



OPEN

A scalable open-source MATLAB toolbox for reconstruction and analysis of multispectral optoacoustic tomography data

Devin O'Kelly^{1,2}, James Campbell III¹, Jeni L. Gerberich¹, Paniz Karbasi², Venkat Malladi³, Andrew Jamieson³, Liqiang Wang² & Ralph P. Mason¹✉

Multispectral photoacoustic tomography enables the resolution of spectral components of a tissue or sample at high spatiotemporal resolution. With the availability of commercial instruments, the acquisition of data using this modality has become consistent and standardized. However, the analysis of such data is often hampered by opaque processing algorithms, which are challenging to verify and validate from a user perspective. Furthermore, such tools are inflexible, often locking users into a restricted set of processing motifs, which may not be able to accommodate the demands of diverse experiments. To address these needs, we have developed a Reconstruction, Analysis, and Filtering Toolbox to support the analysis of photoacoustic imaging data. The toolbox includes several algorithms to improve the overall quantification of photoacoustic imaging, including non-negative constraints and multispectral filters. We demonstrate various use cases, including dynamic imaging challenges and quantification of drug effect, and describe the ability of the toolbox to be parallelized on a high performance computing cluster.

Multispectral optoacoustic tomography (MSOT) is a relatively new imaging modality, which combines optical contrast and ultrasonic resolution in order to provide a highly parametric view of an imaged sample. Due to the high spatiotemporal resolution, MSOT captures vast quantities of information, creating potential challenges: How does one effectively and efficiently analyze such data? Moreover, how does one ensure that the analysis is done in a manner, which is transparent and verifiable, and thus trustworthy, while allowing the flexibility to adjust the manner of processing to accommodate the differential needs of a wide variety of experiments and data acquisition conditions?

To address these issues, we have created an MSOT Reconstruction, Analysis, and Filtering Toolbox (RAFT), enabling users of various technical skill levels to benefit from the rich information present in MSOT data, and to share and compare analyses between sites in a verifiable manner. The package is available as a MATLAB toolbox and a standalone command-line application, with plans to deploy to Docker containers for site- and platform-independence, potential deployment to cloud computing resources, as well as to preserve archival code for execution. The toolbox provides a basis for users to add their own reconstruction and unmixing algorithms, but may also be run in a 'black-box' mode, driven only through external configuration files, thus providing greater repeatability. Scalability is addressed by parallelization within MATLAB, allowing computationally intensive reconstructions to be distributed across several computational workers. A Nextflow wrapper around the MATLAB systems allows multiple pipelines to be run on a single computational cluster simultaneously.

Background

MSOT may be thought of as optically-encoded ultrasound; an object under brief, intense illumination absorbs some of the illuminating energy, converting some of that energy into heat^{1,2}. This heat induces transient thermoelastic expansion, creating a pressure wave, which travels outward from the point of absorption. If this absorption and thermal conversion process occurs in a short enough period of time, the pressure wave is temporally confined, creating a compact wavefront, which can be detected using ultrasonic transducers. The reconstruction

¹Department of Radiology, University of Texas Southwestern Medical Center, 5323 Harry Hines Blvd., Dallas, TX 75390-9058, USA. ²BioHPC, University of Texas Southwestern Medical Center, Dallas, TX, USA. ³Lyda Hill Department of Bioinformatics, University of Texas Southwestern Medical Center, Dallas, TX, USA. ✉email: ralph.mason@UTSouthwestern.edu

of the original pressure image then provides a measure of energy deposition by the original illumination, which may be related across multiple illumination wavelengths to yield an overall spectral image. Knowledge of the endmembers present allows the spectral image to be unmixed into its corresponding endmember images.

The promise of MSOT, given recent advancements in laser tuning controls, data acquisition bandwidth, and ultrasound transducer design, is that it can provide optical contrast with ultrasound resolution. Indeed, the modality is highly scalable across various temporal and spatial regimes: It may be used for super-resolution imaging at the scale of tens of nanometers^{3–5}, is becoming popular for preclinical and small animal investigations^{6–11} and has been deployed for clinical usage on human subjects^{12–18}. The rate of imaging is fundamentally limited by two parameters: The signal to noise ratio achievable by a given photoacoustic imaging technology, and the exposure limits defined by guiding agencies. In practice, the field assumes a maximum energy deposition of 20 mJ/cm² at skin surface, and therefore many low-energy laser shots may be substituted for few high-energy laser shots^{19,20}.

Beyond the nature of the acquisition devices themselves, there are numerous methods by which photoacoustic images may be reconstructed; these range from direct inversion algorithms analogous to the ‘delay-and-sum’ approaches of ultrasound²¹, to analytical inversions valid under specialized geometries^{22,23}, to model-based approaches analogous to those used in CT and MRI^{22,24–29}. Still further methods use the time-reversal symmetry of the governing equations and numerical simulation to determine the original photoacoustic energy distribution^{30,31}. These all operate under the governing photoacoustic equations. Under pulsed laser light and the assumption of ideal point transducers, the photoacoustic equations can be written as an optical component (1) and an acoustic component (2)³²:

$$p_0(\vec{r}, \vec{\lambda}) = p(t = 0, \vec{r}, \vec{\lambda}) = H(\vec{r}, \vec{\lambda}) = \Gamma(\vec{r})\phi(\vec{r}, \vec{\lambda})\mu_a(\vec{r}, \vec{\lambda}) = \Gamma(\vec{r})\phi(\vec{r}, \vec{\lambda}) \sum_{i=1}^{N_c} C_i(\vec{r})\varepsilon_i(\vec{\lambda}) \quad (1)$$

$$p_d(t, \vec{r}_d) = \frac{\partial}{\partial t} \left[\frac{t}{4\pi} \iint_{|\vec{r}_d - \vec{r}| = v_s t} p_0(\vec{r}) d\Omega \right] \quad (2)$$

Here, p denotes pressure values, assumed to instantaneously take on the values $p_0(\vec{r})$ at each location \vec{r} throughout the imaging region at time $t = 0$. $H(\vec{r})$ denotes the heating function, which is dependent on the Gruneisen parameter Γ describing the efficiency of conversion from light to heat at each point, the light fluence $\phi(\vec{r}, \vec{\lambda})$ defined throughout the imaging region for each illumination wavelength λ , and the absorption coefficient $\mu_a(\vec{r}, \vec{\lambda})$, which describes the conversion of light fluence to absorbed energy at each point and each wavelength. In turn, the absorption coefficient is described by the absorption spectra $\varepsilon_i(\vec{\lambda})$ of all endmembers present at each point, weighted by their concentration at that point $C_i(\vec{r})$ (Eq. 1).

Once light has been absorbed and converted into acoustic waves, these waves spread outward from the point of origin according to the wave equation with wave-fronts traveling at the speed of sound v_s in the medium, eventually reaching the detection transducers located at \vec{r}_d , each of which has a solid angle $d\Omega$ describing the contribution of each imaged point to the overall measured signal at time t (Eq. 2). This integration over the 2D detection surface is denoted by the double integral in Eq. 2.

Once data have been acquired using transducer geometry \vec{r}_d and sampled at \vec{t} , the time-dependent photoacoustic data must be reconstructed into the corresponding image. Thorough overviews of the state of the field have been compiled by various groups^{1,33–35}. Reconstruction may be accomplished through either direct or inverse methods; the former encompasses such approaches as backprojection algorithms, while the latter encompasses any approach using a model of image formation mapping from $H(\vec{r})$ to $p_d(\vec{t}, \vec{r}_d)$ and minimizing a cost function. For any given model or reconstruction approach, a variety of parameters and constraints may be applied to the process to enforce certain conditions such as non-negativity, or to regularize the reconstruction process and emphasize certain aspects of the data.

A general representation of the reconstruction process in discrete form is given by the optimization problem

$$\vec{I} = \underset{\vec{I}}{\operatorname{argmin}} \left\| M\vec{I} - \vec{d} \right\| + \lambda \left\| R(\vec{I}, \vec{d}) \right\| \quad (3)$$

where \vec{I} represents the reconstructed image, \vec{d} the acquired data, M the forward model mapping \vec{I} to \vec{d} , R a regularization function which is a function of \vec{I} and/or \vec{d} , and λ a regularization parameter which adjusts the influence of the regularization term in the value of the objective function. Reconstruction may be performed using any norm, though the L_2 Euclidean norm is most commonly used.

A plethora of pre- and post-processing approaches may be added to the analysis chain; reconstruction may be performed on the individual images, or jointly among several multispectral images simultaneously. Unmixing may be performed prior to, or after, reconstruction, and both reconstruction and unmixing may themselves be combined into a single operation to yield unmixed images directly from multispectral ultrasound data.

It is important that analyses be transparent, verifiable, and validatable. Many approaches may give ‘suitable’ results in a qualitative sense, but applying further quantitative analysis to these less than quantitative intermediate results can lead to spurious analyses and interpretation. It is thus key that assumptions be explicit for any analysis, and that the analytical provenance for a given analyzed datum be recorded. As a salient example, bandpass filtering of the recorded ultrasound signal and deconvolution of the transducer impulse response are common

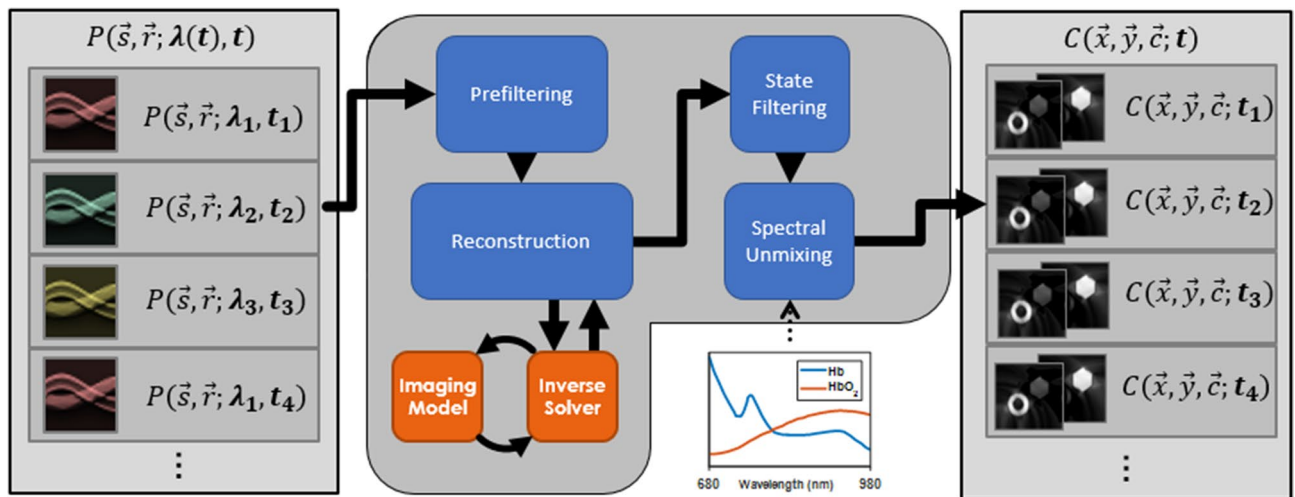


Figure 1. Example pipeline structure. A stream of individual single-wavelength photoacoustic data frames is transformed through a cascade of processing actions to yield a series of spectrally unmixed images at each point in time. The pipeline takes advantage of the known acquisition geometry of the system to perform reconstruction, and user-defined spectral endmembers to perform unmixing. Figure created using a combination of MATLAB and Microsoft PowerPoint.

preprocessing steps³⁶, but many deconvolution methods require an assumption of white noise, which is violated if bandpass filtering is performed prior to deconvolution.

