

DIFFUSION ACTIVE LEARNING: TOWARDS DATA-DRIVEN EXPERIMENTAL DESIGN IN COMPUTED TOMOGRAPHY

Anonymous authors

Paper under double-blind review

ABSTRACT

We introduce *Diffusion Active Learning*, a novel approach that integrates a generative diffusion model with sequential experimental design to adaptively acquire data for solving inverse problems in imaging. We first pre-train an unconditional diffusion model on domain-specific data. The diffusion model is aimed to capture the structure of the underlying data distribution, which is then leveraged in the active learning process. During the active learning loop, we use the forward model of the inverse problem together with the diffusion model to generate conditional data samples from the posterior distribution, all consistent with the current measurements. Based on the generated samples we quantify the uncertainty in the current estimate in order to select the most informative next measurement. We showcase the proposed approach for its application in X-ray computed tomography imaging. Our results demonstrate significant reductions in data acquisition requirements (*i.e.*, lower X-ray dose) and improved image reconstruction quality across several real-world tomography datasets.

1 INTRODUCTION

Computed Tomography (CT) is an imaging technique for reconstructing objects from X-ray projection data. The concept of tomography originated with the work of Radon in 1917, however, it was not until 1971 that Godfrey Hounsfield and Allan Cormack developed the first practical CT scanner. Recent advances at large synchrotron facilities have pushed CT resolution into the nanometer range, enabling novel scientific applications, such as the inspection of composite materials, quality control of computer chips, and revealing cellular structure in biological tissues. Achieving a resolution as low as several nanometers requires image acquisition times of up to several days (Aidukas et al., 2024). In such settings, the X-ray dose deposited onto the sample becomes a key limiting factor, causing radiation damage which ultimately limits the achievable resolution (Howells et al., 2009). Learning based reconstruction methods and data-driven ‘smart’ acquisition techniques are a promising avenue to improve data efficiency, allowing for high-resolution reconstructions with lower acquisition times and reduced X-ray dose.

Mathematically, tomographic projections are described by the Radon transform (Deans, 1983) of the object. In the corresponding inverse problem, multiple 2D projections from different angles are combined into a single 3D reconstruction of the object (Fig. 1). However, traditional reconstruction algorithms such as filtered back-projection or iterative reconstruction schemes only make use of information contained in the measurements (Kak & Slaney, 2001), and neglect additional structure in the data distribution. For example, tomographic scans of computer chips or composite materials display highly regular structures (see Fig. 3). Such regularity can be learned from prior data sets and facilitate improved learning-based reconstruction algorithms.

Moreover, standard CT acquisition scans are acquired by either rotating the object or the CT scanner over an equidistant angle grid. Uniform scanning neglects any potential structure in the data distribution, that could allow to achieve better reconstructions by scanning the most informative angles. Active learning is a subfield of machine learning that studies algorithms for adaptive data acquisition, with the goal to obtain the most informative data points (Settles, 2009). However, the adoption of active learning in applications remains challenging, with literature highlighting instances

054 where active learning methods have shown limited effectiveness compared to conventional uniform
055 or i.i.d. training schemes (e.g., Lowell et al., 2019). This is due to the many underlying challenges
056 in sequential decision-making algorithms, such as the need for uncertainty quantification, as well as
057 added computational and implementation complexity. The effectiveness of active learning further
058 relies on exploiting inherent structure in the data distribution and on statistical modeling assumptions
059 (Balcan et al., 2010), which are often difficult to specify for the domain of interest.

060 In this work, we set to address the outlined challenges in computed tomography by combining
061 a learned generative prior with data-driven experimental design. We introduce *Diffusion Active*
062 *Learning* (DAL), which combines generative diffusion models with active learning to solve inverse
063 problems in imaging. To leverage the structure of the data distribution, we pre-train the diffusion
064 model on data collected from slices of tomographic reconstructions (e.g., image slices of integrated
065 circuits or composite materials). In the active learning loop, we generate samples from the posterior
066 distribution of the diffusion model, conditioned on the measurements collected so far. Based on the
067 generated samples, we quantify the uncertainty in the solution of the inverse problem, and use it to
068 select where to sample next. By repeating this process, and as more measurements are incorporated,
069 we effectively constrain the Diffusion posterior distribution until it collapses to a deterministic,
070 data-consistent final estimate.

