the high spatiotemporal resolution and imaging depth of ultrasonic imaging, enabling visualization of tissues in deeper regions compared to purely optical high-resolution techniques [1]. The ability of PACT to reveal both structural and functional information, such as vascular connectivity and tissue oxygen saturation, facilitates a broad range of applications, including those in neuroscience. Transcranial PACT brain imaging is one such application that is being actively explored for preclinical [2–4] and clinical [5] applications.

The skull introduces significant challenges for accurate image reconstruction in transcranial PACT of human subjects. As the photoacoustic (PA) pressure field that is induced within brain tissue propagates outward through the skull, it undergoes not only acoustic absorption and scattering but also wave mode conversion at the soft tissue–skull interfaces, with longitudinal-to-shear conversion at the first interface and shear-to-longitudinal conversion at the second interface [6]. These physical factors introduce pronounced aberrations in the PA pressure

field. If these aberrations are not properly accounted for during image reconstruction, artifacts are introduced that can degrade the quality of the reconstructed images [7].

To mitigate the effect of skull-induced aberrations in transcranial PACT, optimization-based reconstruction methods (OBRMs) have been developed [7,8]. These methods are based on the acoustic–elastic wave equation and can account for the relevant three-dimensional (3D) wave physics. While the methods make use of the viscoelastic properties of the skull, only estimates of these quantities are available in practice. Inaccuracies in the assumed skull properties lead to modeling errors that can yield artifacts in the reconstructed images [9]. Such modeling errors can be mitigated by certain OBRMs that jointly estimate the skull properties along with the initial pressure distribution [9]. However, in any case, OBRMs are iterative in nature, where at each iteration the 3D acoustic–elastic wave equation must be solved at least twice. This can result in long image reconstruction times, even if high-performance computing platforms are utilized [10]. This is particularly undesirable in functional or dynamic brain imaging applications, where

* Correspondence to: Department of Bioengineering, 1102 Everitt Lab, 1406 W. Green St, Urbana, IL 61801, USA.

E-mail address: maa@illinois.edu (M.A. Anastasio).

<https://doi.org/10.1016/j.pacs.2025.100698>

Received 19 December 2024; Received in revised form 8 February 2025; Accepted 9 February 2025

Available online 20 February 2025

2213-5979/© 2025 The Authors. Published by Elsevier GmbH. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

many image volumes must be reconstructed and therefore rapid image reconstruction is essential.

Data-driven approaches offer an alternative to overcoming some of the limitations of OBRMs, including long reconstruction times. For example, an image-domain learning-based image reconstruction method for 3D transcranial PACT was reported [11]. The method utilized an image-to-image translation network to map a preliminary estimate of the initial pressure distribution that contained aberration-induced artifacts to a higher-quality one. However, that study did not address inaccuracies in skull property estimation and lacked validation with experimental data.

In another approach, Zhang et al. [12] proposed a data-domain learning-based method to mitigate skull-induced image aberrations. Their learning-based method was trained on two-dimensional (2D)-simulated photoacoustic measurements using simplified 2D numerical head phantoms, where the skull was modeled as a homogeneous viscoelastic medium surrounding the perimeter of the brain. However, the method demonstrated a limited ability to mitigate aberrations in the cortical brain region, which is the primary region of interest for transcranial PACT. The stylized 2D study and nonphysical light delivery designs considered may have contributed to this. Additionally, the study lacked validation with experimental data.

In summary, there remains a need to develop and validate data-driven approaches for transcranial PACT that: (1) address the inherently 3D nature of the problem; (2) can mitigate the impact of uncertainty in the assumed skull properties of the to-be-imaged subject; (3) reduce reconstruction times compared to OBRMs; and (4) are validated using experimental data to ensure practical applicability.

This study aims to develop and investigate a data-driven method for rapid image reconstruction in 3D transcranial PACT that addresses the unmet needs stated above. In the proposed method, a preliminary estimate of the 3D initial pressure distribution is reconstructed using an approximate yet computationally efficient method based on the adjoint of the PACT forward operator [13]. Notably, the adjoint operator utilizes the provided skull geometry and nominal values of the viscoelastic skull properties. The preliminary estimate of the initial pressure distribution is then mapped to a high-quality, deaberrated one through a deep neural network-based image restoration model. The method is systematically investigated in virtual imaging studies that involve ensembles of 3D numerical head phantoms (NHPs) [11,14]. Additionally, the method is evaluated by use of experimental data that were acquired by use of a physical head phantom containing a human skull.

The remainder of the article is structured as follows. In Section 2, an overview of the imaging physics and existing image reconstruction methods for transcranial PACT is provided. In Section 3, the proposed learning-based reconstruction method is described. The design of the virtual imaging studies to evaluate the proposed method is elaborated in Section 4, with the results presented in Section 5. The experimental study using an *ex-vivo* human skull is presented in Section 6. Finally, a discussion of the provided results and conclusions of our studies are provided in Sections 7 and 8.

2. Background

2.1. Transcranial PACT imaging model

Let the photoacoustically-induced stress tensor at location $\mathbf{r} \in \mathbb{R}^3$ and time $t \in [0, \infty)$ be defined as

$$\boldsymbol{\sigma}(\mathbf{r}, t) \equiv \begin{bmatrix} \sigma^{11}(\mathbf{r}, t) & \sigma^{12}(\mathbf{r}, t) & \sigma^{13}(\mathbf{r}, t) \\ \sigma^{21}(\mathbf{r}, t) & \sigma^{22}(\mathbf{r}, t) & \sigma^{23}(\mathbf{r}, t) \\ \sigma^{31}(\mathbf{r}, t) & \sigma^{32}(\mathbf{r}, t) & \sigma^{33}(\mathbf{r}, t) \end{bmatrix}, \quad (1)$$

where $\sigma^{ij}(\mathbf{r}, t)$ denotes the stress in the i th direction acting on the plane perpendicular to the j th direction within the object. Let $\mathbf{v}(\mathbf{r}, t) \in \mathbb{R}^3$ represent the vector-valued acoustic particle velocity. In a heterogeneous

isotropic elastic medium with a diffusive acoustic absorption coefficient distribution $\alpha(\mathbf{r})$, the propagation of $\sigma^{ij}(\mathbf{r}, t)$ and $\mathbf{v}(\mathbf{r}, t)$ can be modeled by the following elastic equations [15–18]:

$$\partial_t \mathbf{v}(\mathbf{r}, t) + \alpha(\mathbf{r})\mathbf{v}(\mathbf{r}, t) = \frac{1}{\rho(\mathbf{r})}(\nabla \cdot \boldsymbol{\sigma}(\mathbf{r}, t)) \quad (2)$$

$$\partial_t \boldsymbol{\sigma}(\mathbf{r}, t) = \lambda(\mathbf{r})\text{tr}(\nabla \mathbf{v}(\mathbf{r}, t))\mathbf{I} + \mu(\mathbf{r})(\nabla \mathbf{v}(\mathbf{r}, t) + \nabla \mathbf{v}(\mathbf{r}, t)^T),$$

subject to the initial conditions,

$$\boldsymbol{\sigma}_0(\mathbf{r}) \equiv \boldsymbol{\sigma}(\mathbf{r}, t)|_{t=0} = -p_0(\mathbf{r})\mathbf{I}, \quad \mathbf{v}(\mathbf{r}, t)|_{t=0} = 0. \quad (3)$$

Here, $p_0(\mathbf{r})$ denotes the photoacoustically-induced initial pressure distribution, $\rho(\mathbf{r})$ represents the medium's density distribution, and $\lambda(\mathbf{r})$ and $\mu(\mathbf{r})$ denote the distributions of Lamé's first and second parameters, respectively. The shear and longitudinal wave speeds are given by $c_s(\mathbf{r}) \equiv \sqrt{\mu(\mathbf{r})/\rho(\mathbf{r})}$ and $c_l(\mathbf{r}) \equiv \sqrt{(\lambda(\mathbf{r}) + 2\mu(\mathbf{r}))/\rho(\mathbf{r})}$, respectively. In a fluid medium such as soft tissue, shear waves can be neglected. In this case, the second Lamé parameter $\mu(\mathbf{r})$ is set to zero, and therefore $c_s(\mathbf{r}) = 0$. The operator $\text{tr}(\cdot)$ in (2) calculates the trace of a matrix, and $\mathbf{I} \in \mathbb{R}^{3 \times 3}$ is the identity matrix. The explicit derivation of these equations is detailed in [19].

Assuming idealized ultrasound transducers, in a PACT experiment, the measured pressure data $p(\mathbf{r}_0, t) = -\frac{1}{3}\text{tr}(\boldsymbol{\sigma}(\mathbf{r}_0, t))$ are recorded at transducer locations $\mathbf{r}_0 \in \partial\Omega$, where $\partial\Omega \subset \mathbb{R}^3$ denotes the measurement aperture. In its discrete-to-discrete (D-D) form, the forward operator that maps the sought-after initial pressure distribution to the measured pressure data can be expressed as [10,13]:

$$\mathbf{g} = \mathbf{M}\mathbf{H}_u^{\text{wave}}\mathbf{p}_0 \equiv \mathbf{H}_u\mathbf{p}_0. \quad (4)$$

Here, the vector \mathbf{g} denotes the lexicographically ordered discrete measurement data acquired at all transducer locations, the vector \mathbf{p}_0 represents a finite-dimensional approximation of the initial pressure distribution, $\mathbf{H}_u^{\text{wave}}$ is the discrete wave propagation operator, and \mathbf{M} represents the discrete sampling operator. In this work, the action of $\mathbf{H}_u^{\text{wave}}$ was computed using the 10th-order staggered grid [20] finite-difference time-domain (FDTD) scheme that involved an absorbing boundary condition. The overall forward operator \mathbf{H}_u represents the composition of the sampling and wave propagation operators. Above, the subscript u is used to emphasize the dependence of the forward operator on the medium's acoustic and elastic property distributions, denoted as $u(\mathbf{r}) \equiv [\rho(\mathbf{r}), \lambda(\mathbf{r}), \mu(\mathbf{r}), \alpha(\mathbf{r})]$. Note that, in practice, only estimates of these property distributions are typically available, which will be denoted by $\hat{u}(\mathbf{r}) = [\hat{\rho}(\mathbf{r}), \hat{\lambda}(\mathbf{r}), \hat{\mu}(\mathbf{r}), \hat{\alpha}(\mathbf{r})]$ and referred to as nominal values.

2.2. Transcranial PACT image reconstruction

Many existing PACT image reconstruction methods do not account for shear waves, which can result in image distortions [6] when applied to transcranial PACT measurement data. The time-reversal method addresses medium viscoelastic heterogeneity and can incorporate shear wave physics [21,22]. However, its effectiveness relies on accurate prior knowledge of the skull's properties. In addition, time-reversal methods are unsuitable for use with acoustically lossy media, for which the time reversibility property does not hold. Furthermore, due to anatomical constraints, transcranial imaging typically employs a hemispherical array covering only the sides and top of the head. Accurate image reconstruction from such data generally requires the use of iterative methods [10,23].