It is also important that a given analysis be repeatable and consistent across datasets; indeed, if one has a dataset that benefitted from a particular analytical chain, then one would like to easily apply the same processing parameters to other datasets under similar acquisition conditions. However, given the diversity of possible acquisition settings, an analysis must be sufficiently flexible to accommodate the details of each particular dataset, such as the order of wavelengths sampled.

Numerous tools are available to address individual components of this process; K-wave is an excellent toolbox developed over the past several years by Treeby and Cox, largely intended to simulate pressure fields in ultrasound and photoacoustic imaging contexts³⁷. Toast++ and NIRFAST are toolboxes to simulate light transport efficiently^{38,39}, while MCML addresses light transport from a Monte-Carlo perspective⁴⁰. FIELD-II simulates acoustic sensitivity fields^{41,42}. Given the existence of such tools, it is key that they can be effectively leveraged into the new RAFT framework.

Lastly, for any amount of computational power or configurability, it is critical that the use of such tools be straightforward and consistent, and that any changes made to a codebase do not interfere with established analyses or reduce the overall quality of results. It is therefore key that there be means by which the framework is tested and updated automatically, or at least to a level of convenience enabling update and maintenance.

We address several of these problems and believe that this new tool will provide a foundation for the future development of photoacoustic imaging, and that further development of the RAFT will continue to expand its capabilities.

MSOT-RAFT structure. RAFT operates in a data-driven manner, directed by a configuration file which describes the parameters and methods used to perform a given analysis, and which acts as a record of how a given dataset is processed. Defaults are provided to enable immediate usage, but users may modify these defaults according to their preferences. For settings which are explicitly dependent on metadata, such as the number of samples acquired or the location of transducers, a metadata population step loads such information into a standard form, thereby generalizing the pipeline's application beyond a specific manufacturer's technology. Multiple extensions of the processing pipeline to different photoacoustic data are possible through modification with different loaders.

Data frames are loaded using a memory-mapped data interface, allowing large datasets to be handled without loading the entire dataset into memory at one time. This enables the processing to be performed even on workstations with minimal available memory. The action is handled by a loader, which is instantiated using the metadata information associated with the MSOT dataset, giving it a well-defined mapping between a single scalar index and the corresponding data frame. Each raw data frame is associated with acquisition metadata, such as the temperature of the water bath or laser energy, and is output to the next processing step.

Processing proceeds via the transformation of data frames, which are arrays of data associated with a particular, well-defined coordinate system. An example is the data frame acquired from a single laser pulse and the resulting acoustic acquisition. Data frames may also have additional, contextual data, such as the time of acquisition or excitation wavelength. Figure 1 illustrates a commonly-used processing topology, with an overall effect of transforming a series of single-wavelength photoacoustic data frames into a series of multi-component image frames.

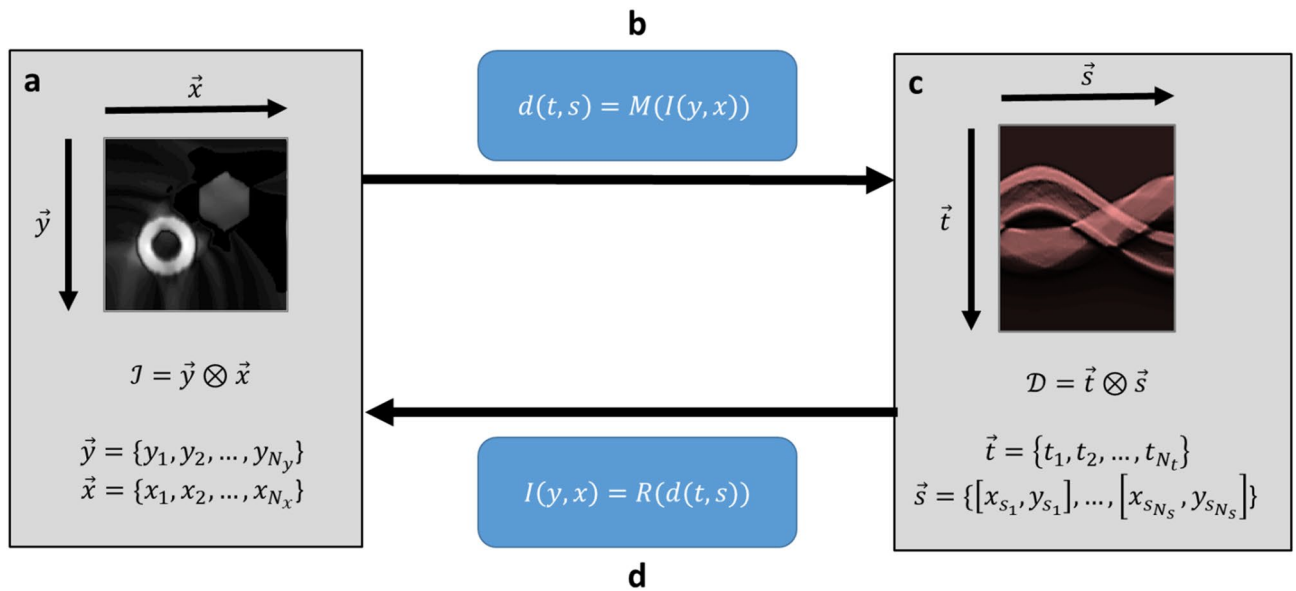


Figure 2. Image-data mappings. An image (a) consisting of some explicit values, is organized according to the coordinate system \mathcal{T} . (b) Through the transformation effected by the forward model operator M , one determines (c) the photoacoustic data in coordinate system \mathcal{D} . In practice, one acquires the data and (d) seeks to reconstruct the corresponding image through the action of some reconstruction operator R . Iterative schemes are often favored for the reconstruction process, and so R is implicitly dependent on the operator M . Figure created using a combination of MATLAB and Microsoft PowerPoint.

Coordinates explicitly describe the structure of the data contained within data frames. Coordinates may be vectors of scalars or vectors of vectors. An example of the former is the use of two scalar coordinates to describe the X and Y position of a given pixel vertex (Fig. 2a, x and y coordinates), while an example of the latter is an index of transducers, each of which has an associated X,Y position (Fig. 2c, \vec{s} coordinate). Transformations between these spaces are effected by the forward model M (Fig. 2b) and the inverse solution operator given in Eq. 3 (Fig. 2d).

Some processing methods require that the procedure maintain memory of its state, e.g., in the use of recursive filters or online processing. We therefore implemented the filters using the MATLAB System Object interface, which provides a convenient abstraction to describe such operator mappings, where the transformation may have some time-dependent internal structure. Other processing methods benefit from the assumption that each data frame within a dataset may be treated independently, such as in the case of reconstruction, and the use of an object-oriented design ensures that the processing of such datasets can be effectively parallelized.

Implementation. The most up-to-date version of the RAFT is available at <https://doi.org/10.5281/zenodo.4658279>, which may have been updated since the time of writing.

The RAFT is designed in a highly polymorphic manner, allowing different methods to be applied to each step, and for the overall pipeline topology to change. As an example, one can perform multispectral state estimation on the raw data prior to reconstruction, in contrast to the topology shown in Fig. 1, where the multispectral state estimation occurs after reconstruction. For succinctness, we will only be illustrating the pipeline topology shown in Fig. 1.

Following its loading into memory and calibration for laser energy variations, each raw data frame is pre-conditioned by subtracting the mean of each transducer's sampled time course, deconvolving the transducer impulse response using Wiener deconvolution⁴³, bandpass filtering the signal, correcting for wavelength-dependent water absorption with the assumption of Beer's law attenuation⁴⁴, and interpolating the data frame into the coordinate system expected by the reconstruction solver.

Reconstruction proceeds by assuming the independence of each data frame, and is thus parallelized across an entire dataset through the use of a number of distributed workers, each performing the same processing on a distinct region of the overall dataset. During initialization, the reconstruction system creates a model operator, which is used during operation to reconstruct each frame into the target image coordinate system. Each frame is then written to disk until processing completes.

Spectral unmixing is accomplished by providing a stream of single-wavelength data frames to a multispectral state filter, which estimates the true multispectral image at each point in time⁴⁵. This multispectral image estimate is then provided to another solver system which inverts the mixing model, derived from the known wavelength space and the assumed endmembers present.

To readily accommodate future methods, we established a consistent input configuration for processing steps. Reconstruction systems are constructed by pairing an inverse solver to a forward model: The forward model represents the mapping from some input space (here, an image) to some output space (here, a frame of

photoacoustic data), while the inverse solver takes the model as an argument and attempts to reduce some objective function subject to some input arguments. Models have a standardized initialization signature, requiring an input coordinate system, an output coordinate system, and a set of model-specific parameters, which, in the case of MSOT, necessarily includes the speed of sound of the medium.

Pipeline parallelization. MSOT data processing is computationally intensive, and the associated large problem sizes incur substantial processing time. In biological research, where cohorts of animals may be assessed multiple times over the course of a study, this processing time can result in prohibitively long experimental iterations, hindering effective development of methods. If there is no co-dependence among individual acquisitions, it is possible to parallelize computationally intensive steps, and particularly desirable to do so when such steps require a long time to process^{46,47}. We therefore implemented the toolbox using MATLAB's object-oriented functionalities to enable one processing system to be copied among an arbitrary number of parallel workers, and to demonstrate the benefits of parallelization in the case of reconstruction. When many datasets are to be processed using the same configuration, it is desirable to parallelize the processing at the scale of datasets. To this end, we implemented a Nextflow wrapper around the toolbox to illustrate a possible scenario of deploying the toolbox on a computational cluster. With a cluster scheduler, an arbitrary number of datasets can be processed simultaneously, enabling substantial horizontal scalability.

Comparison of processing methods. There is a practically infinite number of combinations and permutations of different reconstruction methods, preconditioning steps, solver settings, cost functions, and analyses, resulting in a highly complex optimization problem. By providing means to generate arbitrary amounts of test data in a parametric fashion, and by describing processing pipelines using well-defined recipe files, we enable optimization of the complex parametric landscape describing all possible processing pipelines. We illustrate the use of the pipeline to assess the effects of different processing approaches in providing an end analysis.

Methods

Results were processed using MATLAB 2017b (MathWorks, Natick, MA), though we include continuous integration testing intended to provide consistency between versions. All figures were created using a combination of MATLAB and Microsoft PowerPoint.

All animal work was conducted under animal protocol (APN #2018-102344-C) approved by the UT Southwestern Institutional Animal Care and Use Committee. All animal work was conducted in accordance with the UT Southwestern Institutional Animal Care and Use Committee guidelines as well as all superseding federal guidelines and in conformance with ARRIVE guidelines.

