TASK-BASED ADAPTIVE TRANSMIT BEAMFORMING FOR EFFICIENT ULTRASOUND QUANTIFICATION

Oisín Nolan, Wessel L. van Nierop, Louis D. van Harten, Tristan S.W. Stevens, Ruud J.G. van Sloun

Dept. of Electrical Engineering, Eindhoven University of Technology, The Netherlands

ABSTRACT

Wireless and wearable ultrasound devices promise to enable continuous ultrasound monitoring, but power consumption and data throughput remain critical challenges. Reducing the number of transmit events per second directly impacts both. We propose a task-based adaptive transmit beamforming method, formulated as a Bayesian active perception problem, that adaptively chooses where to scan in order to gain information about downstream quantitative measurements, avoiding redundant transmit events. Our proposed Task-Based Information Gain (TBIG) strategy applies to any differentiable downstream task function. When applied to recovering ventricular dimensions from echocardiograms, TBIG recovers accurate results using fewer than 2% of scan lines typically used, showing potential for large reductions in the power usage and data rates necessary for monitoring. Code is available at `https://github.com/tue-bmd/task-based-ulsa`.

***Index Terms*—** Transmit beamforming, active perception, subsampling, cognitive ultrasound

1. INTRODUCTION

Ultrasound is a popular modality for medical imaging, offering high temporal resolution, cost-effectiveness, and versatility [1]. Another notable advantage is portability, with recent research showing promising results for wireless and wearable ultrasound technology [2, 3], for example, a wearable ultrasound patch used for continuous monitoring at the ICU. The past years have shown great development in ultrasound patch hardware [4], including electronics miniaturization and skin adhesive materials. Yet, power consumption and data communication bottlenecks remain important challenges for long-term monitoring [3]. These, in turn, are strongly dependent on the number of transmit events per second (pulse repetition interval), with more events typically yielding higher image quality at the expense of the aforementioned energy and communication bottlenecks. Reducing these while retaining performance can be achieved via smart acquisition strategies and powerful reconstruction algorithms [5]. We note further that while a typical ultrasound exam involves acquiring images of the anatomy, the aim in continuous monitoring is to track some quantitative parameters of the anatomy. In this paper, we leverage this fact to develop an algorithm that adaptively steers the ultrasound transmit beam to image only parts of the anatomy that contain information about a quantitative parameter of interest, D_t. In this way, we directly reduce the required number of transmit events, thereby reducing power usage and data rates, while acquiring all relevant information for accurate estimation of D_t.

We model this as a Bayesian *active perception* problem [6], wherein the ultrasound probe becomes a sensing agent that tracks a probability distribution representing its beliefs about the state of D_t over time. It is active in the sense that it *chooses* which scan lines to acquire in order to minimize its own uncertainty about D_t. The agent does this in an iterative fashion, alternating between action and perception in what is referred to as a *perception-action loop*. We refer to this process as *task-based adaptive transmit beamforming*, given that transmit pattern is optimized with respect to some downstream measurement task, with the sole criterion for the measurement task being that it is a differentiable function of the ultrasound image. Our contributions are summarized as follows:

- We derive and implement a novel perception-action loop that drives generic task-based transmit beamforming strategies to recover quantitative parameters with minimal transmit events.
- We evaluate our model using measurements quantifying ventricular hypertrophy produced by the EchoNet-LVH model [7], showing that our algorithm can recover accurate estimates of the measurement signal using a small fraction ($\sim$2%) of the scan lines typically used for imaging.

2. RELATED WORK

A related research direction uses compressed sensing to reduce the data rates required for ultrasound imaging [8]. A number of these works use supervised machine learning methods to recover beamformed images from subsampled channel data using fixed or random subsampling strategies

This work was supported by the European Research Council (ERC) under the ERC starting grant nr. 101077368 (US-ACT).