To address these issues, OBRMs for transcranial PACT have been proposed [7,9], in which estimates of \mathbf{p}_0 are computed by minimizing an objective function. A commonly employed objective function is of penalized least squared (PLS) form:

$$\hat{\mathbf{p}}_{0,\text{OBRM}} = \arg \min_{\mathbf{p}_0} \frac{1}{2} \|\mathbf{g} - \mathbf{H}_{\hat{u}}\mathbf{p}_0\|_2^2 + \gamma R(\mathbf{p}_0), \quad (5)$$

where $\hat{\mathbf{p}}_{0,\text{OBRM}}$ is the estimated initial pressure distribution, $\mathbf{H}_{\hat{u}}$ is the D-D forward operator in (4) computed with the nominal acoustic and

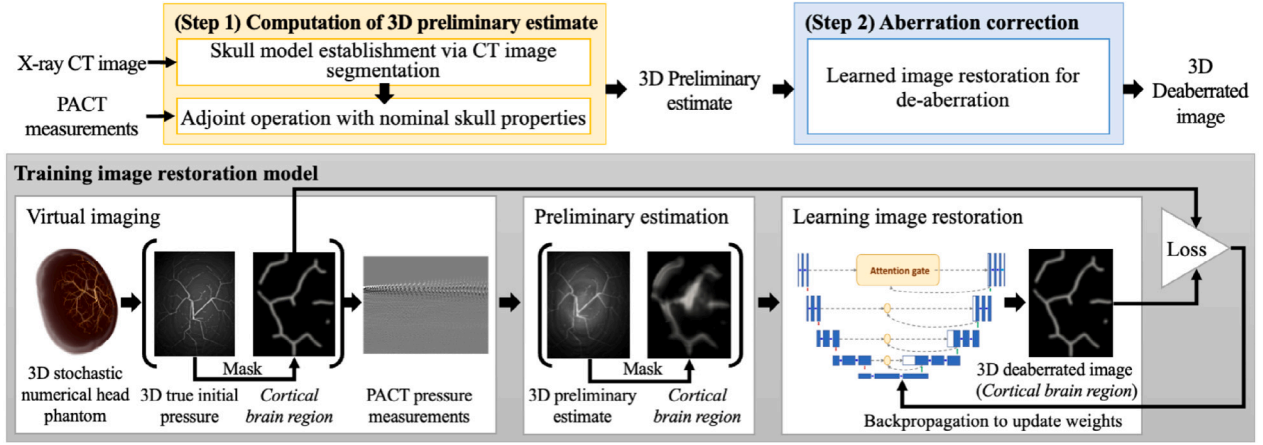


Fig. 1. The proposed reconstruction method (top), which consists of two steps: Computation of a 3D preliminary estimate by applying the adjoint operation with nominal skull properties to the measured data (Step 1) and learned image restoration, where the aberrated cortical brain region within the preliminary estimate of the entire head phantom (the input to the restoration model) is mapped to a deaberrated, high-quality image of the cortical brain region (Step 2). The learning-based method is trained on synthetic data generated from virtual imaging of stochastic numerical head phantoms (bottom).

elastic property distributions $\hat{u}(\mathbf{r})$, γ is a regularization parameter, and $R(\mathbf{p}_0)$ is the penalty term. Gradient- or proximal-based methods can be employed to find approximate solutions of (5), which can yield accurate estimates of the initial pressure distribution [24]. However, inaccuracies in the assumed property distributions $\hat{u}(\mathbf{r})$ can compromise aberration correction and result in image artifacts. Moreover, OBRMs are computationally intensive due to their iterative nature.

In this work, an OBRM for transcranial PACT is implemented as a reference method, against which the proposed data-driven method is compared. The method is designed to solve (5) for the case where the penalty term is specified as the total variation semi-norm [25,26]. The fast iterative shrinkage-thresholding algorithm [26] is employed to approximately solve the corresponding optimization problem. In the special case where $\gamma = 0$ (no penalty term), Nesterov accelerated gradient descent [27] is employed as the optimization method.

3. Learned image reconstruction method

A subject-specific, learning-based, image reconstruction method for 3D transcranial PACT is proposed. It is assumed that the geometry of the skull is accurately known from adjunct X-ray CT or magnetic resonance imaging [28] studies of the subject. The proposed method is designed to mitigate the impact of forward modeling errors caused by differences between the assumed nominal skull viscoelastic properties and the true properties of the subject's skull. The method will also provide a dramatic reduction in image reconstruction times as compared to conventional OBRMs.

The proposed method involves two steps, as illustrated in Fig. 1 (top). In Step 1, the acoustic and elastic property distributions in the subject's head are numerically modeled by use of a parameterized version of the skull geometry and nominal values of the acoustic and elastic properties of the skull and surrounding soft tissues [11,14]. Based on this head model, a 3D preliminary estimate of the initial pressure distribution $\tilde{\mathbf{p}}_0$ is reconstructed by applying the adjoint of the forward operator in (4) to the measured data:

$$\tilde{\mathbf{p}}_0 = \mathbf{H}_a^\dagger \mathbf{g}, \quad (6)$$

where \mathbf{H}_a^\dagger denotes the adjoint of the forward operator based on nominal properties $\hat{u}(\mathbf{r})$. Although the adjoint operator is not an accurate image reconstruction operator, it can partially compensate for skull-induced aberrations and can be computed relatively efficiently. Most importantly, it transforms the measured data into the image domain, streamlining the subsequent processing step.

In Step 2, a learned image restoration model is employed to mitigate the artifacts in $\tilde{\mathbf{p}}_0$ and yield a final high-quality estimate of the initial pressure distribution. This is achieved as

$$\hat{\mathbf{p}}_{0,\text{learn}} = A_\theta(\tilde{\mathbf{p}}_0), \quad (7)$$

where A_θ denotes an image restoration network that is parameterized by weights θ and maps the preliminary estimate $\tilde{\mathbf{p}}_0$ to the deaberrated, higher-quality, estimate $\hat{\mathbf{p}}_{0,\text{learn}}$ as depicted in Fig. 1 (bottom).

The adopted architecture for the image restoration model was a 3D U-Net [29] with several modifications. Attention gates [30] were incorporated to focus on the cortical brain region. The number of output features in the first convolution layer was lowered from 64 to 24 to reduce the required GPU memory. Batch normalization [31] was applied after each convolution to accelerate training and mitigate overfitting [32]. Inspired by the VGG [33] and ResNet [34] models, the max-pooling and upsampling layers were replaced with strided convolution layers, each with a stride of 2.

The image restoration model is trained to minimize, with respect to θ , the mean squared error (MSE) loss function:

$$L(\theta) = \frac{1}{2} \sum_{k=1}^K \|A_\theta(\tilde{\mathbf{p}}_0^k) - \mathbf{W} \odot \mathbf{p}_0^k\|_F^2. \quad (8)$$

Here, K represents the number of training samples, $\|\cdot\|_F$ denotes the Frobenius norm, \mathbf{W} is the masking matrix used to emphasize the cortical region (the primary region of interest for transcranial PACT) during the loss computation, \odot denotes the element-wise product, and \mathbf{p}_0^k and $\tilde{\mathbf{p}}_0^k$ denote the 3D target and preliminary estimate of initial pressure distribution for the k th training sample, respectively.

The image restoration model is trained in a fully supervised manner, which requires the true (non-aberrated) initial pressure distribution \mathbf{p}_0 as the target. By building well-characterized physical phantoms that incorporate clinical variability, it is conceptually possible to experimentally acquire pairs of target initial pressure distributions and corresponding PACT measurement data. However, generating a large number of such phantoms and measurement data required for training is impractical due to the significant cost and time involved.

To make the generation of large training datasets feasible, stochastic 3D NHPs [11,14] and high-fidelity simulation of the PACT data acquisition process are utilized in this study. The k th synthetic data sample consists of $\{u^k, \mathbf{p}_0^k, \mathbf{g}^k = \mathbf{H}_{u^k} \mathbf{p}_0^k, \tilde{\mathbf{p}}_0^k = \mathbf{H}_a^\dagger \mathbf{g}^k\}$, where u^k corresponds to the acoustic and elastic property distributions of the k th NHP. The quantity \mathbf{p}_0^k denotes the target image (i.e., ground truth initial pressure distribution), while the restoration model input $\tilde{\mathbf{p}}_0^k$ is computed via (6)

from \mathbf{g}^k . The generation of \mathbf{p}_0^k and \mathbf{g}^k , along with the virtual imaging process involved, is detailed in Sections 4.1 and 4.2. The flexibility of configurable stochastic 3D NHPs [11,14] and virtual imaging facilitates the generation of synthetic data ensembles under various scenarios, ranging from ideal to more practical conditions. This capability enables the systematic investigation of the proposed method in virtual imaging studies, as in Section 4.4.

To assess its practical applicability and effectiveness of the proposed method, the learning-based reconstruction method is further evaluated by use of experimental phantom data. Specifically, the method is trained exclusively on synthetic data and employed directly to reconstruct images from experimental data. The experimental data were acquired by imaging a physical head phantom using an existing 3D PACT imaging system [35]. Details regarding the experimental validation and results are provided in Section 6.

4. Virtual imaging studies

This section provides details of the virtual (i.e., *in-silico*) imaging studies conducted for the systematic investigation of the proposed learning-based reconstruction method. The establishment of the stochastic 3D NHPs for transcranial PACT and the virtual imaging process performed to generate synthetic datasets are detailed in Sections 4.1 and 4.2, respectively. The training procedure for the learning-based method, which utilizes subject-specific synthetic data, is described in Section 4.3. Details of the virtual studies conducted to assess the reconstruction accuracy, computation time, and generalizability of the learning-based method are provided in Section 4.4.

4.1. Generation of stochastic numerical head phantoms

Stochastic 3D NHPs were employed to generate a subject-specific synthetic dataset for model training, providing ground truth that is necessary for supervised learning but is typically unavailable in experimental scenarios. A single NHP consisted of an optical phantom and an acoustic phantom to model optical and acoustic/elastic properties, respectively, of the head. The NHP was composed of soft tissues outside the skull (scalp vessels and scalp tissue), the skull, and soft tissues inside the skull (cortical vasculature and extravascular tissue within the cortical brain region). The procedure for generating subject-specific stochastic 3D NHPs is outlined below. Additional details can be found in [11,14].

The subject's skull geometry was assumed to be known from adjunct CT or magnetic resonance imaging (MRI). In this study, the skull geometry was modeled as a three-plate, three-layer structure [9] derived from adjunct CT data. Similarly, the geometry of soft tissues outside the skull was also assumed to be known, as these tissues can often be reconstructed with high fidelity in practice without accounting for the skull. In contrast, the geometry of cortical vasculature was stochastically synthesized [11,14,36], which reflects that the configuration of the cortical vessels in the subject is unknown. In summary, in each realization of the subject-specific NHP, the geometry of the cortical vasculature was distinct, while the geometry of the skull and tissues outside the skull remained fixed and specific to the to-be-imaged subject.