071 Using a diffusion model as a prior for the active learning enables several key advantages: (i) Unlike
072 traditional regularizers (e.g., Total Variation), the learned prior is data-dependent and captures
073 problem specific structure. (ii) The structure can be extracted from (large-scale) prior data, avoiding
074 to manually specify intricate regularities in the statistical modeling approach. (iii) The diffusion
075 model successfully captures multi-modal, highly structured distributions (e.g., natural images). This
076 is unlike, for example, the Gaussian Laplace approximation used in prior works (Antoran et al.,
077 2023; Barbano et al., 2022a), which are inherently unimodal.

078 We demonstrate the effectiveness of our approach on three real-world tomography datasets and
079 conduct an extensive evaluation of different acquisition strategies. Our results shows that DAL
080 provides significant improvements in reconstruction quality with fewer measurements compared to
081 conventional uniformly sampled projection measurements. For the scientific datasets benchmarked
082 in this paper, we achieve the same average Peak Signal-to-Noise Ratio (PSNR) in the reconstruction
083 with up to four times fewer measurements, which translates to up to $4\times$ reduction in X-ray dose
084 (see Table 1). Furthermore, we achieve this with more than $2\times$ computational speed-ups compared
085 to the second-fastest baseline (Barbano et al., 2022a) (see Figure 5)

086 **Remark:** In a sparse reconstruction settings, diffusion-based approaches can introduce artifacts or
087 hallucinations as not enough data is measured; preventing a direct use in high-stakes setting such
088 as medical diagnosis (medical X-ray scans are also very fast, and therefore not considered a direct
089 application of the proposed approach). The sparse reconstruction setting, however, is inherently
090 ill-posed, and as such, a learned prior provides a good trade-off between provable accuracy and
091 data-efficiency.

092 1.1 RELATED WORK

093
094 There is a long history of work that uses ideas from computer vision to improve image quality in
095 tomography and medical imaging, see e.g., (Li et al., 2022; Parvaiz et al., 2023) for comprehen-
096 sive surveys. Building on ideas from image denoising and segmentation (Ronneberger et al., 2015),
097 Ulyanov et al. (2018) propose the *Deep Image Prior* (DIP), which utilizes the structure of a randomly
098 initialized convolutional network (U-net) as a prior for image reconstruction tasks. This approach
099 demonstrated that even without pre-training, deep neural network architectures can provide an im-
100 plicit bias that improves image reconstruction quality. The DIP methodology has influenced various
101 subsequent works that incorporate deep learning architectures in inverse problem settings, e.g., using
102 pre-training and an initial reconstruction (Bagger et al., 2020; Barbano et al., 2022b).

103 The success of diffusion-based models in image reconstruction has also been extended to solve in-
104 verse problems. When working with ill-posed inverse problems and using sparse measurements, the
105 goal is to use the diffusion model as a prior to fill in the missing information in the reconstruction.
106 However, sampling from the posterior distribution conditioned on the measurements is a challeng-
107 ing task. As the number of measurements may be different every time, techniques for conditional
sampling used in common text-to-image model have not gained traction for inverse problem solv-




Figure 1: Left: An illustration of X-ray tomography with rotated sample and projection measurements on the detector. Right: 3D reconstruction can be done slice by slice of each vertical sample layer, simplifying tomographic reconstruction to a series of 2D reconstructions from 1D projections.

ing. Nonetheless, Song et al. (2021) apply score-based generative models specifically to sparse reconstruction in medical CT imaging. They sample from the posterior distribution by steering the diffusion process using the measurements. This decouples the training and inference processes and is adopted by most methods hereafter. Their technique is however limited to linear models and tends to fail as measurements become noisy. Chung et al. (2022) introduced the concept of *Diffusion Posterior Sampling*. By modifying the reverse diffusion process, they sample from the posterior distribution, significantly enhancing the reconstruction quality from incomplete or noisy data. Song et al. (2023) extend the use of diffusion models by introducing latent space diffusion models and the concept of *Hard Data Consistency*. This approach solves an optimization problem to align generated samples with observed data, thereby ensuring that reconstructions adhere closely to the measurements. Their method shows notable improvements in handling ill-posed inverse problems, particularly in medical imaging contexts. Lastly, Barbano et al. (2023) proposed a steerable conditional diffusion model designed to adapt to out-of-distribution scenarios in imaging inverse problems. Their method ensures that the generative process remains robust even when the data deviates from the training distribution.