We verified the ability of the pipeline to perform reconstructions through the use of several testing schemes. First, we implemented analytical data generation, both single-wavelength and multispectral, in order to test the numerical properties of models and reconstructions. Data were generated consisting of random paraboloid absorbers across a field of view, along with the corresponding photoacoustic data, for a variety of pixel resolutions. For each model, 50 random images were generated with random numbers of sources, random sizes, and random locations at each of the chosen resolutions, and the correspondence of each model's forward data to the analytical data was quantified. The assumption of point detectors allows one to find the analytical signal expected at a given sampling location, and the assumption of linearity allows one to calculate the total signal from several non-overlapping sources. This provides validation of imaging models, by comparing the known image of sources against the model output. We use this process to demonstrate the potential use of the RAFT as a process optimization tool, informing selections of methods and parameters under different use-cases.

To test the performance of the different models and solvers at different imaging resolutions, we generated numeric data of paraboloid absorbers. The span of each paraboloid was constrained to lie in the field of view of the reconstructed image. Given a known set of paraboloid parameters, it is possible to construct the corresponding projection of the paraboloids onto the image space and the data space. Given the known source image, we compared the output of each model to the ground truth data. To compare solvers, given the output data, we compared the reconstructed image against the known ground truth image, creating an error image (Eq. 4). We quantified the mean bias (Eq. 5), average L_1 (Eq. 6) and L_2 norms (Eq. 7), scaled by the number of pixels in an image frame or reconstruction:

$$e(x, y) = I_{known}(x, y) - I_{recon}(x, y) \quad (4)$$

$$bias = \frac{\sum_{y \in \mathcal{Y}} \sum_{x \in \mathcal{X}} e(x, y)}{N_x N_y} \quad (5)$$

$$\|\tilde{e}\|_{L_1} = \frac{\sum_{y \in \mathcal{Y}} \sum_{x \in \mathcal{X}} |e(x, y)|}{N_x N_y} \quad (6)$$

$$\|\tilde{e}\|_{L_2} = \frac{\sqrt{\sum_{y \in \mathcal{Y}} \sum_{x \in \mathcal{X}} e(x, y)^2}}{N_x N_y} \quad (7)$$

The bias of the reconstruction signifies the tendency of the reconstruction process to over- or under-estimate the pixel values of the reconstructed image—bias values close to 0 indicate that the reconstructed values are on average equal to the true values. The L_1 and L_2 norms reflect the error in the reconstruction—lower values of each indicate that the solution converges to a more precise estimate of the true image values.

We additionally quantified the Structural Similarity Index (SSIM)⁴⁸, a measure of image similarity, as a scale-invariant figure of merit:

$$SSIM_{x\hat{x}} = \frac{(2\mu_x\mu_{\hat{x}} + C_1)(2\sigma_{x\hat{x}} + C_2)}{(\mu_x^2 + \mu_{\hat{x}}^2 + C_1)(\sigma_x^2 + \sigma_{\hat{x}}^2 + C_2)} \quad (8)$$

The average norms and the SSIM were calculated for the error images for each generated phantom image, and this process was repeated $N = 50$ times at each resolution tested ($N_x = N_y \in [30, 50, 100, 150, 200, 250, 300, 350, 400]$). For the solver comparisons, we additionally calculated the number of negative pixels in each image as well as the relative residual at convergence for each method, with the relative residual given by:

$$r_{rel} = \frac{\|M\bar{I} - \vec{d}\|}{\|\vec{d}\|} \quad (9)$$

To demonstrate application to biological data, we analyzed two preclinical imaging scenarios: The first (Dataset A) was a dynamic observation of a gas challenge followed by administration of a vascular disrupting agent (VDA). A male NOD-SCID mouse (Envigo) with a human PC3 prostate tumor xenograft implanted subcutaneously in the right aspect of the back was subjected to a gas breathing challenge, while continuously anesthetized with 2% isoflurane. Prior to imaging, the animal was shaved and depilated around the imaging region to avoid optical or acoustic interference from fur. Gas flow was maintained at 2L/min throughout the imaging session and mouse was allowed to equilibrate in the imaging chamber for at least 10 min before measurements commenced. Initial air breathing was switched to oxygen at 8 min, back to air at 16 min, and again to oxygen at 24 min. At 34 min, the animal was given an intraperitoneal injection of 120 mg/kg combretastatin A-4 phosphate (CA4P) *in situ*^{49–51}, and observed for a further 60 min. Imaging proceeded by sampling the wavelengths [715, 730, 760, 800, 830, 850] nm, with each wavelength oversampled 6 times, but not averaged. Overall, 52,091 frames of data were acquired.

The second scenario (Dataset B) examined a single gas challenge, wherein a female nude mouse (Envigo) was implanted with a human MDA-MB-231 breast tumor xenograft in the right dorsal aspect of the lower mammary fat pad. The gas challenge consisted of 11 min breathing air, 8 min breathing oxygen, followed by 10 min breathing air. Imaging proceeded by sampling the wavelengths [715, 730, 760, 800, 830, 850] nm, with each wavelength oversampled 2 times but not averaged. Overall, 11,184 frames were acquired.

We tested the scalability of the pipeline by performing reconstruction of two experimentally-acquired datasets of 52,091 frames (Dataset A) and 989 frames (Dataset C, acquired during a tissue-mimicking phantom experiment), using two distinct indirect methods and one direct method. Reconstruction was performed using either the Universal Backprojection (BP) algorithm²¹, chosen for its intrinsic universality, or the closely related direct interpolated model matrix inversion²⁸ (dIMMI) or curve-driven model matrix inversion²⁶ (CDMMI) models. The two models differ in how they discretize the photoacoustic imaging equations. The dIMMI model uses a piecewise-linear discretization with a defined number of points sampled along the wavefront, and uses bilinear interpolation to assign weights to nearby pixels. CDMMI, in contrast, assumes a spherical wavefront and exactly calculates the arc-length of the wavefront within each pixel. The solvers used for reconstruction were either MATLAB's built-in LSQR function with default parameters, or the non-negative accelerated projected conjugate gradient (nnAPCG) method⁵². The total processing time for reconstructing all frames of Dataset A was determined using $N = [4, 8, 12, 16, 20, 24]$ distributed workers, while $N = [1, 2, \dots, 24]$ distributed workers were applied to Dataset C.

We used the RAFT to execute two distinct processing pathways on each biological dataset, using the package's default parameter settings except where otherwise specified: The first, so-called 'unconstrained' method, used the dIMMI model and an LSQR solution of the imaging equations, and a sliding-window multispectral state estimation filter to produce a complete multispectral image at each point in time.

Unmixing was then performed using the pseudoinverse of the mixing equations:

$$I(\vec{r}, \vec{\lambda}) = \begin{pmatrix} \vec{I}_{\lambda_1}(\vec{r}) \\ \vdots \\ \vec{I}_{\lambda_{N_\lambda}}(\vec{r}) \end{pmatrix} = \begin{pmatrix} I_{\lambda_1, r_1} & \cdots & I_{\lambda_1, r_{N_r}} \\ \vdots & \ddots & \vdots \\ I_{\lambda_{N_\lambda}, r_1} & \cdots & I_{\lambda_{N_\lambda}, r_{N_r}} \end{pmatrix} \approx \mu_a(\vec{r}, \vec{\lambda}) = \sum_{i=1}^{N_c} C_i(\vec{r}) \varepsilon_i(\vec{\lambda}) = E\vec{C} \quad (10)$$

$$= \begin{pmatrix} \varepsilon_{\lambda_1, C_1} & \cdots & \varepsilon_{\lambda_1, C_{N_c}} \\ \vdots & \ddots & \vdots \\ \varepsilon_{\lambda_{N_\lambda}, C_1} & \cdots & \varepsilon_{\lambda_{N_\lambda}, C_{N_c}} \end{pmatrix} \begin{pmatrix} c_{C_1, r_1} & \cdots & c_{C_1, r_{N_r}} \\ \vdots & \ddots & \vdots \\ c_{C_{N_c}, r_1} & \cdots & c_{C_{N_c}, r_{N_r}} \end{pmatrix}$$

$$\hat{C}(\vec{r}, \vec{c}) \approx E^+ \vec{I}(\vec{r}, \vec{\lambda}) \quad (11)$$

Here, $\vec{I}_j(\vec{r})$ represents the spectral image at each wavelength λ_j and each pixel location \vec{r} . Variations in fluence are assumed to be negligible, and so the multispectral image is assumed to be approximately equivalent to the actual absorption image μ_a . The mixing matrix E is constructed by taking the molar absorption coefficients of each endmember C_i at each wavelength λ_j . The unmixed concentration image is thus derived by multiplying $I(\vec{r}, \vec{\lambda})$ on the left by the pseudoinverse E^+ , itself calculated by using MATLAB's *pinv* function. In this work, the endmembers were assumed to be only oxyhemoglobin and deoxyhemoglobin, with values derived from the literature⁵³.

The second, so-called “constrained” approach, used the CDMMI method and an nnAPCG solution of the imaging equations to provide a non-negativity constraint, and an $\alpha\beta$ Kalata filter⁴⁵ for the multispectral state estimation at each point in time to provide a kinematic constraint on the time-evolution of the signal. Unmixing was then performed using the nnAPCG solver applied to the mixing equations. The second imaging scenario was additionally processed using a second constrained approach, using an α Kalata filter. The α filter is very effective at reducing noise over time, but is unable to reliably track rapid dynamic changes in the underlying signal, causing a lagged response. The $\alpha\beta$ filter, by contrast, has lower noise-suppression properties but is much more able to follow signal dynamics. This is in essence due to the fact that the α filter assumes the underlying signal is static.

All approaches were preconditioned by first subtracting the mean from each transducer's pressure data, Wiener deconvolution of the impulse response function, bandpass filtering between 50 kHz and 7 MHz, and correction by the water attenuation coefficient for a given frame's acquisition wavelength, assuming a general path length of 3 cm. All images were reconstructed at 200×200 resolution and a 2.5 cm field of view unless otherwise specified. The effects of the combination of these steps can be seen in Supplementary Fig. 1.

We assessed the variation in quantitation by several approaches. For Dataset A, we quantified the indicator function $H^-(x)$ of negative pixels throughout, *i.e.*, the number of times each pixel in the reconstructed images attained a negative value throughout the imaging session. The occurrence of negative pixels creates interpretation difficulties, and it is desirable to have few or no negative pixels in a dataset. The unmixed values of hemoglobin ($[\text{Hb}]^{\text{MSOT}}$), oxyhemoglobin ($[\text{HbO}_2]^{\text{MSOT}}$), total hemoglobin ($[\text{Hb}_{\text{Tot}}]^{\text{MSOT}}$), and oxygen saturation ($\text{SO}_2^{\text{MSOT}}$) were compared between the two methods using a randomly-sampled binned Bland–Altman analysis with 100×100 bins, randomly sampling 10% of all pixels throughout the image time course. In Dataset A, the significance of the difference in each unmixed parameter before and after drug was quantified, using a two-population T-test and $N = 100$ frames between [19900:20000] frames and [51900:52000], corresponding to 10 s immediately before administration of the drug and 10 s immediately before the end of the imaging session. The T-score itself was used as the figure of merit. We additionally quantified the relative effect of the CA4P response by using the initial gas challenge as a measure of scale, relating the Cohen's D of the CA4P response to the Cohen's D of the transition from oxygen-air at 16 min. This was done based on the presumption that the initial gas challenge provides indications of the patency of vasculature, which should be related to the response to CA4P. Details of this analysis are presented in Supplementary Information.