To form the optical phantom for each NHP, nominal values of tissue-specific optical properties from the literature [11,14] were used. The assignment of optical properties was a deterministic process, in which the same tissues in different NHPs were assigned the same optical property values. To form the acoustic phantom for each NHP, the acoustic properties of the soft tissues were assumed to be the same as those of water. On the other hand, varying degrees of stochasticity were considered when assigning viscoelastic properties to the skull, corresponding to scenarios where the amount of uncertainty in the nominal property values differed. Specifically, three distinct statistical assumptions regarding the skull's viscoelastic properties were

Table 1

Nominal Acoustic and Elastic Properties of Skull and Soft Tissues in the *In-silico* Study [9].

Medium	ρ [kg/mm ³]	c_l [mm/ μ s]	c_s [mm/ μ s]	α [1/ μ s]
Soft Tissues (Water)	1000	1.48	0	0
Diploë	1000	2.10	1.40	0.75
Frontal Bone	1850	2.03	1.40	0.5
Left Parietal	1850	2.76	1.50	0.5
Right Parietal	1850	2.55	1.45	0.5

employed, leading to three variants of NHPs.

In the first variant, the viscoelastic properties of each skull plate and layer were deterministically assigned the nominal values specified in Table 1. In this case, the corresponding skull components in different acoustic phantoms were assigned the same value. This represents an idealized scenario where the nominal skull properties are assumed to be accurate estimates of the true subject properties. The ensemble of NHPs created in this way is referred to as **Dataset 1**.

In the second variant of NHPs, the viscoelastic properties of each skull plate and layer were assigned stochastically. In this case, the corresponding skull components in different acoustic phantoms were assigned, in general, different values that were specified by sampling from a distribution. Specifically, the properties of each skull plate and layer were modeled as normally distributed, with nominal values listed in Table 1 as mean values and standard deviation values of 5% of the nominal values. The ensemble of NHPs created this way is referred to as **Dataset 2**. In the third variant of NHPs, an identical procedure was followed but the standard deviations were increased to 10% of the nominal values. This represents a scenario in which there is an increased degree of uncertainty associated with the assumed nominal property values. In this case, the ensemble of NHPs is referred to as **Dataset 3**.

Datasets 2 and 3 are relevant to practical scenarios where the nominal (but not necessarily true) skull properties are known, but variations in these values are considered when generating training data. A learned reconstruction method trained using such data may exhibit a certain robustness to errors in the assumed nominal skull properties. This will be investigated in the virtual imaging studies described in Section 4.4.

A total of 1000 NHPs were generated for each of the three datasets. The generated NHPs had dimensions of $500 \times 154 \times 385$ voxels ($187.5 \times 57.75 \times 144.375$ mm³) with a voxel size of 375 μ m. Additional details related to NHP construction can be found elsewhere [11,14].

4.2. Virtual transcranial PACT imaging

A 3D physics-based simulation framework was utilized to virtually image the NHPs in Datasets 1–3 and produce pairs of initial pressure maps and corresponding synthetic PA measurement data. This framework involved simulating optical illumination to generate the 3D optical fluence distribution, followed by simulations of acoustic and elastic wave propagation and sampling associated with data acquisition.

The virtual imaging was performed by emulating a real-world 3D imaging setup [7]. A schematic of the simulated PACT imaging system is illustrated in Fig. 2. For light delivery, an optical laser cone beam was employed, positioned above the head, and directed downward. The cone beam had a half-angle of 46° and a pulse energy of 2.2 J. Acoustic measurements were acquired using an arc-shaped ultrasonic transducer probe that rotated around the head in 2012 discrete steps, forming a hemispherical measurement geometry. The probe, with a radius of 13 cm, was equipped with 256 transducer elements, resulting in a total of 515,072 transducer element locations. Data acquisition spanned 160 μ s following optical exposure, with a sampling frequency of 30 MHz, yielding 4800 time samples.

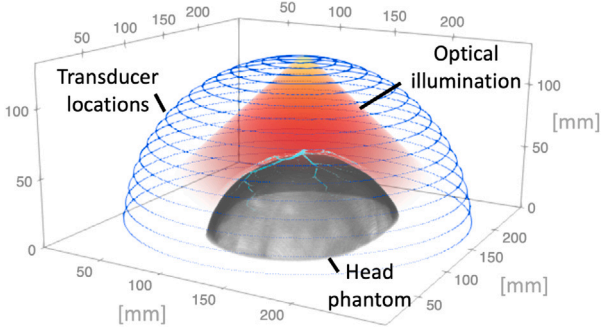


Fig. 2. Schematic of the PACT imaging system, illustrating the placement of the numerical head phantom within the simulated measurement geometry.

Based on the configuration described above, optical illumination was simulated using the Monte Carlo eXtreme (MCX) library [37] to compute the 3D optical fluence distribution and the resulting initial pressure distribution. An example of the simulated initial pressure distribution is illustrated in Fig. 3. In this example, the initial pressure in the scalp vasculature was an order of magnitude higher than that in the cortical vasculature due to optical attenuation. This difference in magnitude required the application of field-of-view (FOV) masking to effectively visualize the cortical brain region. Subsequently, the D-D imaging operator \mathbf{H}_u defined in (4) was utilized to compute noiseless pressure time-traces at each transducer location. The imaging operator was implemented using the FDTD open source library, Devito [38,39]. The simulation domain for both the optical and acoustic physics simulations had a voxel size of 375 μm and dimensions of $704 \times 352 \times 704$ voxels.

To simulate noisy PA pressure data, electronic noise was simulated and added to the noiseless pressure data. The noise was modeled as additive, independent, and identically distributed Gaussian white noise, with its standard deviation, i.e., noise level, determined based on the signal-to-noise ratio (SNR) of the real imaging system [5]. Using the maximum pressure amplitude observed across the entire ensembles of cortical vasculature as the signal strength, the computed SNR was 1.83. Two additional noise levels corresponding to SNRs of 0.92 and 0.37 were also considered.

4.3. Implementation and training

The adjoint of the forward operator \mathbf{H}_u^\dagger employed in Step 1 of the proposed method was implemented using the Devito framework [38, 39]. The learning-based image restoration model A_θ in Step 2 was implemented using Pytorch [40]. To train the image restoration model, the MSE loss function in (8) was minimized by use of the Adam optimizer with decoupled weight decay regularization [41]. The dataset was divided into a training set for updating model weights and a validation set for calculating the validation loss at each epoch. Training was terminated when the validation loss did not decrease for 10 consecutive epochs.

All computations were performed on a compute node equipped with an Intel Xeon Gold 5218 CPU and eight Nvidia Quadro RTX 8000 GPUs. For the evaluation of computation speed across different reconstruction methods, only a single GPU was employed.

4.4. Study designs

Studies were conducted to systematically evaluate the performance of the learning-based method under different assumptions regarding the subject's skull properties. In **Study 1**, the subject's skull properties were assumed to be known. This represents an idealized scenario. **Study 2** addressed a more practical scenario in which the subject's

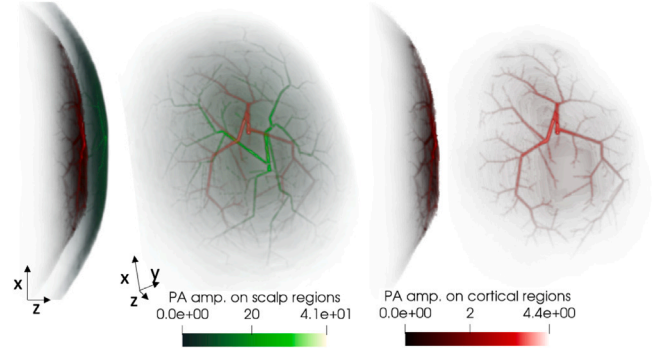


Fig. 3. 3D rendering of the initial pressure distribution without (left) and with cortical field-of-view masking (right) of one example NHP in virtual imaging studies. In the right image, the field-of-view was cropped from the left image to visualize the cortical brain region. Paraview [42] was used for volume rendering.

true skull properties were unavailable but were known statistically. Specifically, the properties of each skull plate and layer were modeled as normally distributed, with nominal values listed in Table 1 defining the means and standard deviation values specified as 5% of the nominal values. Both training and testing data were created by use of the same statistical skull model and were therefore statistically consistent. Similarly, in **Study 3**, the true skull properties remained unknown and were prescribed statistically. However, unlike in Study 2, an additional scenario was considered in which the assumed statistical skull model employed to create the training data was not correctly specified. Specifically, the statistical properties of the training and testing data differed due to inaccurately specifying the degree of uncertainty in the nominal skull property values when creating the training data. Finally, in **Study 4**, the model trained in Study 2 was re-evaluated on different testing data in which the previously employed nominal optical property values of scalp vessels differed from the true values. This mismatch in the scalp vessel optical properties can arise due to variations in the imaging setup and physiological conditions.

Note that Studies 1 and 2 correspond to in-distribution assessments, in which the testing data characteristics matched those of the training data. In contrast, Studies 3 and 4 were conducted under more practical and challenging out-of-distribution conditions, where the testing data were statistically distinct from the training data.

Both visual and quantitative assessments were performed. Paraview [42] was used to visualize reconstructed initial pressure distributions within cortical brain regions through 3D volume rendering. Quantitative assessments were conducted using MSE and structural similarity index (SSIM). Ensemble-averaged evaluation metrics were computed for Studies 1–3.

4.4.1. Study 1 implementation details

Dataset 1 and the associated simulated measurement data were employed in this study. Four versions of the learning-based method were individually trained on noiseless PA measurements and noisy measurements at three distinct noise levels, corresponding to SNRs of 0.37, 0.92, and 1.83. For each model, 400 data pairs of initial pressure maps and corresponding pressure measurements were used for training, while 60 pairs from the same dataset were used for validation. Quantitative evaluation was conducted on a test set of 20 NHPs.

The images reconstructed using the learning-based method were compared to those reconstructed using the universal backprojection (UBP) method, the adjoint operator (the same operator used in Step 1 of the learning-based method), and the OBRM defined in Section 2.2. In all reconstruction methods that demanded skull properties, the nominal values in Table 1 were employed. In all cases, the images retained the same dimensions and voxel size as those of the simulation domain.