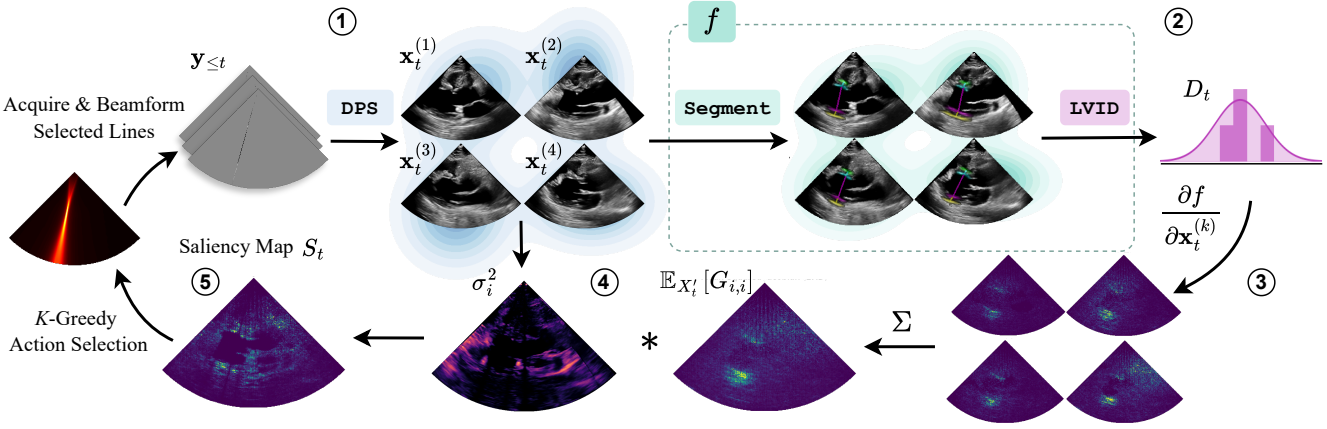

Fig. 1: Diagram illustrating single iteration of the task-based perception-action loop using EchoNetLVH segmentation for the downstream task. ① Generate a set of posterior samples from the sparse acquisition using Diffusion Posterior Sampling (DPS). ② Pass each posterior sample $\mathbf{x}_t^{(i)}$ through the downstream task model f to produce samples from the downstream task distribution. ③ Compute the Jacobian matrix using each of the posterior samples as inputs. ④ Average those Jacobian matrices and multiply them with the pixel-wise variance of the input images to produce the downstream task saliency map. ⑤ Apply K-Greedy Minimization to select K scan line locations for the next acquisition.

that are independent of the data, learning only the reconstruction model [9, 10]. Huijben *et al.* [11] propose a method for learning both the sampling and reconstruction models in an end-to-end manner, where the reconstruction model can be some downstream function of the fully-sampled image. Due to the supervised nature of these methods, however, they need to be trained on datasets consisting of the target subsampling rates and downstream tasks, effectively requiring a re-training of any learned downstream task model, and making them susceptible to out-of-distribution errors if the subsampling rate or task change at test time. Furthermore, the black-box nature of both the sampling policy limits the explainability of the resulting subsampling strategies. An alternative approach is that of *cognitive ultrasound* [6], where acquisition is driven by active perception In this paradigm, the acquisition process is formulated as a perception–action loop, in which beamforming or sampling decisions are chosen to minimize uncertainty about a latent parameter of interest, typically using an explicit observation model for perception, and a white-box action selection function, enhancing interpretability. Federici *et al.* [12] exemplify this approach by designing an active perception algorithm for adaptive beam steering to track the fetal heart and estimate heart rate from power Doppler observations. While the objective is similar to ours – recovering a downstream measurement – their observation model is tailored to Doppler acquisitions and is not applicable to tasks defined on B-mode images. Van Nierop *et al.* [5] propose a cognitive beamforming algorithm that adaptively selects transmit patterns to minimize uncertainty about the reconstructed B-mode image, leveraging state of the art generative image models to implement perception on images. In this work, we build upon the cognitive ultrasound framework introduced by van Nierop *et al.*, but move beyond image fidelity as the sole target, generalizing the approach to downstream quantification functions defined on B-mode images. In our experiments, we benchmark against their method, dubbed General Information Gain (GIG).

3. METHOD

In this section, we derive a novel perception-action loop for task-based transmit beamforming, where the goal is ultimately to minimize uncertainty about D_t. We first introduce a way to quantify that uncertainty, which we show can be decomposed into a sum over a saliency map defined in the image domain, identifying the degree to which each pixel in the image space X_t contributes to the uncertainty in D_t, thereby indicating tissue locations that should be targeted in the next transmit event. Fig. 1 provides a visual overview.