The MDA-MB-231 gas challenge was assumed to have a rectangular input pO_2 waveform. We quantified the centered correlation of each pixel in each channel for each method against this known input function, and calculated the mean and variance of the correlation within tumor, spine, and background ROIs, as well as an ROI covering the whole animal cross-section. Statistical significance of differences between each method's correlation coefficient was quantified using Fisher's Z transformation on the raw correlation coefficients and a two-population Z-test. Significance was set at $\alpha = 0.01$, while strong significance was set at $\alpha = 0.0001$.

We additionally tested the ability of each of the processing approaches (Unconstrained, Constrained + $\alpha\beta$, Constrained + α) to provide results which were amenable to further analysis, in this case fitting a 7-parameter exponential model (Supplementary Methods) to the $[\text{HbO}_2]^{\text{MSOT}}$ time course in each pixel for each method.

Computation was performed using the UT Southwestern BioHPC computational resource; timing was calculated using computational nodes with between 256 and 384 GB of RAM and 24 parallel workers.

Results

When comparing CDMMI to DIMMI in terms of accurate prediction of model output data, both methods performed with statistical equivalence at high resolutions (Fig. 3). At lower resolutions, there was a notable increase in the modelling error of the CDMMI method. This may be attributed to the problem of aliasing; the DIMMI method allocates additional sampling points along the integral curve during model calculation, acting as an anti-aliasing filter, which reduces the incidence of high-frequency artifacts exacerbated by the differentiation step. However, the CDMMI model seems to indicate a trend towards statistically improved performance at higher resolutions.

Reconstruction performance contrasts the modelling results (Figs. 4 and 5); when images were reconstructed using the LSQR method (Fig. 4) and considering the L_1 and L_2 norms (Fig. 5d, e), the CDMMI method had statistically superior performance over DIMMI ($N < 250$), though at higher resolutions the difference was negligible. In contrast, when considering the relative residual (Fig. 4b), DIMMI provided superior reconstruction performance. Similar results were seen for the nnAPCG method (Fig. 5), though we note the minimum relative residual, L_1 norm, and L_2 norm were achieved at a lower resolution ($N = 150$). This discrepancy between the resolutions of minimum relative residual or norm potentially indicates a need for additional iterations at higher resolutions for the nnAPCG method. We note, however, that the presence of a minimum indicates that the RAFT itself may be optimized through the use of randomly-generated data to provide optimal reconstruction settings.

The choice of distinct reconstruction methods has a salient impact on the quality of the reconstruction, as well as the sensitivity of the method to solver parameters. The solution process in MSOT does not necessarily converge after infinite iterations due to ill conditioning of singular values, and so must be halted at an early stage, or a truncated singular value decomposition (T-SVD) be used as a regularization process. This explains the higher




Figure 3. Comparison of dIMMI and CDMMI models. The dIMMI model has statistically superior performance in accurately modelling the photoacoustic imaging process for low image resolutions (a, b) under both the L_1 and L_2 norms. At high resolutions, CDMMI trended towards superior performance over dIMMI, though does not consistently achieve statistical separation. (c) Both dIMMI and CDMMI have improved SSIM scores as a function of pixel resolution; CDMMI achieves slightly better performance at high resolutions. Figure created using a combination of MATLAB and Microsoft PowerPoint.

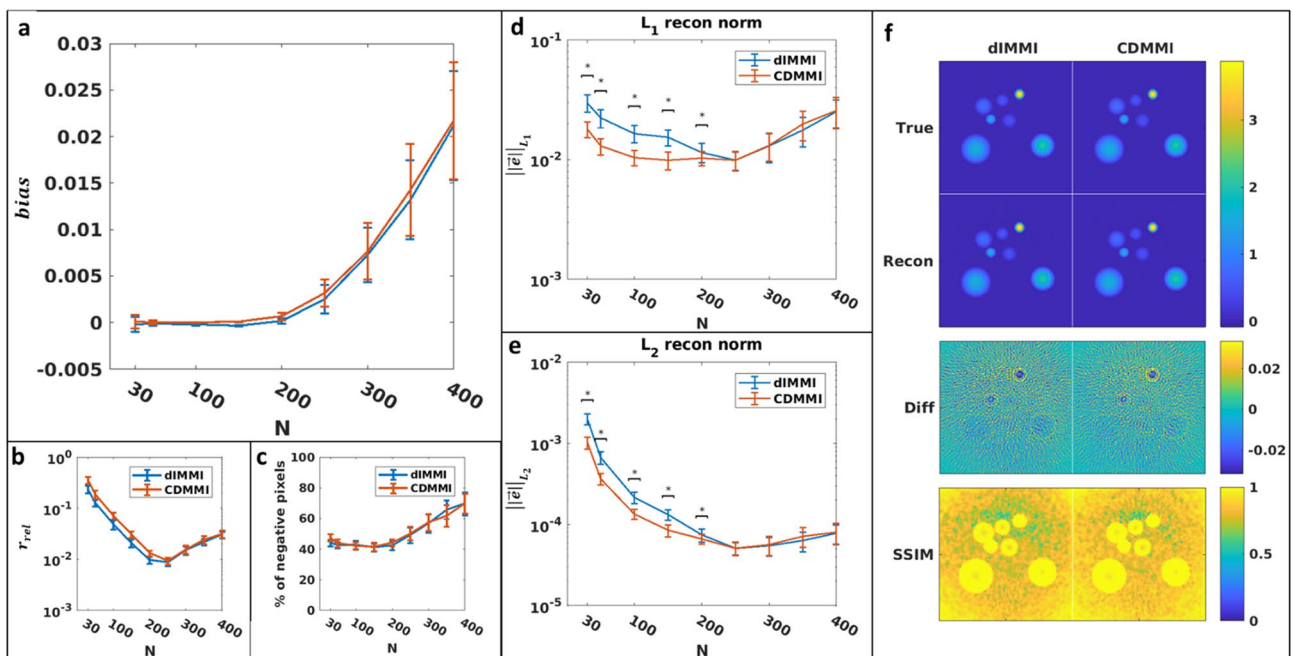


Figure 4. Comparison of dIMMI- and CDMMI-based reconstructions using the LSQR reconstruction algorithm. (a) Both models have a tendency towards increasing absolute error as a function of image size, though (b) the relative residual for CDMMI is higher than dIMMI for low image resolutions, converging at higher resolutions. (c) Both models incur similar numbers of negative pixels in their reconstructions. These trends are reversed, however, when considering the reconstruction L_1 norm (d) and L_2 norm (e) for each model, with CDMMI outperforming dIMMI when reconstructing at low resolutions; again, both methods perform similarly at high resolutions. (f) depicts the ground truth, reconstructed, difference, and SSIM images for each of the methods. Both methods appear to perform similarly well, and the introduction of reconstruction artifacts as seen in the SSIM images affects both to comparable degrees. Figure created using a combination of MATLAB and Microsoft PowerPoint.

variation in the various error metrics shown in Figs. 4 and 5—LSQR internally uses a more numerically stable algorithm and converges more efficiently than the fundamentally nonlinear nnAPCG method. Together with convergence, a similar question applies to analytical quality; although the absolute error of a given reconstruction may be low in a numerical or simulated environment, artifacts may appear in real data due to nonlinearities and modelling inaccuracies, which are not fully captured by the forward models. LSQR, despite its favorable numerical properties, provides no guarantees of converging to a sensible value, and due to the non-local nature of the photoacoustic imaging equations, errors or nonphysical values in one pixel will affect the reliability of values in other pixels. Due to this combination of factors, we recommend the use of the nnAPCG method when possible, as it provides an inherently constrained reconstruction suitable for use in further analyses, particularly regarding spectral unmixing. Indeed, the use of a non-negative analysis constrains the possible $\text{SO}_2^{\text{MSOT}}$ values

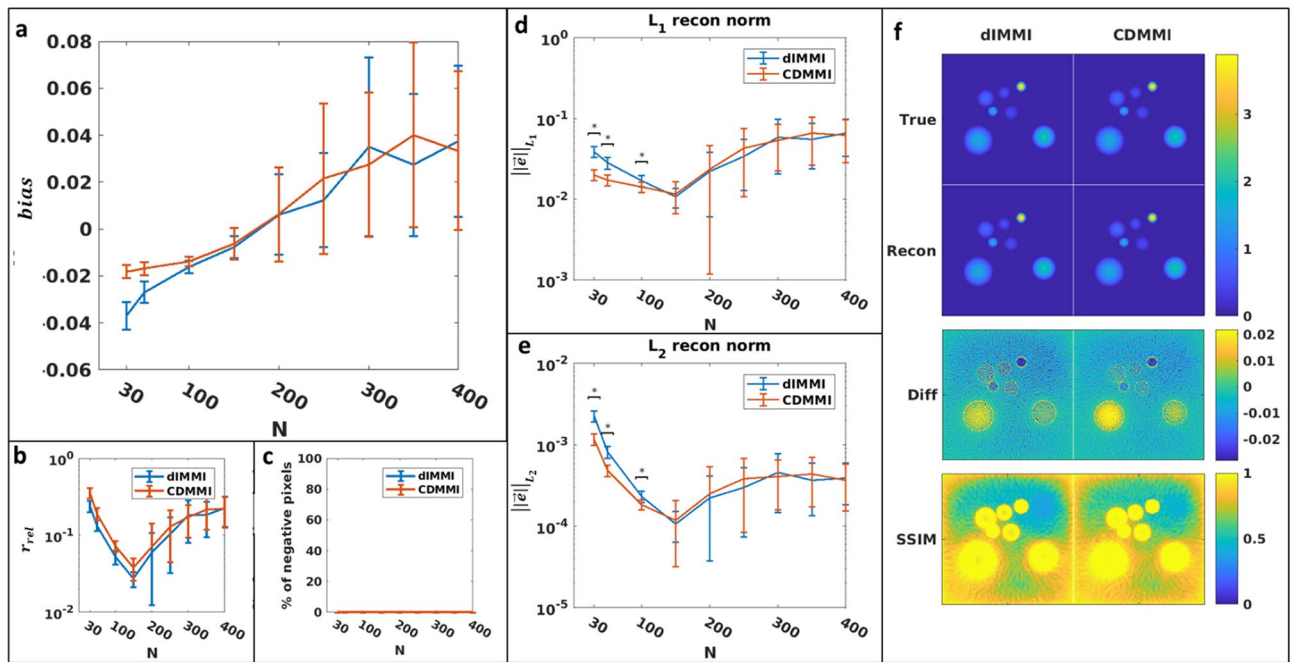


Figure 5. Comparison of dIMMI- and CDMMI-based reconstructions using the nnAPCG reconstruction algorithm. (a) Both models tend towards increasing absolute error as a function of image size though (b) the relative residual for CDMMI is higher than dIMMI for low image resolutions, converging at higher resolutions. (c) Neither model incurs negative pixels due to the non-negative constraint. When considering the reconstruction L_1 norm (d) and L_2 norm (e) for each model, CDMMI results in a lower norm at low resolutions; again, both methods perform similarly at high resolutions. (f) depicts the ground truth, reconstructed, difference, and SSIM images for each of the methods. Both models underestimate the true image intensity for high-intensity objects when reconstructed using nnAPCG, causing mismatches towards the center of each object. Figure created using a combination of MATLAB and Microsoft PowerPoint.

to the range of $[0,1]$ instead of $[-\infty, \infty]$, which may result from various combinations of positive and negative values of oxy- and deoxy-hemoglobin.