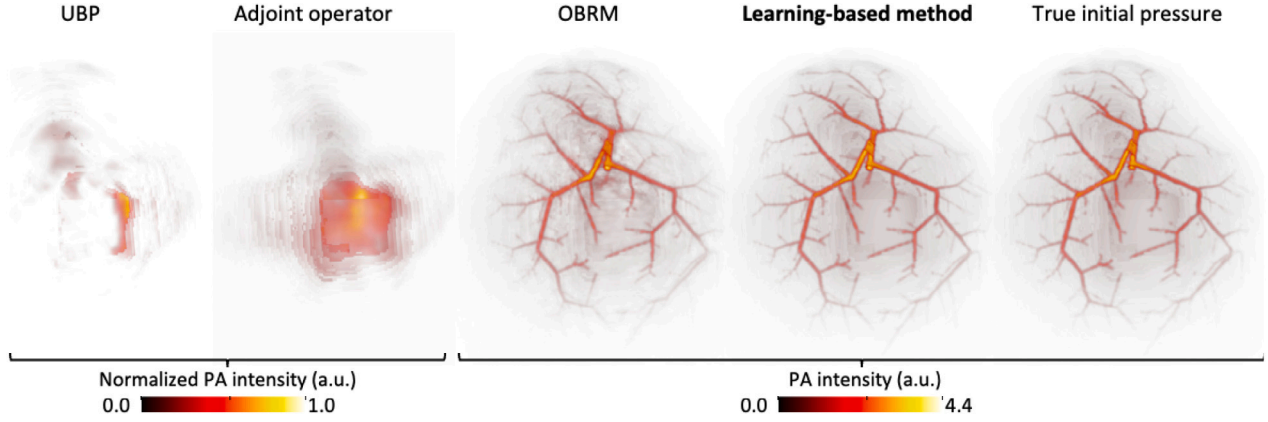


Fig. 4. The results of the virtual imaging Study 1 when skull properties were known accurately. The estimated initial pressure distributions reconstructed from noiseless PA measurements using the UBP method, the adjoint operator, the OBRM, and the learning-based method (first to fourth from left) are presented. The true initial pressure distribution (far right) is provided as a reference. When the skull properties were accurately known, the learning-based method (fourth) effectively compensated for artifacts caused by skull-induced aberrations, achieving image quality comparable to the OBRM (third).

In the UBP method, the speed-of-sound value that yielded the lowest MSE between the normalized true and estimated initial pressure distributions within the whole head region was selected through parameter scanning. The stopping criterion for the OBRM required that the squared Frobenius norm of the difference between two successive iterations $\|\mathbf{p}_0^{(i)} - \mathbf{p}_0^{(i-1)}\|_F^2$, normalized by its maximum $\max_{1 \leq i} \|\mathbf{p}_0^{(i)} - \mathbf{p}_0^{(i-1)}\|_F^2$ over the previous iterations, fell below a user-defined threshold of $\epsilon = 2 \times 10^{-4}$. This criterion is equivalent to monitoring the norm of the gradient in smooth optimization [43]. The regularization parameter was determined through parameter scanning to provide the lowest MSE between the true and estimated initial pressure distributions within the skull. While this criterion is impractical as the true distribution is typically unknown, it was intentionally applied in virtual imaging studies to provide the most accurate possible reference estimate. The regularization parameter was estimated for each measurement noise level considered. During the optimization, a positivity constraint was applied. Due to the high computational burden of the OBRM, quantitative assessments were performed on a test set of 20 NHPs at only a single noise level corresponding to an SNR of 1.83.

4.4.2. Study 2 implementation details

Dataset 2 and the associated noisy measurement data corresponding to an SNR of 1.83 were employed in this study. The OBRM, defined in Section 2.2, with the regularization parameter determined in Study 1, served as the reference method. The learning-based method was trained on 900 data pairs and validated with 60 pairs from the same dataset. For image reconstruction with both the OBRM and learning-based method, nominal skull property values (i.e., ensemble mean values) were employed. Quantitative evaluation metrics were calculated from the reconstructed images of 20 test NHPs to compare the performance of the learning-based method to that of the OBRM.

4.4.3. Study 3 implementation details

Datasets 1–3 and the associated noisy measurement data corresponding to an SNR of 1.83 were employed in this study. As defined in Section 4.1, these data sets varied in the level of uncertainty associated with the skull properties in the NHPs: Datasets 1 (no uncertainty), 2 (moderate uncertainty), and 3 (high uncertainty). Three versions of the learning-based method were trained separately on measurement data corresponding to each of the three datasets. The trained models were then tested across all three datasets, resulting in nine groups of reconstructed images. As in Study 2, the nominal values of the skull properties were used for image reconstruction. Quantitative measures were calculated using 20 test NHPs from each dataset.

4.4.4. Study 4 implementation details

In practice, the optical properties of scalp vessels can depend on factors such as the laser wavelength used and physiological conditions, including oxygen saturation. To examine the impact of inaccuracies in the assumed scalp optical properties on the learned reconstruction method established in Study 2, testing data from twelve additional NHPs were generated from a single NHP in Dataset 1. These additional NHPs differed in the optical absorption coefficients of their scalp vessels. Specifically, the absorption coefficient value used in the NHPs differed by 5% increments from each other, ranging from -30% to $+30\%$, relative to the nominal value specified in [11,14] that was employed in Studies 1–3. These NHPs were virtually imaged, and pressure data were simulated at a noise level corresponding to an SNR of 1.83. Qualitative and quantitative assessments of the learning-based method's robustness to uncertainties in the optical absorption properties of the scalp vessels were conducted based on the images reconstructed from these data.

5. Virtual imaging results

5.1. Study 1 results

As described in Section 4.4.1, the learning-based image reconstruction method was first evaluated under idealized conditions in which the subject's skull properties were known. Fig. 4 shows estimates of the 3D initial pressure distributions within the cortical brain region, reconstructed from noiseless measurements. From left to right, the estimates correspond to the UBP method (first), the adjoint operator (second), the OBRM defined in Section 2.2 (third), and the learning-based method (fourth). The true initial pressure distribution is provided on the far right of the figure. When using the UBP method, designed for an acoustically homogeneous and lossless inviscid medium, the cortical vasculature was strongly contaminated by image artifacts resulting from skull-induced aberrations. The preliminary estimate computed using the adjoint operator in Step 1 of the learning-based method also exhibited significant artifacts. In contrast, the OBRM successfully mitigated the artifacts caused by the skull-induced aberrations, reconstructing the cortical vasculature with high fidelity when the skull properties were accurately known. The learning-based method produced results comparable to the OBRM, demonstrating similar effectiveness in compensating for skull-induced aberrations.

One major advantage of the learning-based method over the OBRM is its substantially reduced computation time. The computation times

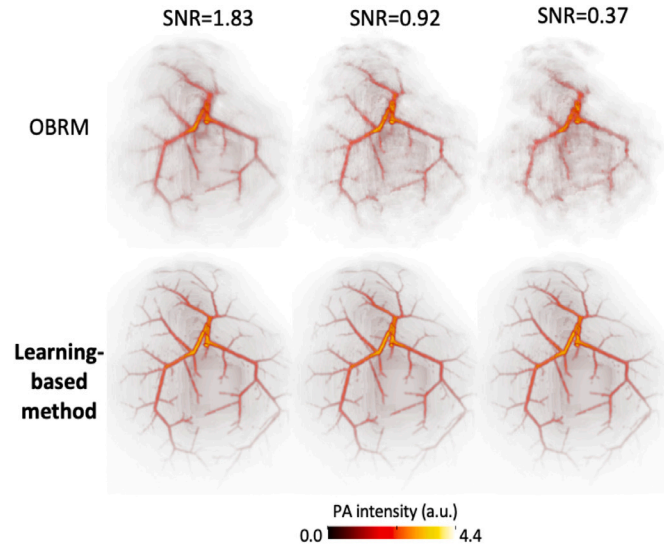


Fig. 5. The results of Study 1 when measurement noise was considered. The estimated initial pressure distributions reconstructed using the OBRM (first row) and the learning-based method (second row) are presented at different noise levels corresponding to SNRs of 1.83, 0.92, and 0.37 (left to right). The learning-based method (second row) demonstrated relatively consistent image quality across varying noise levels, unlike the OBRM (first row).

Table 2
Computation Time Comparison in *In-silico* Study.

Method	Computation time		
	Prelim. Est.	Deaberration	Total
Learning-based method	0.258 h	<1 s	0.258 h
	Per iteration	# of iterations	Total
OBRM	0.639 h	150	96 h

Table 3
Ensemble-averaged MSE and SSIM Comparison in Study 1.

Method	Noise level	MSE ($\times 10^{-4}$)	SSIM ($\times 10^{-2}$)
Learning-based method	Noiseless	1.73 ± 0.27	99.776 ± 0.015
	SNR=1.83	1.71 ± 0.26	99.767 ± 0.016
	SNR=0.92	1.82 ± 0.31	99.766 ± 0.019
	SNR=0.37	2.16 ± 0.36	99.704 ± 0.026
OBRM	SNR=1.83	16.83 ± 1.26	92.856 ± 0.111

for producing the results in Fig. 4 are summarized in Table 2. Once trained, the learning-based method was 350 times faster than the OBRM in image reconstruction. Although reconstruction times for the OBRM can vary depending on the imaging system configuration, the Lipschitz constant, and stopping criteria, the learning-based method is still expected to offer substantial acceleration. Establishing the training dataset (500 data pairs) and training the method required approximately 258 h and 40 h, respectively, but these only need to be performed once for a given subject. Further discussion is provided in Section 7.

Images reconstructed from noisy measurements are provided in Fig. 5. While the quality of the images produced by the OBRM (top row) degraded as the noise level increased, the quality of the images produced by the learning-based method (bottom row) remained high across all three noise levels.

To further quantify the quality of images reconstructed by use of the learning-based method, the ensemble-averaged MSE and SSIM values are provided in Table 3. As the noise levels increased, the images reconstructed using the learning-based method showed a slight rise in MSE and a decrease in SSIM, indicating degradation. For each noise

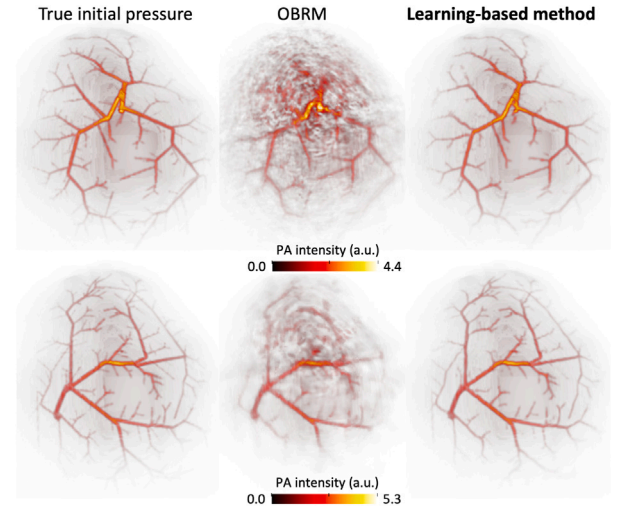


Fig. 6. The results of Study 2 when inaccuracies were present in the assumed nominal skull properties. True (left) and estimated initial pressure distributions reconstructed using the OBRM (middle) and the learning-based method (right) are shown for two example NHPs (first and second rows). The OBRM-reconstructed images (middle) exhibited significant degradation, with artifacts caused by modeling errors due to skull property mismatches. In contrast, the learning-based method (right) effectively mitigated these artifacts, demonstrating enhanced robustness to inaccuracies in the assumed skull properties.