3.1. Perception

The goal of the perception step is to infer a probability distribution over possible values for D_t given the scan lines acquired so far. Our perception step closely follows that of van Nierop *et al.* [5], implementing Bayesian inference via the Diffusion Posterior Sampling (DPS) algorithm [13], which uses a diffusion model [14] to generate samples from the posterior distribution over fully-observed images given partial observations of recent frames $\mathbf{x}_t \sim p(X_t \mid \mathbf{y}_{\leq t})$. Once a

set of posterior samples $\{\mathbf{x}_t^{(i)}\}_{i=1}^{N_p}$ has been generated at time t, the downstream task distribution $D_t \mid \mathbf{y}_{\leq t}$ is approximated by a set of samples $\{\mathbf{d}_t^{(i)} = f(\mathbf{x}_t^{(i)})\}_{i=1}^{N_p}$ produced by passing each posterior image sample through the downstream task model f.

3.2. Action

In order to drive action selection, we first quantify the uncertainty $U(D_t)$ in the state of D_t using the scalar quantity $\mathbb{E}\left[||D_t - \mathbb{E}[D_t]||_2^2\right] = \text{tr}(C_t^D)$ measuring the expected Euclidean distance of the downstream task variable D_t from its mean, with C_t^D being the covariance of $D_t|\mathbf{y}_{\leq t}$. Then, by approximating the downstream task function f with a first-order Taylor expansion, we have that $C^D \approx JC^X J^\top$, where J is the Jacobian matrix of f with reference point $X_t' \sim p(X_t \mid \mathbf{y}_{\leq t})$, X_t and X_t' are independent random variables sampled from the posterior, and C_t^X is the posterior covariance of the images $X_t \mid \mathbf{y}_{\leq t}$. We use a random variable as the reference point to ensure that the Taylor expansion is computed at valid input points, rather than, for example, the posterior mean $\mathbb{E}[X_t \mid \mathbf{y}_{\leq t}]$, which may be off-manifold and out-of-distribution for the downstream task model. If we approximate C_t^X as diagonal, we can then show that $U(D_t)$ can be computed as a sum of saliency-weighted variances from the pixels in the input image X_t:

$$\begin{aligned} U(D_t) &= \mathbb{E}\left[||D_t - \mathbb{E}[D_t]||_2^2\right] \\ &= \text{tr}(C^D) \\ &\approx \text{tr}(JC^X J^\top) \\ &= \text{tr}(J^\top J C^X) && \text{(cyclic property of trace)} \\ &= \text{tr}(GC^X) && (G := J^\top J) \\ &= \sum_i G_{i,i}\sigma_i^2 && (\sigma_i^2 := C_{i,i}^X) \end{aligned} \tag{1}$$

This decomposition therefore becomes a sum of pixel-wise variances in X_t, denoted σ_i^2, weighted by diagonal entries of the Gram matrix $G_{i,i}$. Recall now that the Jacobian matrix is a function of the random variable X_t'. This means that $U(D_t)$, is also a random variable that is a function of X_t'. In order to get a scalar objective to optimize, we therefore take the expected value of $U(D_t)$, denoted $\bar{U}(D_t)$, employing the law of the unconscious statistician (LOTUS) [15]:

$$\begin{aligned} \bar{U}(D_t) &= \mathbb{E}_{U(D_t)}\left[U(D_t)\right] \\ &= \mathbb{E}_{X_t'}\left[U(D_t)\right] && \text{(LOTUS)} \\ &= \mathbb{E}_{X_t'}\left[\text{tr}(GC^X)\right] \\ &= \text{tr}(\mathbb{E}_{X_t'}[G]C^X) && \text{([16])} \\ &= \sum_i \mathbb{E}_{X_t'}[G_{i,i}]\sigma_i^2 \\ &\approx \frac{1}{N}\sum_i \sum_j G_{i,i}\big|_{\mathbf{x}=\mathbf{x}^{(j)}}\sigma_i^2 && \text{(Monte Carlo)} \end{aligned} \tag{2}$$

In other words, we average the Jacobian matrices computed with a set of reference points $\{\mathbf{x}^{(j)}\}_{j=1}^{N_p}$. In practice, we can also estimate σ_i^2 as the empirical variance from the same set of samples. The diagonal of G then contains the L2-norms of the columns of this averaged Jacobian, intuitively quantifying the overall impact that the i^{th} input pixel in X_t has in determining the value of D_t.