Figure 6 illustrates differences in reconstructed performance for three individual pixels in distinct regions of a tumor in a mouse, each showing specific cases within a single study where the constrained approach provides benefits. Though the shapes of the time courses of each parameter within each pixel are consistent between methods, it is clear that the constrained method produces less noisy data. This is due in large part to the use of the $\alpha\beta$ multispectral Kalata filter, which reduces inter-frame noise⁴⁵. The unconstrained approach (red) also demonstrates the risks of using inappropriate analyses when processing data. The $[\text{Hb}]^{\text{MSOT}}$ values in the tumor pixel were negative for a large portion of the experiment (Fig. 6, Pixel 1), which resulted in $[\text{SO}_2]^{\text{MSOT}}$ values exceeding 1. Such pixels can corrupt both spatial and temporal averages, but are not readily compensated; at the same time, the responses of these pixels have resolvable structure, indicating that there is biologically useful data present. The use of the constrained analysis thus provides more complete and more reliable analyses.

As applied to Dataset A, the effects of processing choice are salient: Fig. 7 shows the total number of negative-valued pixels throughout the study. The constrained analysis yields no negative values, while the unconstrained analysis results in negative values in various locations through various channels. Even within the tumor bulk, there are substantial regions where the values of $[\text{HbO}_2]^{\text{MSOT}}$, $[\text{Hb}]^{\text{MSOT}}$, and $\text{SO}_2^{\text{MSOT}}$ attain negative values. The presence of these values would compromise biological inference due to the non-physicality of negative values. Overall differences between the methods for each channel are shown in the Bland–Altman plots (Fig. 8). There is generally good agreement between each, though the constrained method tends to provide lower estimates than the unconstrained method. The presence of an overall diagonal structure reflects the non-negativity of the constrained analysis.

When examining changes in response to treatment, the constrained approach leads to greater significance when comparing each of the parameter values before and 60 min after administration of CA4P (Fig. 9), allowing the resolution of significant changes even in low-signal areas. Similar results were seen when considering the relative effect of the drug administration normalized by the oxygen-air transition response (Fig. 10). There is a large region of anomalous response in the $[\text{Hb}]^{\text{MSOT}}$ and $[\text{HbO}_2]^{\text{MSOT}}$ channels, signified by the magenta arrows in Fig. 10, for both the constrained and unconstrained methods. These variations are attributable to variations in $[\text{Hb}_{\text{tot}}]^{\text{MSOT}}$, which may themselves be due to temperature-dependent signal changes in the animal^{54,55} or systemic blood pressure effects due to CA4P administration⁵⁶. As a result, $\text{SO}_2^{\text{MSOT}}$ provides a more reliable metric of variation due to its intrinsic calibration against $[\text{Hb}_{\text{tot}}]^{\text{MSOT}}$. Similarly, $\text{SO}_2^{\text{MSOT}}$ is less sensitive to variations in light fluence, likely due to variations in light spectrum as a function of wavelength being weaker throughout the bulk of the animal than the variations in light intensity. The unconstrained approach results in a greater

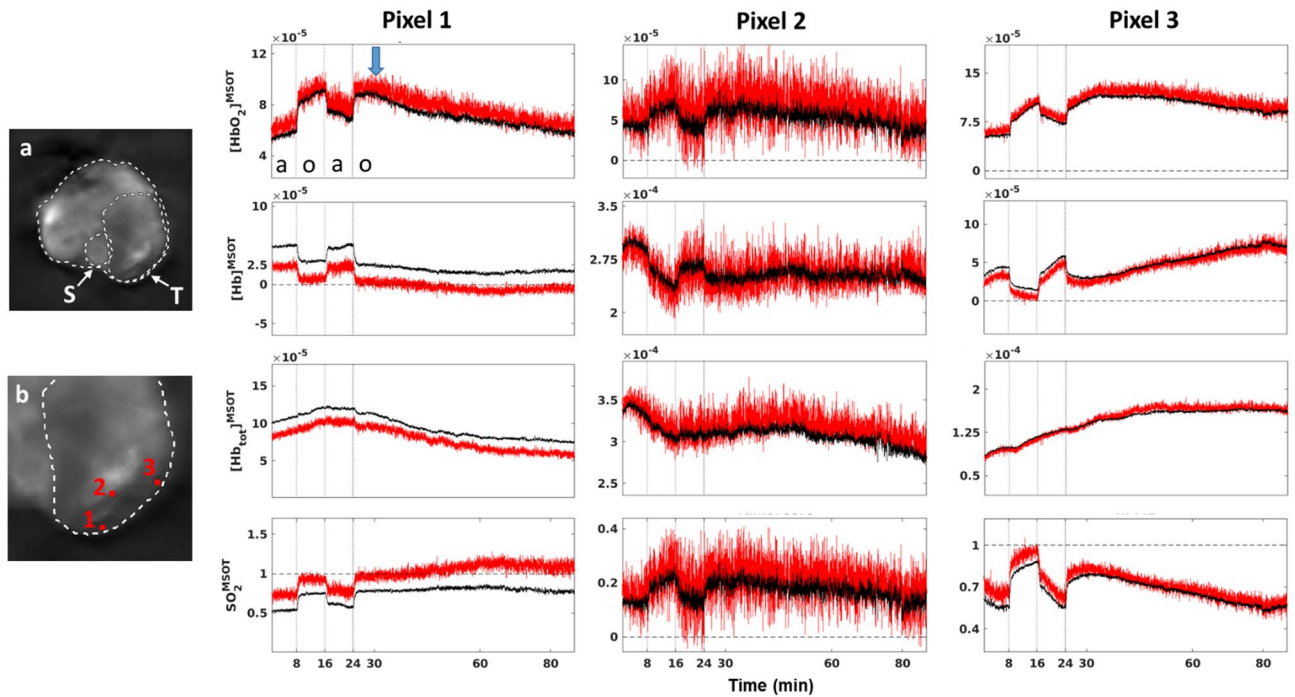


Figure 6. Comparison of different processing approaches implemented using the MSOT-RAFT for analysis of Dataset A. (a) Mean cross-sectional $[Hb_{tot}]^{MSOT}$ image. S: Spine, T: Tumor. (b) Expanded view of tumor periphery, with pixels of interest highlighted; relative size of highlighted squares is larger than a single pixel to improve visibility. The unconstrained approach (red) is substantially noisier and regularly incurs non-physical negative values in the low-signal outer tumor (Pixel 1, $[Hb]^{MSOT}$), which result in spurious values exceeding 1 for downstream SO_2^{MSOT} calculation (Pixel 1, SO_2^{MSOT}). Poor SNR as seen in Pixel 2 is managed with an $\alpha\beta$ filter (black), making the transitions between gases much more conspicuous. The $\alpha\beta$ filter preserves the dynamical structure of each pixel as well, as seen in the transitions of Pixel 3. Figure created using a combination of MATLAB and Microsoft PowerPoint.

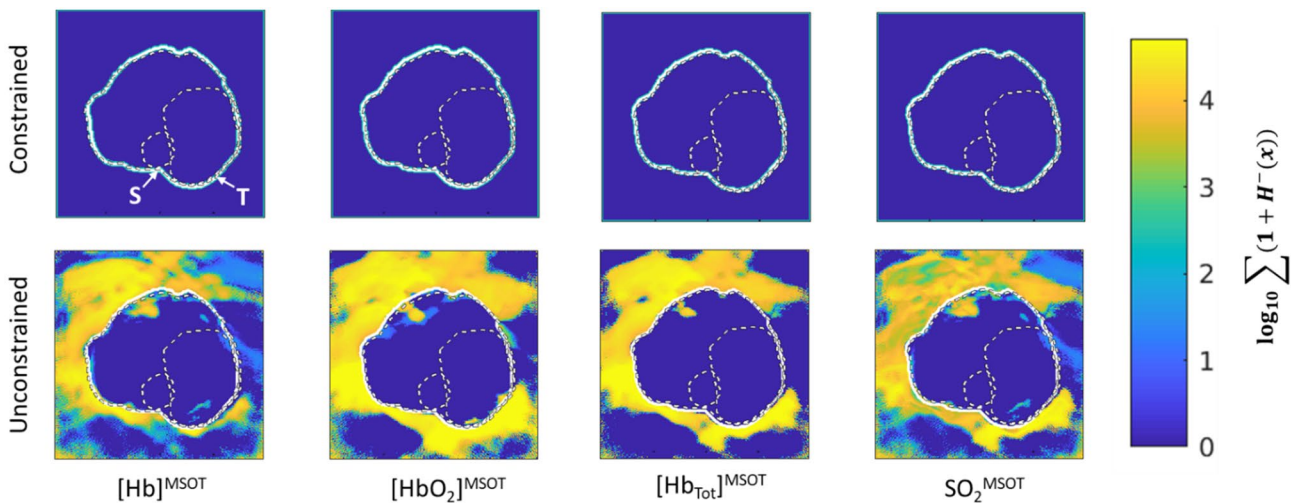


Figure 7. Total occurrence of negative pixels in each channel for the constrained and unconstrained analyses on Dataset A. Animal outline shown, S: Spine, T: Tumor. The constrained analysis results in no negative pixels at any point throughout the imaging time course, reflecting the consistent preservation of the non-negative constraint. The unconstrained analysis, by contrast, develops a large number of negative pixels, including numerous regions within the animal. Figure created using a combination of MATLAB and Microsoft PowerPoint.