Table 4
Ensemble-averaged MSE and SSIM Comparison in Study 2.

Method	MSE ($\times 10^{-4}$)	SSIM ($\times 10^{-2}$)
Learning-based method	4.43 ± 1.60	99.466 ± 0.137
OBRM	69.12 ± 36.07	83.156 ± 3.264

level, the learning-based method outperformed the OBRM.

5.2. Study 2 results

As described in Section 4.4.2, the performance of the learning-based method was assessed when the subject's skull properties were only known statistically. The nominal property values assumed by the OBRM will generally differ from the actual skull property values of the subject, and this can cause artifacts in the reconstructed images. Two examples are shown in Fig. 6, where the images reconstructed using the OBRM exhibited strong artifacts that significantly deteriorated image quality. In contrast, the learning-based method effectively mitigated these artifacts, although minor degradation in the reconstructed thin vasculature was observed.

The ensemble-averaged MSE and SSIM values are presented in Table 4. These data corroborate that, in the presence of skull property mismatches, the learning-based method significantly outperformed the OBRM.

5.3. Study 3 results

As described in Section 4.4.3, the accuracy of the learning-based reconstruction method was evaluated under conditions where the statistical properties of the training and testing data differed. This condition arose from inaccurately specifying the extent of uncertainty in the nominal skull property values when creating the training data. Three versions of the learning-based method were trained by use of Datasets 1–3 and the associated noisy measurement data, and employed to reconstruct images from noisy measurement data from a Dataset 2 sample. The resulting images, which are shown in Figs. 7 (a)–(c), correspond to scenarios of underestimated, accurately estimated, and

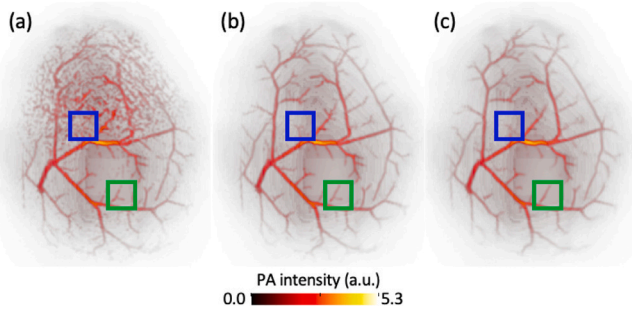


Fig. 7. The reconstructed initial pressure distribution of Study 3 when uncertainty in the distributions of human skulls properties were unknown. The results were reconstructed using the learning-based method under scenarios of (a) underestimated, (b) correctly estimated, and (c) overestimated uncertainties in the skull's viscoelastic properties. The blue and green boxes highlight regions where differences between the scenarios are noticeable. When the uncertainty in the skull properties was underestimated, the learning-based method introduced image artifacts, as shown in (a).

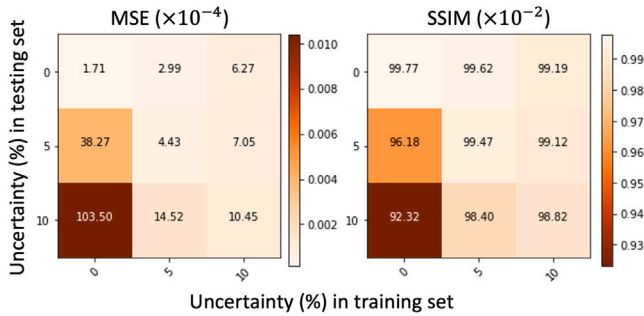


Fig. 8. Ensemble-averaged MSE (left) and SSIM (right) calculated from nine groups of initial pressure distributions reconstructed using the learning-based method. Brighter colors in both color bars indicate higher accuracy. Diagonal components represent cases where the level of skull property uncertainty (i.e., skull property statistics) was accurately estimated, while the bottom-left and top-right corners represent scenarios with underestimated and overestimated uncertainty, respectively. When skull property statistics are not well characterized, underestimating uncertainty may lead to significant image quality degradation, as reflected in the MSE and SSIM metrics.

overestimated skull property uncertainty, respectively. Compared to the accurately estimated scenario in Fig. 7 (b), underestimating skull property uncertainty resulted in artifacts in the reconstructed image, as illustrated in Fig. 7 (a). These artifacts were particularly noticeable in the frontal bone area, where relatively large mismatches in properties were observed between the training and test sets. In contrast, when the skull property uncertainty was overestimated, only minor image degradation was observed, as shown in Fig. 7 (c). Such degradation primarily affected thin vasculature, as highlighted by blue and green boxes in Fig. 7 (c).

As explained in Section 4.4.3, nine groups of images were reconstructed using the learning-based method. Each group contains 20 estimated initial pressure distributions, representing a particular combination of training and testing conditions across the three levels of skull property uncertainty reflected by Datasets 1–3. Fig. 8 shows the ensemble-averaged MSE (left) and SSIM (right) of the nine groups. Diagonal components represent situations where the skull property uncertainty was matched between training and testing data, while the bottom-left and top-right corners represent situations where the skull property uncertainty was underestimated or overestimated, respectively. The performance of the learning-based method for the underestimated uncertainty group was inferior to that in both the properly estimated and overestimated groups. Notably, the MSE and SSIM results from the overestimated group were comparable to those from the

accurately estimated group. These findings align with the observations in Fig. 7, suggesting that accounting for greater uncertainty when producing training data may be prudent when the statistics of human skull properties are not well understood.

5.4. Study 4 results

As described in Section 4.4.4, the performance of the learning-based reconstruction method established in Study 2 was re-evaluated when applied to noisy measurement data produced by NHPs whose scalp vessels had optical absorption properties that differed from those in the training data. When the variations in the scalp vessel absorption coefficients of the testing data were $\pm 10\%$ relative to the nominal value, the images reconstructed using the learning-based method (Figs. 9 (c) and (e)) exhibited comparable visual image quality to the one in which the optical properties of the training and testing data were fixed at the nominal value (Fig. 9 (d)). However, noticeable artifacts emerged in the reconstructed images when the variation reached $\pm 20\%$ or more, as illustrated in Figs. 9 (a), (b), (f), and (g).

The MSE and SSIM of the reconstructed images are shown in Fig. 10. For reference, the MSE and SSIM values from Study 2 are also included (solid line: mean; dashed lines: $\pm\sigma$). Performance degradation was evident as the difference in the optical absorption coefficient of the scalp vasculature increased, indicated by rising MSE values and decreasing SSIM values. This trend is consistent with subjective assessments of the images in Fig. 9. For a $\pm 10\%$ variation, the metrics remained within the reference range, suggesting that the learning-based method, despite being trained with fixed nominal scalp optical properties, demonstrated some tolerance to changes in the scalp properties. The relevance of this to physiological variability in transcranial PACT is discussed in Section 7.

6. Validation using experimental phantom data

The learning-based reconstruction method was further validated by use of experimental data corresponding to a physical head phantom that contained a human skull.

6.1. Experimental setup

A physical head phantom based on an adult human skull was employed as displayed in Fig. 11. To simulate cortical vasculature, a black wire was placed beneath the skull. Three variations of the head phantom were created by altering the shapes of wire mimicking the cortical vasculature. The wire was positioned at a distance from the skull ranging from 3 to 16 mm. The extravascular tissue within the cortical brain region was represented by water. No scalp tissues were included, except for fiducial photoacoustic markers implanted on the skull to ensure accurate spatial registration with the adjunct CT image. The phantom was imaged by use of the experimental setup illustrated in Fig. 11. To improve the SNR of the PA signal from the black wire, an internal light source emitting diffused light at a wavelength of 1064 nm was utilized to illuminate the region of interest from inside the skull, as illustrated in Fig. 11. Details of the imaging system and data acquisition procedures can be found in [35]. Reference images were obtained by acquiring additional PACT measurements without the skull, with each wire positioned identically to its location in the physical head phantom.

6.2. Application of the learning-based method

As explained in Section 3, the proposed learning-based method, trained on synthetic data, was validated by use of the experimental data. To generate synthetic training data that were representative of the physical head phantom, 3D NHPs were created by registering the adjunct CT data based on the fiducial markers on the skull of the physical head phantom. In the 3D NHPs, cortical vessels were

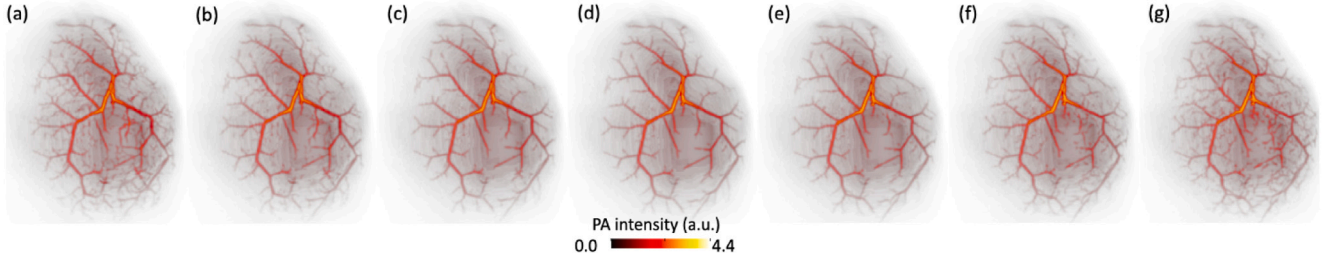


Fig. 9. Initial pressure distributions reconstructed using the learning-based method for NHPs with the optical absorption coefficient values of the scalp vasculature adjusted to -30%, -20%, -10%, 0%, +10%, +20%, and +30% (a to g) relative to the nominal value specified in [11,14]. Additional image artifacts appeared when the absolute variation in the optical absorption coefficient of the scalp vasculature reached or exceeded 20%.

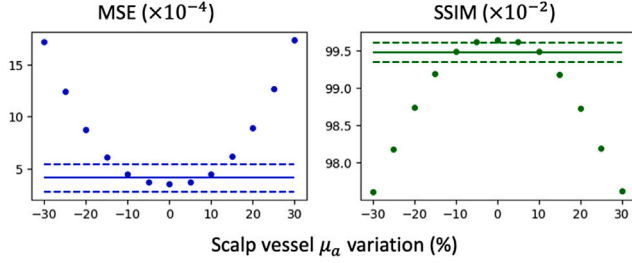


Fig. 10. MSE (left) and SSIM (right) of the reconstructed images for different optical absorption coefficients of scalp vessels. The solid line represents the mean metric value for the learning-based method in Study 2, with the dashed lines indicating the mean \pm standard deviation. The learning-based method demonstrated robustness to $\pm 10\%$ variations in the optical absorption coefficients of scalp vessels, despite being trained with fixed optical properties.