Our goal is then to choose a set of scan line locations A_t for the next transmit event that will minimize $\bar{U}(D_t)$. The focusing angles affect $\bar{U}(D_t)$ via σ_i^2, specifically in that observing the i^{th} pixel eliminates its variance, setting $\sigma_i^2 = 0$. Along with this effect, the generative perception model may resolve uncertainty in other pixels whose values can be inferred from the new observation, and certain other pixel variances may increase due to unpredictable temporal dynamics. Accounting for the entire impact of all candidate measurement actions on $\bar{U}(D_t)$ is therefore computationally challenging, potentially requiring the agent to simulate its perception step for the measurements resulting from each possible choice of A_t. Instead, we opt for an approximate greedy algorithm, *K-Greedy Minimization*, introduced by van Nierop *et al.* [5], using the saliency map $S_t = [\mathbb{E}_{X_t'}[G_{1,1}]\sigma_1^2, \ldots, \mathbb{E}_{X_t'}[G_{i,i}]\sigma_i^2, \ldots, \mathbb{E}_{X_t'}[G_{N,N}]\sigma_N^2]$ as the input, rather than a pixel-wise entropy map. The application of K-Greedy Minimization then returns a set of K scan line locations to be transmitted at time $t+1$, completing the action step.

4. RESULTS

In order to evaluate the performance of our algorithm in a practical setting, we apply it to the task of recovering measurements of heart dimensions typically taken from the parasternal long-axis (PLAX) view in echocardiography; in particular, the Left Ventricular Inner Diameter (LVID). We use the EchoNetLVH model [7] to produce heatmaps for the location of each of the measurement anchor points, with the center of mass of each heatmap identifying its anchor point, and the Euclidean distance between anchor points producing the eventual measurement. The measurement process producing $\mathbf{y}_{\leq t}$ is simulated using a masking operation, and a diffusion model trained over sequences of 3 frames from the EchoNetLVH dataset is used for the perception step, following the architecture of van Nierop *et al.* [5]. This was implemented using the `zea` python package [17].

We compare our perception-action loop introduced in Section 3, Task-Based Information Gain (TBIG), to the perception-action loop introduced by van Nierop *et al.* [5], General Information Gain (GIG). The perception model in both methods uses a SeqDiff [18] initialization step of $\tau_{\texttt{SeqDiff}} = 50$, with a total of 500 diffusion steps. The aim of the experiments is to evaluate whether explicitly optimizing the selected transmit angles for a downstream task results in a significantly improved estimate of the downstream param-