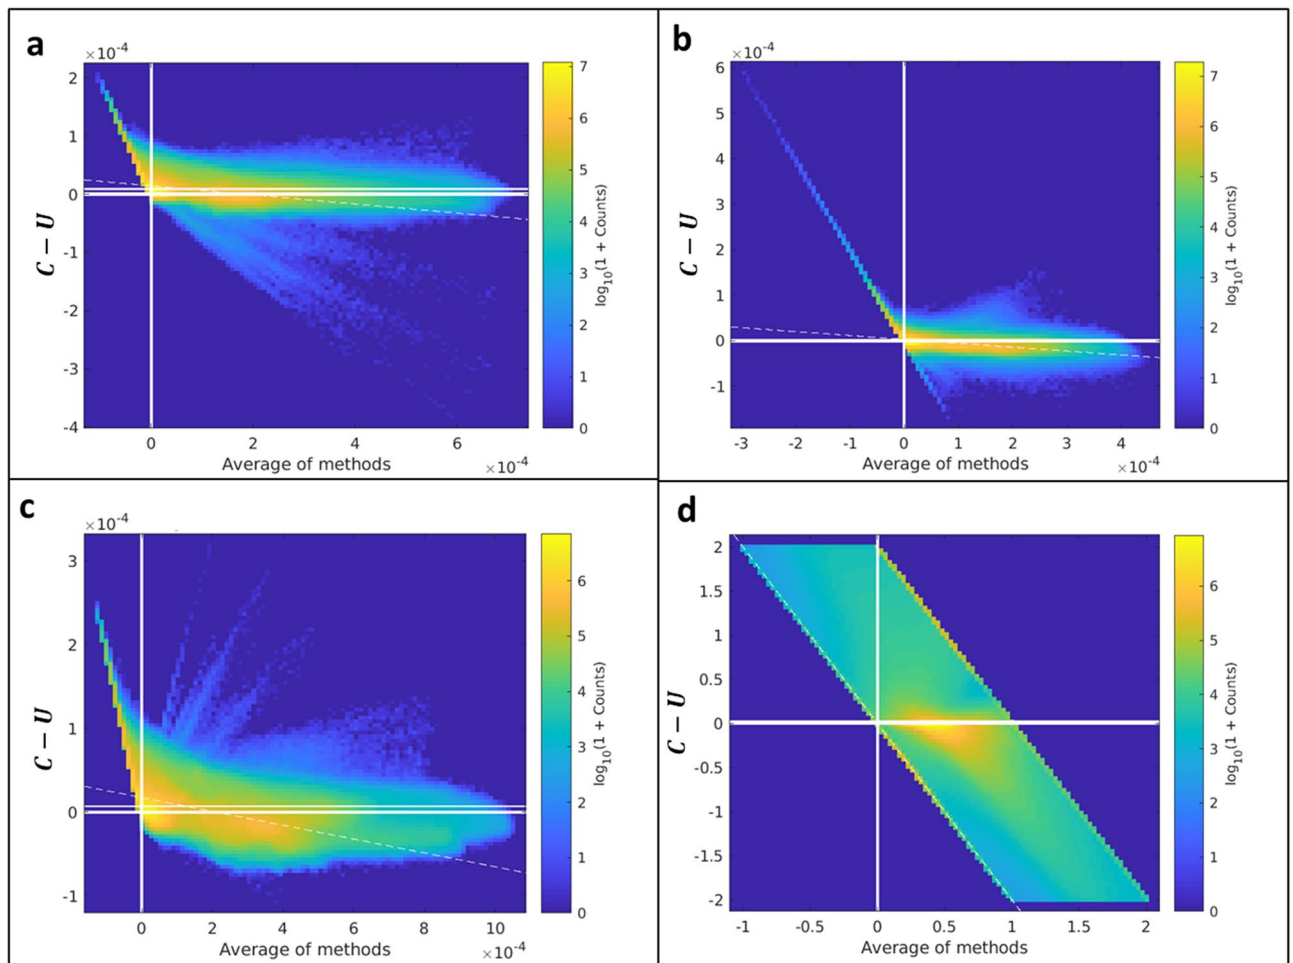


Figure 8. Bland–Altman plots between the constrained and unconstrained analyses on Dataset A. (a) $[\text{HbO}_2]^{\text{MSOT}}$, (b) $[\text{Hb}]^{\text{MSOT}}$, (c) $[\text{Hb}_{\text{Tot}}]^{\text{MSOT}}$, (d) $\text{SO}_2^{\text{MSOT}}$. There is generally good agreement throughout all channels, though there appears to be a consistently lower estimate of all parameters using the constrained approach. The appearance of diagonal structure in (a–c) is due to the non-negative constraint, while the band of finite width in (d) is the result of the constrained method's SO_2 estimates being confined to $[0,1]$. Figure created using a combination of MATLAB and Microsoft PowerPoint.

number of small-scale anomalous responses (Fig. 10, white arrows) due to the non-physical values of various parameters in those pixels. In contrast, the relative effect of the constrained $\text{SO}_2^{\text{MSOT}}$ is much more consistent, showing vessel-like structures in the response within the tumor bulk, with a dramatically reduced occurrence of anomalous responses.

Increases in correlation against the known pO_2 waveform were seen for both the α and $\alpha\beta$ constrained methods when compared to the unconstrained approach (Figs. 11, 12). Though the distributions of correlation were similar (Fig. 11), the constrained methods improved the positive correlation with $[\text{HbO}_2]^{\text{MSOT}}$ and $\text{SO}_2^{\text{MSOT}}$ and the negative correlation with $[\text{Hb}]^{\text{MSOT}}$, while having a relatively mild effect on $[\text{Hb}_{\text{Tot}}]^{\text{MSOT}}$ (Fig. 12). Despite the visual similarities in correlation values between α and $\alpha\beta$ filtered timeseries (Fig. 11), the $\alpha\beta$ filter has the advantage of quickly following dynamic changes, while the α filter naturally suffers from a lag time after such changes. Nevertheless, the α filter suppresses noise effectively, leading to the improved correlations shown in Fig. 12. The improvements in correlation were particularly significant in the HbO_2 and SO_2 channels, possibly reflecting the suitability of these channels for measuring response to gas challenge. The constrained methods were also able to achieve better model fits, as seen in Fig. 13. The unconstrained approach provided results which were generally consistent with the constrained approaches, but resulted in many noncausal switching time values in the tumor region.

The process of reconstruction benefits from parallelization, enabling faster processing. As shown in Fig. 14, even a few additional logical cores, as is available on most modern processors, provided dramatically improved performance. The model-based approaches generally benefited more from this parallelization; due to the larger proportion of the processing time per-frame taken up by the solution process itself, increased parallelization provides greater benefit to the model-based reconstructions. Diminishing returns with increasing numbers of processors are attributable to network and hard-disk limitations. We note that the scaling performance is comparable for each method, indicating that the relative overhead of parallelization for each method is comparable.




Figure 9. T-scores of difference in each parameter before and after CA4P administration in Dataset A. Across all channels, the constrained method results in T scores with substantially increased magnitude, corresponding to greater significance. The large positive region in the upper-left of the $[\text{Hb}]^{\text{MSOT}}$, $[\text{HbO}_2]^{\text{MSOT}}$, and $[\text{Hb}_{\text{Tot}}]^{\text{MSOT}}$ is likely due to physiological effects unrelated to the local vascular effects of CA4P or the choice of processing. This effect is normalized in the ratiometric calculation of $\text{SO}_2^{\text{MSOT}}$. Figure created using a combination of MATLAB and Microsoft PowerPoint.




Figure 10. Relative Cohen's d (Supplementary Information) of drug administration calibrated against oxygen-air transition during initial gas challenge in Dataset A. In both constrained and unconstrained $[\text{Hb}]^{\text{MSOT}}$ and $[\text{HbO}_2]^{\text{MSOT}}$ images, there is a large region of anomalous relative effect (magenta arrows), likely due to changes in measured $[\text{Hb}_{\text{Tot}}]^{\text{MSOT}}$ over the course of the experiment. This large region is absent in the $\text{SO}_2^{\text{MSOT}}$ analyses due to the ratiometric calculation. More localized artifacts such as those highlighted in the unconstrained analysis (white arrows) are due to modelling or reconstruction inaccuracies, and so are not compensated by calculating $\text{SO}_2^{\text{MSOT}}$. Figure created using a combination of MATLAB and Microsoft PowerPoint.

Discussion

The MSOT-RAFT provides a common platform for reconstructing optoacoustic tomography data in a variety of scenarios. The package is open-source and may be scaled to accommodate the computing resources available. Configuration of the overall pipeline during preprocessing enables the straightforward usage of default processing approaches, while simultaneously providing the flexibility for more sophisticated modification of the processing path. We note that this publication is a static record, and recommend examining the Zenodo repository for any updates.




Figure 11. Correlation images against known inhaled O_2 time course for Dataset B. U: Unconstrained analysis. C- $\alpha\beta$: Constrained analysis using $\alpha\beta$ Kalata filter. C- α : Constrained analysis using α Kalata filter. The use of the constrained approach results in improved Pearson's correlation (ρ) with the gas challenge pO_2 time course across all channels. Figure created using a combination of MATLAB and Microsoft PowerPoint.

The MSOT-RAFT is accessible through a variety of interfaces, whether as a MATLAB toolbox or as a compiled executable library usable through other methods, with imminent deployment of Docker images capable of running on Singularity-enabled high performance computing environments. This allows it to be flexibly used in a broad array of contexts, whether running on an individual machine or as a component of a much larger distributed processing pipeline. There are numerous additional factors to consider; light fluence $\phi(\vec{r}, \vec{\lambda})$ is inhomogeneous throughout the imaging region, owing to the removal of photons via absorption by superficial layers. Fluence is also inhomogeneous through wavelength, due to the generally inhomogeneous distribution of endmembers within superficial layers. These additional considerations necessitate additional processing steps, which could be conveniently added through the modular structure of the framework.

The end goal of imaging is the ability to quantitatively resolve the spatial distribution of a variety of parameters, and the reconstruction procedures presented may be augmented for truly quantitative measurements. The reconstructed photoacoustic image is an inversion of a linear map representing the time-distance relationship under the assumption of zero acoustic attenuation and homogenous speed of sound. As others have noted, improved image quality can be attained with spatially-variant time-distance relationships^{57,58}. The determination of the local speed of sound could be performed through Bayesian-type methods⁵⁸, geometric simplification, or adjunct imaging such as transmission-reflection ultrasound⁵⁹.

Though the MSOT-RAFT is presently developed for the input of photoacoustic imaging data acquired using the iThera MSOT imaging systems, we note that this is a question of input format; additional manufacturers and even experimental systems may generate data, which could be successfully reconstructed after suitable input formatting. We anticipate that the MSOT-RAFT will provide an effective starting point for future developments. In particular, the modular structure of the pipeline, and the end-to-end comparisons of reconstruction quality, enable external optimization of the entire system, so as to achieve optimal performance. Since the RAFT is parameter-driven, it could be tuned using a hyperparameter optimizer such as Spearmin in order to create more effective reconstruction pipelines^{60–62}. The resulting optimized pipeline could then be recorded and shared, enabling more rapid dissemination of successful processing motifs.

The RAFT is broadly broken down into modular steps, and they do not all need to be performed using the toolbox itself. Indeed, one could use the RAFT for a large portion of the analysis, and inject additional processing to the analytical chain. This provides for extensive future developments, for example adding fluence correction prior to spectral unmixing.