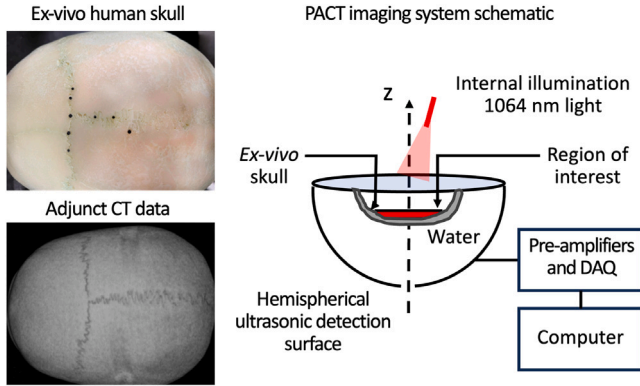


Fig. 11. A photograph and adjunct CT image of the ex-vivo human skull used in the experimental study (left) and a schematic of the PACT imaging system (right).

synthesized within the region corresponding to the region of interest, where the black wire was positioned in the physical head phantom. The optical and acoustic phantoms for each NHP were formed by assigning nominal property values from the literature [11,14,44] to each tissue type. Particularly, the viscoelastic properties of each skull plate and layer were modeled as normally distributed with means equal to the nominal values specified in Table 5 and standard deviation values equal to 10% of their respective nominal values. In Table 5, the density and wave speed values of the skull plate layer were adopted from [44] as specified in [35].

As detailed in Section 6.1, the physical head phantoms included the wire and optical fiducial markers, with the remaining medium consisting of the skull and water, and internal illumination was employed in the experiment. This configuration led to minimal optical attenuation within the water region. Based on these experimental characteristics, the initial pressure distributions corresponding to the NHPs

Table 5

Acoustic and Elastic Properties of Skull and Soft Tissues in the Ex-vivo Study.

Medium	ρ [kg/mm ³]	c_l [mm/ μ s]	c_s [mm/ μ s]	α [1/ μ s]
Soft Tissues (Water)	1000	1.465	0	0
Diploë	1000	2.1	1.4	0.75
Frontal Bone	1800	2.8	1.35	0.5
Left Parietal	1800	2.8	1.35	0.5
Right Parietal	1800	2.8	1.35	0.5

were computed using a two-region approach, with different methods applied to the regions inside and outside the skull. Assuming a uniform optical fluence within the cortical brain region, a constant fluence value was determined as follows to align with the experiment. Five NHPs were randomly chosen, and their estimated initial pressure distributions were computed via (6), assuming an optical fluence value of 1. The ratio between the maximum values of these estimates and those of the experimental data was calculated, and the optical fluence value was adjusted to ensure that the mean of the maximum values was consistent between the synthetic and experimental data. This nominal optical fluence value was then applied to all synthetic data. For regions outside the skull, where optical markers were the primary sources of initial pressure, estimates obtained using the OBRM were adopted. Simulations of acoustic and elastic wave propagation and sampling associated with data acquisition were conducted as described in Section 4.2.

The learning-based method was trained using 400 data pairs of initial pressure maps and corresponding noiseless pressure measurements, while 60 pairs were used for validation. Subsequently, images were reconstructed from experimental pressure data by use of the UBP method, the adjoint operator, the OBRM, and the learning-based method. All but the UBP method (that assumes a homogeneous acoustic medium) used the nominal acoustic and elastic properties of the skull in Table 5. Additionally, in the OBRM, the regularization parameter was determined through parameter scanning and empirically chosen due to the lack of ground truth. Reference images were obtained by reconstructing high-quality images of the cortical targets using the UBP method from experimental pressure data collected without the skull present.

In addition to visual assessment, quantitative evaluation was performed using the Pearson correlation coefficient (PCC) and Dice similarity coefficient (DICE). A challenge in this quantitative assessment was the presence of small registration errors and intensity differences between the reference images acquired without skull and those acquired with the skull present. These intensity differences arose due to variations in the optical fluence field. To mitigate these practical limitations of the experimental setup, the reference image and the corresponding reconstructed images in the presence of the skull were co-registered based on OBRM results using the SimpleITK package [45]. Furthermore, scale-invariant image quality measures were preferred over conventional metrics, such as MSE or SSIM, to account for optical fluence variations. Specifically, PCC was computed as the covariance of

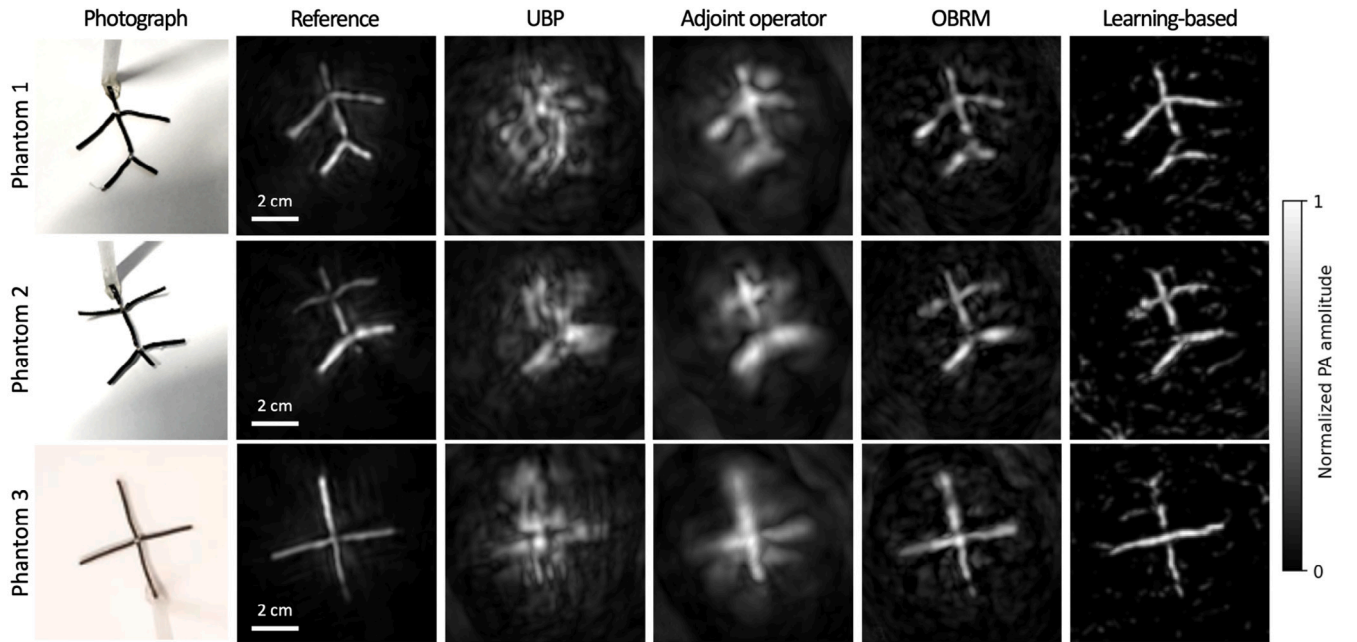


Fig. 12. Ex-vivo experimental results. Photographs of the three black wires included in the phantoms (first column) and images reconstructed from experimental pressure data measured without the skull employing the UBP method (second column) are provided as references. Images reconstructed from experimental pressure data measured with the skull using the UBP method, the adjoint operator, the OBRM, and the learning-based method (third to sixth columns) are presented for comparison. All reconstructed images are displayed as maximum intensity projections along the z -axis. The learning-based method effectively mitigated artifacts caused by skull-induced aberrations and produced reconstructed images visually comparable to those of the finely tuned, state-of-the-art OBRM, while offering a two-order magnitude acceleration.

pixel intensities between the reference image and the image acquired with the skull present, normalized by the product of the standard deviations of the pixel intensities in each image. A PCC value close to 1 indicates that the two images are nearly identical up to a scaling factor. The DICE metric was calculated following segmentation using Otsu's thresholding [46] to specifically assess the similarity of vascular structures while reducing the influence of intensity differences caused by variation in optical fluence distributions.

6.3. Experimental results

Fig. 12 presents the photographs of the cortical-vessel-mimicking structures included in the phantoms (first column), reference images obtained without the skull by use of the UBP method (second column), and images reconstructed with the skull present by use of the UBP method, the adjoint operator, the OBRM, and the learning-based method (third to sixth columns, respectively). Given the large amplitude differences in the initial pressure distributions between the skull-absent and skull-present cases, images are displayed in amplitude-normalized form.

As illustrated in Fig. 12, the images reconstructed using the UBP method (third column) exhibited significant degradation due to skull-induced aberrations, which were partially mitigated in preliminary images estimated using the adjoint operator (fourth column). Both the OBRM with fine-tuned skull properties (fifth column) and the learning-based method (sixth column) effectively reduced these aberration-induced artifacts. Notably, the learning-based method, despite being trained exclusively on simulated data, achieved a level of deaberration comparable to that of the OBRM. The quantitative metric results, presented in Table 6, further confirm that the learning-based method achieved image quality comparable to the OBRM while significantly outperforming both the UBP method and the adjoint operator.

In addition, the learning-based method yielded a two-order of magnitude reduction in reconstruction time compared to the OBRM. The

Table 6

Quantitative Metric Results in the Ex-vivo Study.

Method	Phantom 1		Phantom 2		Phantom 3	
	PCC	DICE	PCC	DICE	PCC	DICE
UBP	0.33	0.01	0.43	0.01	0.28	0.02
Adj	0.46	0.08	0.48	0.20	0.34	0.12
OBRM	0.66	0.59	0.69	0.54	0.60	0.54
Learning-based	0.66	0.62	0.67	0.51	0.59	0.54

Table 7

Computation Time Comparison for the Ex-vivo Study.

Method	Computation time		
	Prelim. Est.	Deaberration	Total
Learning-based method	0.167 h	<1 s	0.167 h
	Per iteration	# of iterations	Total
OBRM	0.375 h	80	30 h

computation times for producing the results in Fig. 12 using the computational hardware specified in the virtual imaging studies are summarized in Table 7. Establishing the training dataset (500 data pairs) in this study required approximately 188 h, with a total training time of approximately 8.7 h. Further discussion is provided in Section 7.

7. Discussion

In this study, a learning-based image reconstruction method was explored for 3D transcranial PACT. The method significantly reduces computational time and demonstrates robustness to inaccuracies in the assumed skull properties. The method was systematically evaluated through virtual imaging studies using stochastic 3D numerical head phantoms. In Study 1, the results showed that the learning-based method performed comparably to a baseline OBRM when the skull properties were accurately known. In Study 2, the method achieved

accurate image reconstructions when inaccurate skull property values were used, a scenario where OBRMs (model-based reconstruction methods) often fail due to model mismatches (Fig. 6). In Studies 3 and 4, the learning-based method also demonstrated robust generalizability to variations in the viscoelastic property statistics and optical properties of the scalp vasculature assumed in the training dataset.

Importantly, the method effectively mitigated skull-induced image artifacts in experimental studies (Fig. 12), despite being trained solely on computer-simulated data. This highlights the potential of the method for practical applications. To the best of our knowledge, this study represents the first experimental validation of a learned image reconstruction method for 3D transcranial PACT.