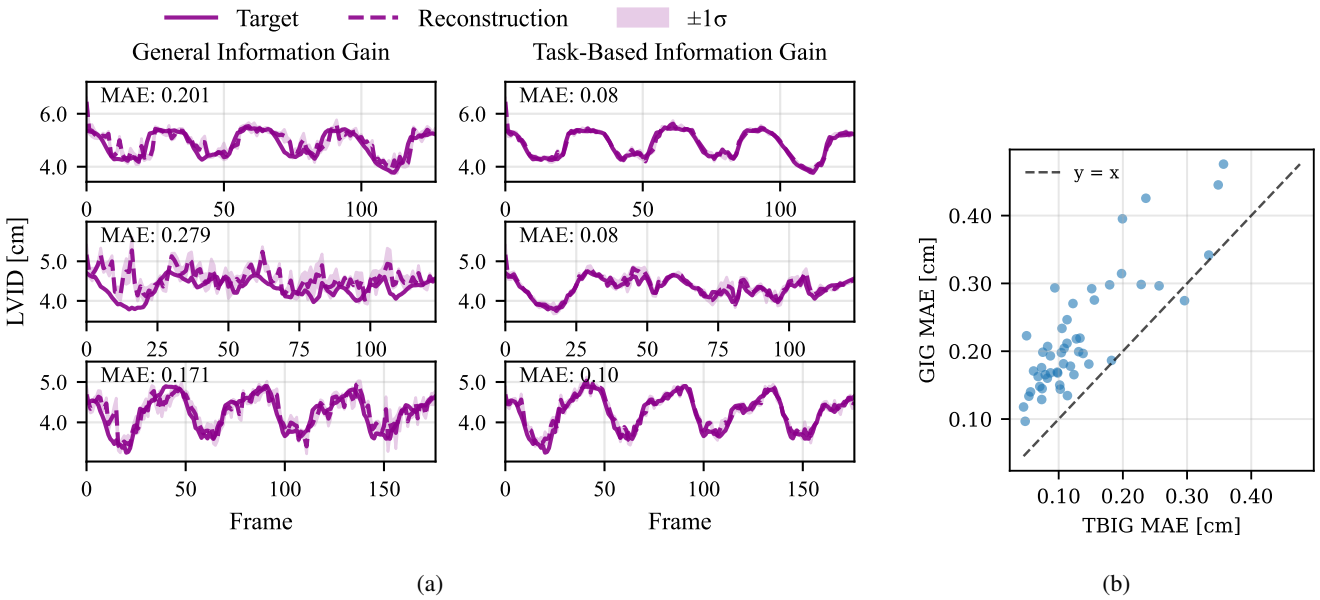

Fig. 2: (a) shows a qualitative comparison of measurement signal recovery for three patients using (i) TBIG and (ii) GIG sampling strategies with 5/256 scan lines. The mean absolute error (MAE) between the target and reconstruction is provided at the top left. The uncertainty for each reconstruction is quantified as the standard deviation of the measurement values estimated from the samples in the belief set computed by applying f to $\{\mathbf{x}_t^{(i)}\}_{i=0}^{N_p}$ for both methods. (b) shows a sample-wise error comparison for each patient in the evaluation set, where it is clear that TBIG almost always outperforms GIG.

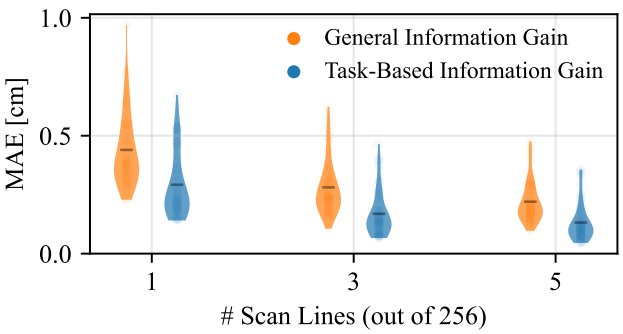

Fig. 3: The distribution of MAE scores between target and reconstructed Left Ventricular Inner Diameter (LVID) time series for the first 100 frames from 50 patients from the EchoNetLVH validation set for both TBIG and GIG strategies.

eters of interest. We compare the two strategies qualitatively and quantitatively in terms of the MAE between the target signal, acquired by applying f to the fully-sampled B-mode images, and the reconstructed signal, acquired by applying the EchoNetLVH model to $\{\mathbf{x}_t^{(i)}\}_{i=1}^{N_p}$, averaging the resulting N_p heatmaps for each anchor point, and extracting the measurements as described above.

In Fig. 2 we compare the downstream measurement signals recovered by TBIG and GIG to the target signals, using a budget of 5/256 scan lines. It is clear that the TBIG objective recovers a significantly more accurate estimate of the target signal in almost all cases. In Fig. 3, we report the distributions of the MAE evaluated on sequences of 100 frames from a set of 50 unseen patients from the EchoNetLVH validation set, using a number of different subsampling rates (1, 3, and 5 lines out of a total of 256).

5. DISCUSSION & CONCLUSION

It is clear from the results that Task-Based Information Gain (TBIG) outperforms General Information Gain (GIG) across all measurement targets and subsampling rates. We observe also that TBIG using only 3 lines outperforms GIG using 5, indicating the potential for a significant reduction in data rate relative to GIG without sacrificing quality. We also highlight that the measurement signals have been recovered here using only a tiny fraction of the scan lines typically used (less than 2%), indicating potential for use in continuous monitoring and other wireless ultrasound technologies. A promising avenue for future work is therefore to apply this method in a continuous monitoring setting, using data from a wearable ultrasound patch, and evaluate its effect on battery life and data throughput.

6. REFERENCES

[1] Vincent Chan and Anahi Perlas, "Basics of ultrasound imaging," in *Atlas of ultrasound-guided procedures in interventional pain management*, pp. 13–19. Springer, 2010.