We plan to extend the filter interface to allow for the calling of external subroutines, such as efficient GPU kernels, compiled C and C++ functions, and various Python scripts. We additionally hope to add support for

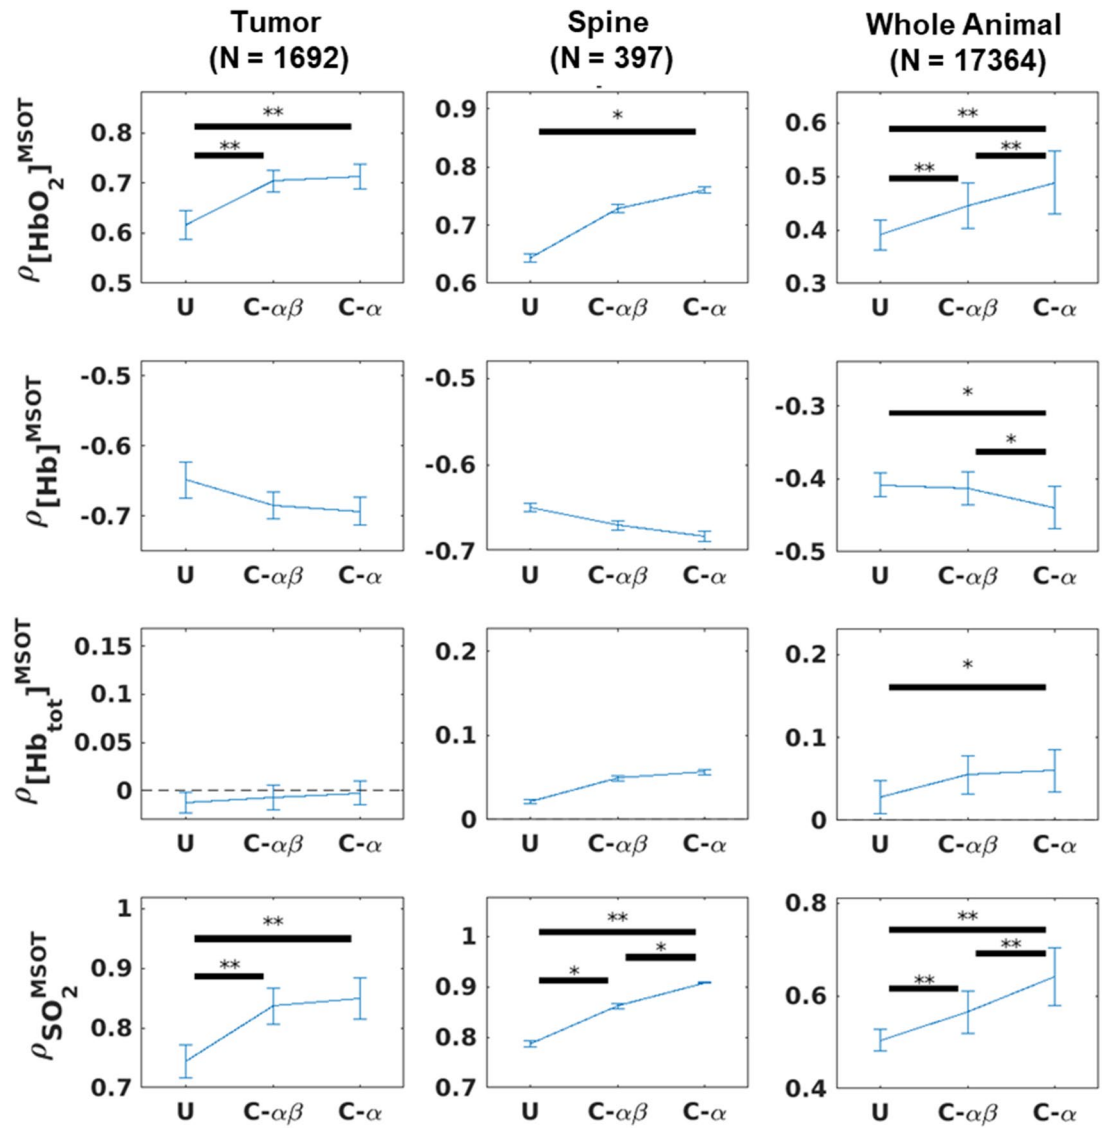


Figure 12. Plots of average correlation for distinct ROIs of Dataset B. A general improvement in correlation is seen using either of the constrained approaches, though the superior noise-rejection properties of the α filter enable it to achieve better overall correlation with the gas challenge pO_2 time course. Figure created using a combination of MATLAB and Microsoft PowerPoint.

standard imaging formats such as DICOM or OME-XML, to enable management of data using PACS systems and inclusion in other studies and databases ranging from the experimental to the clinical.

Conclusion

We have described the implementation and demonstrated the performance of MSOT-RAFT, an open-source toolbox for processing and reconstructing photoacoustic imaging data, and have demonstrated its use for processing and analyzing photoacoustic imaging data.




Figure 13. Selected parameters from fitting Dataset B using 7-parameter monoexponential model (Supplementary Information) using three different processing approaches. U: Unconstrained analysis. $C-\alpha\beta$: Constrained analysis using $\alpha\beta$ Kalata filter. $C-\alpha$: Constrained analysis using α Kalata filter. Although performance of the model fit is similar across much of the imaged area, the unconstrained analysis assigns inaccurate values of switching times to a large region on the interior of the tumor (white arrows). Figure created using a combination of MATLAB and Microsoft PowerPoint.




Figure 14. Reconstruction using different algorithms for varying numbers of parallel workers. Both Dataset A (left) and Dataset C (right) benefit from parallelization, though both show diminishing returns. When reconstruction time is normalized to the time required for a single worker (left) or 4 workers (right) to execute, it becomes clear that both model-based approaches benefit greatly from parallelization, while backprojection (BP) more rapidly saturates due to the greater proportion of its reconstruction which is non-parallelizable. For large numbers of workers, it becomes evident that system overhead becomes restrictive, as there is a decrease in relative speedup. Note separate axes for dIMMI and CDMMI (left) versus BP (right). Figure created using a combination of MATLAB and Microsoft PowerPoint.

Code availability

The most up to date publicly available version of the RAFT, along with examples of usage and test data, can be found at <https://doi.org/10.5281/zenodo.4658279> and is available under the MIT license.

Data availability

All data analyzed in this study are available from the corresponding author (RPM) on reasonable request.

Received: 5 April 2021; Accepted: 17 August 2021

Published online: 06 October 2021

References

1. Beard, P. Biomedical photoacoustic imaging. *Interface focus* **1**, 602–631. <https://doi.org/10.1098/rsfs.2011.0028> (2011).
2. Wang, L. V. *Photoacoustic Imaging and Spectroscopy*. (CRC Press, 2009).
3. Shi, J., Tang, Y. & Yao, J. Advances in super-resolution photoacoustic imaging. *Quant. Imaging Med. Surg.* **8**, 724. <https://doi.org/10.21037/qims.2018.09.14> (2018).
4. Vilov, S. *et al.* Super-resolution photoacoustic and ultrasound imaging with sparse arrays. *Sci. Rep.* **10**, 1–8. <https://doi.org/10.1038/s41598-020-61083-2> (2020).
5. Zhang, P., Li, L., Lin, L., Shi, J. & Wang, L. V. In vivo superresolution photoacoustic computed tomography by localization of single dyed droplets. *Light: Sci. Appl.* **8**, 1–9. doi:<https://doi.org/10.1038/s41377-019-0147-9> (2019).
6. Hupple, C. W. *et al.* A light-fluence-independent method for the quantitative analysis of dynamic contrast-enhanced multispectral optoacoustic tomography (DCE MSOT). *Photoacoustics* **10**, 54–64. <https://doi.org/10.1016/j.pacs.2018.04.003> (2018).
7. Balasundaram, G. *et al.* Noninvasive anatomical and functional imaging of orthotopic glioblastoma development and therapy using multispectral optoacoustic tomography. *Transl. Oncol.* **11**, 1251–1258. <https://doi.org/10.1016/j.tranon.2018.07.001> (2018).
8. Rich, L. J., Miller, A., Singh, A. K. & Seshadri, M. Photoacoustic imaging as an early biomarker of radio therapeutic efficacy in head and neck cancer. *Theranostics* **8**, 2064. <https://doi.org/10.7150/thno.21708> (2018).
9. Hudson, S. V. *et al.* Targeted noninvasive imaging of EGFR-expressing orthotopic pancreatic cancer using multispectral optoacoustic tomography. *Can. Res.* **74**, 6271–6279. <https://doi.org/10.1158/0008-5472.Can-14-1656> (2014).
10. Mallidi, S., Larson, T., Aaron, J., Sokolov, K. & Emelianov, S. Molecular specific optoacoustic imaging with plasmonic nanoparticles. *Opt Express* **15**, 6583–6588. <https://doi.org/10.1364/oe.15.006583> (2007).
11. Mallidi, S., Luke, G. P. & Emelianov, S. Photoacoustic imaging in cancer detection, diagnosis, and treatment guidance. *Trends Biotechnol* **29**, 213–221. <https://doi.org/10.1016/j.tibtech.2011.01.006> (2011).
12. Buehler, A., Kacprowicz, M., Taruttis, A. & Ntziachristos, V. Real-time handheld multispectral optoacoustic imaging. *Opt Lett* **38**, 1404–1406. <https://doi.org/10.1364/OL.38.001404> (2013).
13. Dean-Ben, X. L., Ozbek, A. & Razansky, D. Volumetric real-time tracking of peripheral human vasculature with GPU-accelerated three-dimensional optoacoustic tomography. *IEEE Trans Med Imaging* **32**, 2050–2055. <https://doi.org/10.1109/tmi.2013.2272079> (2013).
14. Dean-Ben, X. L. & Razansky, D. Portable spherical array probe for volumetric real-time optoacoustic imaging at centimeter-scale depths. *Opt Express* **21**, 28062–28071. <https://doi.org/10.1364/OE.21.028062> (2013).
15. McNally, L. R. *et al.* Current and emerging clinical applications of multispectral optoacoustic tomography (MSOT) in oncology. *Clin. Cancer Res.* **22**, 3432. <https://doi.org/10.1158/1078-0432.CCR-16-0573> (2016).
16. Diot, G. *et al.* Multispectral Optoacoustic Tomography (MSOT) of Human Breast Cancer. *Clin. Cancer Res.* **23**, 6912–6922. <https://doi.org/10.1158/1078-0432.ccr-16-3200> (2017).
17. Becker, A. *et al.* Multispectral optoacoustic tomography of the human breast: characterisation of healthy tissue and malignant lesions using a hybrid ultrasound-optoacoustic approach. *Eur. Radiol.* **28**, 602–609. <https://doi.org/10.1007/s00330-017-5002-x> (2018).
18. Dogan, B. E. *et al.* Optoacoustic imaging and gray-scale US features of breast cancers: correlation with molecular subtypes. *Radiology* **292**, 564–572. <https://doi.org/10.1148/radiol.2019182071> (2019).
19. ANSI Standard. Z136. 1–2000, for Safe Use of Lasers. *Published by the Laser* (2000).
20. Thomas, R. J. *et al.* A procedure for multiple-pulse maximum permissible exposure determination under the Z136. 1–2000 American National Standard for Safe Use of Lasers. *J. Laser Appl.* **13**, 134–140. doi: <https://doi.org/10.2351/1.1386796> (2001).
21. Xu, M. & Wang, L. V. Universal back-projection algorithm for photoacoustic computed tomography. *Phys Rev E Stat Nonlin Soft Matter Phys* **71**, 016706. <https://doi.org/10.1103/PhysRevE.71.016706> (2005).
22. Lutzweiler, C., Dean-Ben, X. L. & Razansky, D. Expediting model-based optoacoustic reconstructions with tomographic symmetries. *Med. Phys.* **41**, 013302. <https://doi.org/10.1118/1.4846055> (2014).
23. Wang, K. & Anastasio, M. A. A simple Fourier transform-based reconstruction formula for photoacoustic computed tomography with a circular or spherical measurement geometry. *Phys Med Biol* **57**, N493–499. <https://doi.org/10.1088/0031-9155/57/23/N493> (2012).
24. Dean-Ben, X. L., Buehler, A., Ntziachristos, V. & Razansky, D. Accurate model-based reconstruction algorithm for three-dimensional optoacoustic tomography. *IEEE Trans Med Imaging* **31**, 1922–1928. <https://doi.org/10.1109/TMI.2012.2208471> (2012).
25. Ding, L., Dean-Ben, X. L. & Razansky, D. Real-time model-based inversion in cross-sectional optoacoustic tomography. *IEEE Trans. Med. Imaging* **35**, 1883–1891. <https://doi.org/10.1109/Tmi.2016.2536779> (2016).
26. Liu, H. *et al.* Curve-driven-based acoustic inversion for photoacoustic tomography. *IEEE Trans. Med. Imaging* **35**, 2546–2557. <https://doi.org/10.1109/TMI.2016.2584120> (2016).
27. Oraevsky, A. A., Deán-Ben, X. L., Wang, L. V., Razansky, D. & Ntziachristos, V. Statistical weighting of model-based optoacoustic reconstruction for minimizing artefacts caused by strong acoustic mismatch. **7899**, 789930, doi:<https://doi.org/10.1117/12.874623> (2011).
28. Rosenthal, A., Razansky, D. & Ntziachristos, V. Fast semi-analytical model-based acoustic inversion for quantitative optoacoustic tomography. *IEEE Trans Med Imaging* **29**, 1275–1285. <https://doi.org/10.1109/TMI.2010.2044584> (2010).
29. Zhang, J., Anastasio, M. A., La Riviere, P. J. & Wang, L. V. Effects of different imaging models on least-squares image reconstruction accuracy in photoacoustic tomography. *IEEE Trans Med Imaging* **28**, 1781–1790. <https://doi.org/10.1109/TMI.2009.2024082> (2009).
30. Cox, B. T. & Treeby, B. E. Artifact trapping during time reversal photoacoustic imaging for acoustically heterogeneous media. *IEEE Trans Med Imaging* **29**, 387–396. <https://doi.org/10.1109/TMI.2009.2032358> (2010).
31. Huang, C., Wang, K., Nie, L., Wang, L. V. & Anastasio, M. A. Full-wave iterative image reconstruction in photoacoustic tomography with acoustically inhomogeneous media. *IEEE Trans. Med. Imaging* **32**, 1097–1110. <https://doi.org/10.1109/TMI.2013.2254496> (2013).
32. Xia, J., Yao, J. & Wang, L. V. Photoacoustic tomography: principles and advances. *Electromagn. Waves (Cambridge, Mass.)* **147**, 1, doi:<https://doi.org/10.2528/pier14032303> (2014).