The learning-based method demonstrated a two-order of magnitude reduction in reconstruction times as compared to the OBRM. This is particularly important for practical applications of 3D transcranial PACT, such as functional and dynamic imaging, which require the rapid computation of numerous images. Although preparing the training dataset and training the learning-based method are necessary, these are one-time computational tasks and do not need to be repeated for a given subject and imaging configuration.

The proposed method is subject-specific, designed to fully utilize prior information about the subject (e.g., adjunct CT data) for effective aberration mitigation. However, this subject-specific nature requires establishing a training dataset and training the method for each subject. Fortunately, generating a subject-specific training set is relatively inexpensive, which requires computational resources comparable to performing 6–10 model-based reconstructions in numerical and validation studies. In future studies, the approach could potentially be extended to accommodate stochastic variations in skull and scalp anatomy.

One limitation of the synthetic PACT data used for training the learning-based method is the reduced-dimensional skull model, which assumes piecewise constant properties and may therefore not fully capture the heterogeneity of the viscoelastic properties within each skull plate. This limitation could potentially be addressed by using the adjunct CT data along with nominal average skull properties to estimate more detailed spatial distributions of the skull properties [47]. The extent to which such refinements in skull modeling may improve reconstructed image quality remains a topic for future investigation. A systematic investigation of the impact of skull registration errors on image reconstruction accuracy represents another important topic for future investigation.

The results from virtual imaging Study 4 demonstrated the generalizability of the learning-based method with respect to differences in the scalp optical properties in the training and testing data. This is relevant to practical applications, where variations in the optical absorption coefficient of the scalp vasculature may be present even across different scans of the same *in-vivo* subject. While the results demonstrated that the proposed method is robust to variations of $\pm 10\%$ in the optical absorption coefficient of the scalp vessels (Fig. 10), future work may consider more realistic NHP in which optical properties are spatially varying and randomly sampled within physiological ranges for each tissue type.

Future research should also include investigating the proposed method via more complex *ex-vivo* studies and extending it to *in-vivo* applications. As a next step, physical phantoms incorporating tissues both inside and outside the skull would provide a more realistic representation of the human head. In addition, using tubes filled with blood at varying oxygen saturation levels would allow for a thorough investigation of the proposed method's feasibility for functional imaging.

8. Conclusion

In this study, a learning-based image reconstruction method for 3D transcranial PACT was investigated, demonstrating its capability to compensate for image artifacts caused by skull-induced aberrations.

The proposed method effectively mitigated the impact of modeling errors when the skull's viscoelastic properties were not precisely known and achieved a two-order magnitude acceleration in image reconstruction time compared to conventional model-based approaches. To the best of our knowledge, this study provides the first experimental validation of a learned image reconstruction method for 3D transcranial PACT. This study marks a significant advancement in transcranial PACT and may facilitate the clinical translation of this emerging modality.

CRedit authorship contribution statement

Hsuan-Kai Huang: Writing – original draft, Software, Investigation, Formal analysis. **Joseph Kuo:** Software, Investigation. **Yang Zhang:** Investigation. **Yousuf Aborahama:** Investigation. **Manxiu Cui:** Investigation. **Karteekeya Sastry:** Investigation. **Seonyeong Park:** Writing – review & editing, Investigation. **Umberto Villa:** Writing – review & editing, Supervision, Methodology, Conceptualization. **Lihong V. Wang:** Writing – review & editing, Supervision, Conceptualization. **Mark A. Anastasio:** Writing – review & editing, Supervision, Methodology, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Lihong V. Wang reports a relationship with Microphotoacoustics, Inc. that includes: board membership. Lihong V. Wang reports a relationship with CalPACT, LLC that includes: board membership. Lihong V. Wang reports a relationship with Union Photoacoustic Technologies, Ltd. that includes: board membership. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research used the Delta advanced computing and data resource which is supported by the National Science Foundation, United States (award OAC 2005572) and the State of Illinois. Delta is a joint effort of the University of Illinois Urbana-Champaign and its National Center for Supercomputing Applications. This work was supported in part by the National Institutes of Health, United States under Awards R01 EB031585 and R01 EB034261. This work was also supported in part by National Institutes of Health, United States grants U01 EB029823 (BRAIN Initiative), R35 CA220436 (Outstanding Investigator Award), and R01 CA282505 as well as grant number 2024-337784 from the Chan Zuckerberg Initiative DAF, an advised fund of the Silicon Valley Community Foundation, United States. L.W. has a financial interest in Microphotoacoustics, Inc., CalPACT, LLC, and Union Photoacoustic Technologies, Ltd., which, however, did not support this work.

Data availability

Data will be made available on request.

References

- [1] L.V. Wang, S. Hu, Photoacoustic Tomography: In vivo imaging from organelles to organs, *Science* 335 (6075) (2012) 1458–1462, Publisher: American Association for the Advancement of Science Section: Review.
- [2] X. Wang, Y. Pang, G. Ku, X. Xie, G. Stoica, L.V. Wang, Noninvasive laser-induced photoacoustic tomography for structural and functional in vivo imaging of the brain, *Nature Biotechnol.* 21 (7) (2003) 803–806, Number: 7 Publisher: Nature Publishing Group.
- [3] J. Yao, J. Xia, K.I. Maslov, M. Nasiravanaki, V. Tsytarev, A.V. Demchenko, L.V. Wang, Noninvasive photoacoustic computed tomography of mouse brain metabolism in vivo, *NeuroImage* 64 (2013) 257–266.

- [4] M. Nasiravanaki, J. Xia, H. Wan, A.Q. Bauer, J.P. Culver, L.V. Wang, High-resolution photoacoustic tomography of resting-state functional connectivity in the mouse brain, *Proc. Natl. Acad. Sci.* 111 (1) (2014) 21–26, Publisher: National Academy of Sciences Section: Physical Sciences.
- [5] S. Na, J.J. Russin, L. Lin, X. Yuan, P. Hu, K.B. Jann, L. Yan, K. Maslov, J. Shi, D.J. Wang, C.Y. Liu, L.V. Wang, Massively parallel functional photoacoustic computed tomography of the human brain, *Nat. Biomed. Eng.* 6 (5) (2022) 584–592, Number: 5 Publisher: Nature Publishing Group.
- [6] R.W. Schoonover, M.A. Anastasio, Compensation of shear waves in photoacoustic tomography with layered acoustic media, *J. Opt. Soc. Amer. A* 28 (10) (2011) 2091–2099, Publisher: Optica Publishing Group.
- [7] J. Poudel, S. Na, L.V. Wang, M.A. Anastasio, Iterative image reconstruction in transcranial photoacoustic tomography based on the elastic wave equation, *Phys. Med. Biol.* 65 (5) (2020) 055009, Publisher: IOP Publishing.
- [8] J. Poudel, T.P. Matthews, K. Mitsuhashi, A. Garcia-Urbe, L.V. Wang, M.A. Anastasio, Iterative image reconstruction in elastic inhomogeneous media with application to transcranial photoacoustic tomography, in: *Medical Imaging 2017: Ultrasonic Imaging and Tomography*, Vol. 10139, SPIE, 2017, pp. 79–85.
- [9] J. Poudel, M.A. Anastasio, Joint reconstruction of initial pressure distribution and spatial distribution of acoustic properties of elastic media with application to transcranial photoacoustic tomography, *Inverse Problems* 36 (12) (2020) 124007, Publisher: IOP Publishing.
- [10] J. Poudel, Y. Lou, M.A. Anastasio, A survey of computational frameworks for solving the acoustic inverse problem in three-dimensional photoacoustic computed tomography, *Phys. Med. Biol.* 64 (14) (2019) 14TR01, Publisher: IOP Publishing.
- [11] J. Kuo, Advancing Photoacoustic Neuroimaging Through Deep Learning (Ph.D. thesis), University of Illinois at Urbana-Champaign, 2022.
- [12] F. Zhang, J. Zhang, Y. Shen, Z. Gao, C. Yang, M. Liang, F. Gao, L. Liu, H. Zhao, F. Gao, Photoacoustic digital brain and deep-learning-assisted image reconstruction, *Photoacoustics* 31 (2023) 100517.
- [13] K. Mitsuhashi, J. Poudel, T.P. Matthews, A. Garcia-Urbe, L.V. Wang, M.A. Anastasio, A forward-adjoint operator pair based on the elastic wave equation for use in transcranial photoacoustic computed tomography, *SIAM J. Imaging Sci.* 10 (4) (2017) 2022–2048, Publisher: Society for Industrial and Applied Mathematics.
- [14] H.-K. Huang, J. Kuo, S. Park, U. Villa, L.V. Wang, M.A. Anastasio, A learning-based image reconstruction method for skull-induced aberration compensation in transcranial photoacoustic computed tomography, in: *Photons Plus Ultrasound: Imaging and Sensing 2024*, Vol. 12842, SPIE, 2024, pp. 56–63.
- [15] Z. Alterman, F.C. Karal Jr., Propagation of elastic waves in layered media by finite difference methods, *Bull. Seismol. Soc. Am.* 58 (1) (1968) 367–398.
- [16] J. Virieux, P-SV wave propagation in heterogeneous media; velocity-stress finite-difference method, *Geophysics* 51 (4) (1986) 889–901.
- [17] B. Bolt, *Seismology: Surface Waves and Earth Oscillations*, Elsevier, 2012, Google-Books-ID: X7jLzYY7LBUC.
- [18] R. Madariaga, K. Olsen, R. Archuleta, Modeling dynamic rupture in a 3D earthquake fault model, *Bull. Seismol. Soc. Am.* 88 (5) (1998) 1182–1197.
- [19] J. Poudel, Three-dimensional Image Reconstruction in Transcranial Photoacoustic Computed Tomography (McKelvey School of Engineering Theses & Dissertations), 2019.
- [20] P. Moczo, J. Kristek, M. Galis, P. Pazak, M. Balazovjech, The finite-difference and finite-element modeling of seismic wave propagation and earthquake motion, *Acta Phys. Slovaca. Rev. Tutor.* 57 (2) (2007).
- [21] J.L.B. Robertson, B.T. Cox, J. Jaros, B.E. Treeby, Accurate simulation of transcranial ultrasound propagation for ultrasonic neuromodulation and stimulation, *J. Acoust. Soc. Am.* 141 (3) (2017) 1726–1738.
- [22] J. Robertson, J. Urban, J. Stitzel, B.E. Treeby, The effects of image homogenisation on simulated transcranial ultrasound propagation, *Phys. Med. Biol.* 63 (14) (2018) 145014, Publisher: IOP Publishing.
- [23] M. Anastasio, J. Zhang, E. Sidky, Y. Zou, D. Xia, X. Pan, Feasibility of half-data image reconstruction in 3-D reflectivity tomography with a spherical aperture, *IEEE Trans. Med. Imaging* 24 (9) (2005) 1100–1112, Conference Name: IEEE Transactions on Medical Imaging.
- [24] A. Beck, *Introduction to Nonlinear Optimization: Theory, Algorithms, and Applications with MATLAB*, SIAM, 2014.
- [25] A. Beck, M. Teboulle, Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems, *IEEE Trans. Image Process.* 18 (11) (2009) 2419–2434, Conference Name: IEEE Transactions on Image Processing.
- [26] A. Beck, M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, *SIAM J. Imaging Sci.* 2 (1) (2010) 183–202, Publisher: Society for Industrial and Applied Mathematics.
- [27] Y.E. Nesterov, A method of solving a convex programming problem with convergence rate $O(k^{-2})$, in: *Doklady Akademii Nauk*, vol. 269, Russian Academy of Sciences, 1983, pp. 543–547, Issue: 3.
- [28] E. Deining-Czermak, A. Euler, S. Franckenberg, T. Finkenstaedt, C. Villefort, D. Gascho, R. Guggenberger, Evaluation of ultrashort echo-time (UTE) and fast-field-echo (FRACTURE) sequences for skull bone visualization and fracture detection – A postmortem study, *J. Neuroradiol.* 49 (3) (2022) 237–243.
- [29] O. Ronneberger, P. Fischer, T. Brox, U-Net: Convolutional networks for biomedical image segmentation, 2015, arXiv:1505.04597.
- [30] O. Oktay, J. Schlemper, L.L. Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N.Y. Hammerla, B. Kainz, B. Glocker, D. Rueckert, Attention U-Net: Learning where to look for the pancreas, 2018, arXiv:1804.03999 [cs].
- [31] S. Ioffe, C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shifts, 2015, arXiv:1502.03167.
- [32] Y. Wu, K. He, Group normalization, 2018, pp. 3–19.
- [33] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, 2015, <http://dx.doi.org/10.48550/arXiv.1409.1556>, arXiv: 1409.1556 [cs].
- [34] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, 2016, pp. 770–778.
- [35] Y. Aborahama, K. Sastry, M. Cui, Y. Zhang, Y. Luo, R. Cao, L.V. Wang, De-aberration for transcranial photoacoustic computed tomography through an adult human skull, 2024, arXiv:2404.05937.
- [36] G.D.M. Talou, S. Safaei, P.J. Hunter, P.J. Blanco, Adaptive constrained constructive optimisation for complex vascularisation processes, *Sci. Rep.* 11 (1) (2021) 6180, Number: 1 Publisher: Nature Publishing Group.
- [37] Q. Fang, D.A. Boas, Monte Carlo simulation of photon migration in 3D turbid media accelerated by graphics processing units, *Opt. Express* 17 (22) (2009) 20178–20190, Publisher: Optica Publishing Group.
- [38] M. Louboutin, M. Lange, F. Luporini, N. Kukreja, P.A. Witte, F.J. Herrmann, P. Velsko, G.J. Gorman, Devito (v3.1.0): an embedded domain-specific language for finite differences and geophysical exploration, *Geosci. Model. Dev.* 12 (3) (2019) 1165–1187.
- [39] F. Luporini, M. Louboutin, M. Lange, N. Kukreja, P. Witte, J. Hückelheim, C. Yount, P.H.J. Kelly, F.J. Herrmann, G.J. Gorman, Architecture and performance of Devito, a system for automated stencil computation, *ACM Trans. Math. Software* 46 (1) (2020).
- [40] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, PyTorch: An imperative style, high-performance deep learning library, 2019, <http://dx.doi.org/10.48550/arXiv.1912.01703>, arXiv:1912.01703.
- [41] I. Loshchilov, F. Hutter, Decoupled weight decay regularization, 2019, <http://dx.doi.org/10.48550/arXiv.1711.05101>, arXiv:1711.05101 [cs, math].
- [42] J. Ahrens, B. Geveci, C. Law, ParaView: An end-user tool for large data visualization, in: C.D. Hansen, C.R. Johnson (Eds.), *Visualization Handbook*, Elsevier Butterworth-Heinemann, Burlington, MA, USA, 2005, pp. 717–731, [doi:10.1016/B978-012387582-2/50038-1].
- [43] P.L. Combettes, J.-C. Pesquet, Proximal splitting methods in signal processing, in: H.H. Bauschke, R.S. Burachik, P.L. Combettes, V. Elser, D.R. Luke, H. Wolkowicz (Eds.), *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, Springer, New York, NY, 2011, pp. 185–212.
- [44] S. Na, L.V. Wang, Photoacoustic computed tomography for functional human brain imaging [Invited], *Biomed. Opt. Express* 12 (7) (2021) 4056–4083, Publisher: Optica Publishing Group.
- [45] R. Beare, B. Lowekamp, Z. Yaniv, Image segmentation, registration and characterization in R with SimpleITK, *J. Stat. Softw.* 86 (2018) 1–35.
- [46] N. Otsu, A threshold selection method from gray-level histograms, *IEEE Trans. Syst. Man Cybern.* 9 (1) (1979) 62–66, Conference Name: IEEE Transactions on Systems, Man, and Cybernetics.
- [47] J.-F. Aubry, M. Tanter, M. Pernot, J.-L. Thomas, M. Fink, Experimental demonstration of noninvasive transskull adaptive focusing based on prior computed tomography scans, *J. Acoust. Soc. Am.* 113 (1) (2003) 84–93.