[2] Laura García, Sandra Viciano-Tudela, Jaime Lloret, and José-Manuel Moltó-Jorda, "Survey on wireless ultrasound imaging," *Health and Technology*, pp. 1–12, 2025.

[3] Hao Huang, Ray S Wu, Muyang Lin, and Sheng Xu, "Emerging wearable ultrasound technology," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 71, no. 7, pp. 713–729, 2023.

[4] Muyang Lin, Ziyang Zhang, Xiaoxiang Gao, Yizhou Bian, Ray S Wu, Geonho Park, Zhiyuan Lou, Zhuorui Zhang, Xiangchen Xu, Xiangjun Chen, et al., "A fully integrated wearable ultrasound system to monitor deep tissues in moving subjects," *Nature biotechnology*, vol. 42, no. 3, pp. 448–457, 2024.

[5] Wessel L van Nierop, Oisín Nolan, Tristan SW Stevens, and Ruud JG van Sloun, "Patient-adaptive focused transmit beamforming using cognitive ultrasound," *arXiv preprint arXiv:2508.08782*, 2025.

[6] Ruud JG Van Sloun, "Active inference and deep generative modeling for cognitive ultrasound," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 2024.

[7] Grant Duffy, Paul P Cheng, Neal Yuan, Bryan He, Alan C Kwan, Matthew J Shun-Shin, Kevin M Alexander, Joseph Ebinger, Matthew P Lungren, Florian Rader, et al., "High-throughput precision phenotyping of left ventricular hypertrophy with cardiovascular deep learning," *JAMA cardiology*, vol. 7, no. 4, pp. 386–395, 2022.

[8] Musyyab Yousufi, Muhammad Amir, Umer Javed, Muhammad Tayyib, Suheel Abdullah, Hayat Ullah, Ijaz Mansoor Qureshi, Khurram Saleem Alimgeer, Muhammad Waseem Akram, and Khan Bahadar Khan, "Application of compressive sensing to ultrasound images: a review," *BioMed research international*, vol. 2019, no. 1, pp. 7861651, 2019.

[9] Alon Mamistvalov, Ariel Amar, Naama Kessler, and Yonina C Eldar, "Deep-learning based adaptive ultrasound imaging from sub-nyquist channel data," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 69, no. 5, pp. 1638–1648, 2022.

[10] Shujaat Khan, Jaeyoung Huh, and Jong Chul Ye, "Adaptive and compressive beamforming using deep learning for medical ultrasound," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 67, no. 8, pp. 1558–1572, 2020.

[11] Iris AM Huijben, Bastiaan S Veeling, Kees Janse, Massimo Mischi, and Ruud JG van Sloun, "Learning subsampling and signal recovery with applications in ultrasound imaging," *IEEE Transactions on Medical Imaging*, vol. 39, no. 12, pp. 3955–3966, 2020.

[12] Beatrice Federici, Ruud JG van Sloun, and Massimo Mischi, "Active inference for closed-loop transmit beamsteering in fetal doppler ultrasound," *arXiv preprint arXiv:2410.04869*, 2024.

[13] Hyungjin Chung, Jeongsol Kim, Michael Thompson McCann, Marc Louis Klasky, and Jong Chul Ye, "Diffusion posterior sampling for general noisy inverse problems," in *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. 2023, OpenReview.net.

[14] Jonathan Ho, Ajay Jain, and Pieter Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.

[15] Hossein Pishro-Nik, *Introduction to probability, statistics, and random processes*, Kappa Research, LLC Blue Bell, PA, USA, 2014.

[16] Joram Soch et al., "Statproofbook/statproofbook.github.io: The book of statistical proofs," `https://statproofbook.github.io/`, 2024, Version 2023, Zenodo: `https://doi.org/10.5281/zenodo.4305949`.

[17] Tristan S.W. Stevens, Wessel L. van Nierop, Ben Luijten, Vincent van de Schaft, Oisín I. Nolan, Beatrice Federici, Louis D. van Harten, Simon W. Penninga, Noortje I.P. Schueler, and Ruud J.G. van Sloun, "zea: A Toolbox for Cognitive Ultrasound Imaging," July 2025.

[18] Tristan SW Stevens, Oisín Nolan, Jean-Luc Robert, and Ruud JG Van Sloun, "Sequential posterior sampling with diffusion models," in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–5.