33. Cox, B. T., Arridge, S. R. & Beard, P. C. Estimating chromophore distributions from multiwavelength photoacoustic images. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **26**, 443–455, doi:<https://doi.org/10.1364/josaa.26.000443> (2009).
34. Xu, M. & Wang, L. V. Photoacoustic imaging in biomedicine. *Rev. Sci. Instrum.* **77**, 041101. <https://doi.org/10.1063/1.2195024> (2006).
35. Lutzweiler, C. & Razansky, D. Optoacoustic imaging and tomography: reconstruction approaches and outstanding challenges in image performance and quantification. *Sensors (Basel)* **13**, 7345–7384. <https://doi.org/10.3390/s130607345> (2013).
36. Caballero, M. A. A., Rosenthal, A., Buehler, A., Razansky, D. & Ntziachristos, V. Optoacoustic determination of spatio-temporal responses of ultrasound sensors. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **60**, 1234–1244. <https://doi.org/10.1109/TUFFC.2013.2687> (2013).
37. Treeby, B. E. & Cox, B. T. k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields. *J. Biomed. Opt.* **15**, 021314. <https://doi.org/10.1117/1.3360308> (2010).
38. Dehghani, H. *et al.* Near infrared optical tomography using NIRFAST: Algorithm for numerical model and image reconstruction. *Commun. Numer. Methods Eng.* **25**, 711–732 (2009).
39. Schweiger, M. & Arridge, S. R. The Toast++ software suite for forward and inverse modeling in optical tomography. *J. Biomed. Opt.* **19**, 040801. <https://doi.org/10.1117/1.JBO.19.4.040801> (2014).
40. Wang, L., Jacques, S. L. & Zheng, L. MCML—Monte Carlo modeling of light transport in multi-layered tissues. *Comput. Methods Programs Biomed.* **47**, 131–146 (1995).
41. Jensen, J. A. & Svendsen, N. B. Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **39**, 262–267. <https://doi.org/10.1109/58.139123> (1992).
42. Jensen, J. A. In *10th Nordicbaltic Conference on Biomedical Imaging, vol. 4, Supplement 1, Part 1*: 351–353. (Citeseer).
43. Lu, T. & Mao, H. In *2009 Symposium on Photonics and Optoelectronics*. 1–4 (IEEE).
44. Bigio, I. J. & Fantini, S. *Quantitative Biomedical Optics: Theory, Methods, and Applications*. (Cambridge University Press, 2016).
45. O’Kelly, D., Guo, Y. & Mason, R. P. Evaluating online filtering algorithms to enhance dynamic multispectral optoacoustic tomography. *Photoacoustics* **19**, 100184. <https://doi.org/10.1016/j.pacs.2020.100184> (2020).
46. Hill, M. D. & Marty, M. R. Amdahl’s Law in the Multicore Era. *Computer* **41**, 33–38. <https://doi.org/10.1109/MC.2008.209> (2008).
47. Amdahl, G. M. In *Proceedings of the April 18–20, 1967, Spring Joint Computer Conference*. 483–485. <https://doi.org/10.1145/146544.82.1465560>.
48. Zhou, W., Bovik, A. C., Sheikh, H. R. & Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* **13**, 600–612. <https://doi.org/10.1109/TIP.2003.819861> (2004).
49. Nielsen, T. *et al.* Combretastatin A-4 phosphate affects tumor vessel volume and size distribution as assessed using MRI-based vessel size imaging. *Clin Cancer Res.* **18**, 6469–6477. <https://doi.org/10.1158/1078-0432.ccr-12-2014> (2012).
50. Dey, S. *et al.* The vascular disrupting agent combretastatin A-4 phosphate causes prolonged elevation of proteins involved in heme flux and function in resistant tumor cells. *Oncotarget* **9**, 4090. <https://doi.org/10.18632/oncotarget.23734> (2018).
51. Tomaszewski, M. R. *et al.* Oxygen-enhanced and dynamic contrast-enhanced optoacoustic tomography provide surrogate biomarkers of tumor vascular function, hypoxia, and necrosis. *Can. Res.* **78**, 5980–5991. <https://doi.org/10.1158/0008-5472.CAN-18-1033> (2018).
52. Ding, L., Luis Dean-Ben, X., Lutzweiler, C., Razansky, D. & Ntziachristos, V. Efficient non-negative constrained model-based inversion in optoacoustic tomography. *Phys Med Biol* **60**, 6733–6750, doi:<https://doi.org/10.1088/0031-9155/60/17/6733> (2015).
53. Cheong, W.-F., Prahl, S. A. & Welch, A. J. A review of the optical properties of biological tissues. *IEEE J. Quantum Electron.* **26**, 2166–2185 (1990).
54. Petrova, E., Liopo, A., Oraevsky, A. A. & Ermilov, S. A. Temperature-dependent optoacoustic response and transient through zero Grüneisen parameter in optically contrasted media. *Photoacoustics* **7**, 36–46. <https://doi.org/10.1016/j.pacs.2017.06.002> (2017).
55. Shah, J. *et al.* Photoacoustic imaging and temperature measurement for photothermal cancer therapy. *J. Biomed. Opt.* **13**, 034024. <https://doi.org/10.1117/1.2940362> (2008).
56. Busk, M., Bohn, A. B., Skals, M., Wang, T. & Horsman, M. R. Combretastatin-induced hypertension and the consequences for its combination with other therapies. *Vascul. Pharmacol.* **54**, 13–17. <https://doi.org/10.1016/j.vph.2010.10.002> (2011).
57. Luis Deán-Ben, X., Ntziachristos, V. & Razansky, D. *Effects of Small Variations of Speed of Sound in Optoacoustic Tomographic Imaging*. Vol. 41 (2014).
58. Lutzweiler, C., Meier, R. & Razansky, D. Optoacoustic image segmentation based on signal domain analysis. *Photoacoustics* **3**, 151–158. <https://doi.org/10.1016/j.pacs.2015.11.002> (2015).
59. Merčep, E., Herraiz, J. L., Deán-Ben, X. L. & Razansky, D. Transmission–reflection optoacoustic ultrasound (TROPUS) computed tomography of small animals. *Light: Sci. Appl.* **8**, 18, doi:<https://doi.org/10.1038/s41377-019-0130-5> (2019).
60. Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A. & Talwalkar, A. Hyperband: a novel bandit-based approach to hyperparameter optimization. *J. Mach. Learn. Res.* **18**, 6765–6816 (2017).
61. Hazan, E., Klivans, A. & Yuan, Y. Hyperparameter optimization: a spectral approach. arXiv preprint arXiv:1706.00764 (2017).
62. Bergstra, J. S., Bardenet, R., Bengio, Y. & Kégl, B. In *Advances in Neural Information Processing Systems*. 2546–2554.

Acknowledgements

We would like to acknowledge Stefan Morscher, Neal Burton, Jacob Tippetts, and Clinton Huppel of iThera Medical GmbH, for extensive assistance in organizing and validating the pipeline. We thank David Trudgian and Daniel Moser for their assistance in integrating the pipeline into the high performance computing environment. The research was supported in part by National Institutes of Health (NIH) Grant 1R01CA244579-01A1, Cancer Prevention and Research Institute of Texas (CPRIT) IIRA Grants RP140285 and RP140399 and the assistance of the Southwestern Small Animal Imaging Resource through the NIH Cancer Center Support Grant 1P30 CA142543. DOK was the recipient of a fellowship administered by the Lyda Hill Department of Bioinformatics. The iThera MSOT was purchased under NIH Grant 1 S10 OD018094-01A1.

Author contributions

DO conceived of, designed, and implemented the RAFT, including the MATLAB toolbox and Nextflow wrapper. JC III and JG contributed to the implantation, care, and imaging of the animals included in this study. PK provided valuable guidance relating to the design of the RAFT. VM provided substantial technical assistance in implementing the Nextflow wrapper and continuous integration scheme. AJ provided assistance and code relating to the continuous integration component. LW provided technical oversight and key financial support. RPM provided financial support and supervision of the development process.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-021-97726-1>.

Correspondence and requests for materials should be addressed to R.P.M.

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

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



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

© The Author(s) 2021