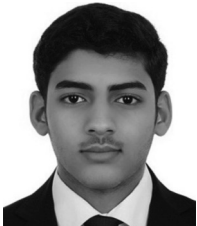
Hsuan-Kai Huang received the B.S. degree in Physics from the National Taiwan University, Taiwan, in 2018, and is a Ph.D. candidate in Electrical and Computer Engineering from the University of Illinois at Urbana Champaign, Urbana, USA. His research interests lie in transcranial photoacoustic computed tomography, machine learning, image reconstruction, and inverse problem.



Joseph Kuo received his B.S. in Electrical Engineering and M.S. in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign in 2020 and 2022, respectively. His research interests include transcranial photoacoustic computed tomography, machine learning, and image reconstruction. Currently, Joseph works at Intel, where he specializes in performance projections and optimizations for deep learning, contributing to advancements in computational efficiency.



Yang Zhang is a Postdoctoral Research Associate at the Caltech Optical Imaging Laboratory, Chreng Department of Medical Engineering, California Institute of Technology (Caltech). He received his Ph.D. from the University of Hong Kong (2018) and M.S. from Peking University (2015), respectively. His research interest is developing optical and ultrasonic technologies for disease screening, cancer diagnosis, and brain function study.



Yousef Aborahama received the B.Sc. degree (summa cum laude) in electrical engineering from the American University of Sharjah, Sharjah, UAE, in 2017. He is currently pursuing the Ph.D. degree in electrical engineering with the University of Toronto. His research interests are in the areas of wireless communications, microwave imaging, computational electromagnetics, and applied mathematics.



Manxiu Cui received his Bachelor degree from the Department of Electronic Engineering, Tsinghua University, Beijing, China in 2020. He is currently a graduate student in Caltech Optical Imaging Laboratory at California Institute of Technology, Pasadena, CA, USA. His research interest is photoacoustic imaging and wavefront shaping.



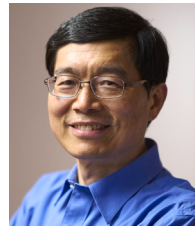
Karteekeya Sastry is a graduate student in the Caltech Optical Imaging Laboratory at the California Institute of Technology. He received his B.Tech and M.Tech degrees in electrical engineering from the Indian Institute of Technology, Madras, India in 2020. His research interests include image reconstruction and signal processing methods in photoacoustic tomography and other related imaging modalities.



Seonyeong Park received the B.S. degree in Electronics, Computer, and Telecommunication Engineering and the M.S. degree in Information and Communications Engineering from Pukyong National University, Busan, Korea, in 2011 and 2013, respectively. She earned the Ph.D. degree in Electrical and Computer Engineering from Virginia Commonwealth University, Richmond, VA, USA, in 2017. She is currently a Research Assistant Professor in the Department of Bioengineering at the University of Illinois Urbana-Champaign, Urbana, IL, USA. Her research interests include photoacoustic computed tomography, numerical modeling of biomedical objects, image reconstruction, and inverse problems.



Umberto Villa received the B.S. and M.S. degrees in Mathematical Engineering from the Politecnico di Milano, Milan, Italy, in 2005 and 2007, respectively, and the Ph.D. degree in Mathematics from Emory University, Atlanta, GA, USA, in 2012. He is a Research Associate Professor at the Oden Institute for Computational Engineering and Science, The University of Texas at Austin, Austin, TX, USA. His research interests lie in the computational and mathematical aspects of large-scale inverse problems, imaging science, and uncertainty quantification.



Lihong Wang edited the first book on photoacoustic tomography. His book entitled "Biomedical Optics: Principles and Imaging," one of the first textbooks in the field, won the 2010 Joseph W. Goodman Book Writing Award. He has published 610 peer-reviewed journal articles and delivered 620 keynote/plenary/invited talks. His Google Scholar h-index and citations have reached 164 and 113K, #1 most cited in optics according to Stanford/Elsevier. His laboratory was the first to report functional photoacoustic tomography, 3D photoacoustic microscopy, photoacoustic endoscopy, photoacoustic reporter gene imaging, the universal photoacoustic reconstruction algorithm, and CUP (world's fastest camera). He chairs the annual conference on Photons plus Ultrasound, the largest conference at Photonics West. He was the Editor-in-Chief of the Journal of Biomedical Optics. He received the NIH Director's Pioneer, NIH Director's Transformative Research, and NIH/NCI Outstanding Investigator awards. He also received the OSA C.E.K. Mees Medal, IEEE Technical Achievement Award, IEEE Biomedical Engineering Award, SPIE Britton Chance Biomedical Optics Award, IPPA Senior Prize and OSA Michael S. Feld Biophotonics Award. He is a Fellow of the AAAS, AIMBE, Electromagnetics Academy, IAMBE, IEEE, NAI, OSA, SPIE as well as a Foreign Fellow of COS. An honorary doctorate was conferred on him by Lund University, Sweden. He was inducted into the National Academy of Engineering.



Mark A. Anastasio received the Ph.D. degree from The University of Chicago, Chicago, IL, USA, in 2001. He is currently a Donald Biggar Willett Professor of Engineering and the Head of the Department of Bioengineering, University of Illinois at Urbana-Champaign, Urbana, IL, USA. His research broadly addresses computational image science, inverse problems in imaging, and machine learning for imaging applications. He has contributed broadly to emerging biomedical imaging technologies that include photoacoustic computed tomography, ultrasound computed tomography, and X-ray phase-contrast imaging. His work has been supported by numerous NIH grants and he served for two years as the Chair of NIH EITA Study Section. Dr. Anastasio is a fellow of the Society of Photo-Optical Instrumentation Engineers (SPIE), American Institute for Medical and Biological Engineering (AIMBE), the International Academy of Medical and Biological Engineering (IAMBE), and the Institute of Electrical and Electronics Engineers (IEEE).