Active learning is a huge field (e.g., Settles, 2009) with roots in Bayesian experimental design (Chaloner & Verdinelli, 1995). The active learning literature has mostly focused on the classification setting, where the goal is to reduce the labelling effort by actively querying the label for the most uncertain data points. These ideas have been applied to deep learning Ren et al. (2021) and medical imaging Budd et al. (2021) as well. Gal et al. (2017) discuss various acquisition strategies for active learning with image data, also applicable in the regression setting relevant to this work. Only few prior works have explored experimental design for computed tomography. Closely related is the work by Barbano et al. (2022a), who introduce Bayesian experimental design for computed tomography. Their main innovation is the use of a Linearized Laplace Deep Image Prior for uncertainty quantification Antoran et al. (2022) to guide the acquisition of measurement data. Although this demonstrates the applicability of sequential experimental design in the CT setting, the evaluation is still limited to a simplistic toy example. Subsequently, Antoran et al. (2023) further scale the Laplace approach using a sampling based technique and demonstrate it on a CT reconstruction task, although not in combination with active learning.

Going well-beyond one-step (greedy) active learning schemes, reinforcement learning (RL) provides a framework to solve a multi-stage planning problem over the combinatorial space of possible experimental designs. This is prominently explored in the related field of magnetic resonance imaging (MRI), see, for example, the works by Zhang et al. (2019); Pineda et al. (2020); Bakker et al. (2020); Jin et al. (2019). Similarly, RL based methods for computed tomography are proposed by Wang et al. (2023); Shen et al. (2022). These approaches primarily address the problem of choosing an optimal sequence of measurements for a given reconstruction method; this comes at the price of increased algorithmic and training complexity compared to greedy active learning schemes. We also remark that unlike MRI, the tomographic forward model is linear, and greedy approaches provably converge (Riquelme et al., 2017). In principle, RL based methods can be combined with a learned prior; this is however beyond the scope of this work.

Algorithm 1: Diffusion Active Learning

Input: k : number of conditional samples, \mathcal{D}_1 : initial set of measurements, \mathcal{M} : pre-trained diffusion model

```

1 for  $t = 1, \dots, n$  do
2   samples:  $\mathbf{x}_t^1, \dots, \mathbf{x}_t^k = \mathcal{M}.\text{conditional\_sampling}(\text{data} = \mathcal{D}_t, \text{num\_samples} = k)$ 
3   mean prediction:  $\bar{\mathbf{x}}_t = \frac{1}{k} \sum_{i=1}^k \mathbf{x}_t^i$ 
4   maximum variance acquisition:  $\psi_t = \arg \max_{\psi \in \Phi} \frac{1}{k} \sum_{i=1}^k \|\mathcal{A}_\psi(\mathbf{x}_t^i) - \mathcal{A}_\psi(\bar{\mathbf{x}}_t)\|^2$ 
5   new measurement:  $\mathbf{y}_{\psi_t} = \mathcal{A}_{\psi_t}(\mathbf{x}^*) + \epsilon_t$ 
6   data update:  $\mathcal{D}_{t+1} = \mathcal{D}_t \cup \{(\psi_t, \mathbf{y}_{\psi_t})\}$ 

```

2 SETTING

In this work, we consider the reconstruction of 2D objects (images) from their 1D projections (measurements). The reconstruction of 3D volumes from 2D projections follows from the same principle. Alternatively, 3D volumes can also be reconstructed by stacking 2D slices. Figure 1 illustrates a typical acquisition setup for 3D volume reconstruction.

Formally, let $\mathbf{x}^* \in \mathbb{R}^{d \times d}$ be a 2-dimensional grayscale image, corresponding to a slice of the object that we aim to reconstruct. In our simplified setup, we assume that the detector has $l \in \mathbb{N}$ pixels, corresponding to the resolution of the observed projection. For a given angle ψ , the observed measurement $\mathbf{y}_\psi \in \mathbb{R}^l$ is given by a forward operator $\mathcal{A}_\psi : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^l$ and a noise vector $\epsilon \in \mathbb{R}^l$,

$$\mathbf{y}_\psi = \mathcal{A}_\psi(\mathbf{x}^*) + \epsilon. \quad (1)$$

The noise is often assumed to be Gaussian or Poisson distributed. For the special case of parallel beam tomography, the forward model $\mathcal{A}_\psi(\mathbf{x})$ is a linear operator that is mathematically determined by the Radon transform (Kak & Slaney, 2001).

For a given set of n projections $\mathcal{Y}_\psi = \{\mathbf{y}_{\psi_1}, \dots, \mathbf{y}_{\psi_n}\}$ (or \mathcal{Y} for simplicity) and measured at angles $\psi = \{\psi_1, \dots, \psi_n\}$, the goal is to reconstruct \mathbf{x} from \mathcal{Y} . Assuming a Gaussian distribution of the noise ϵ , the reconstruction problem can be solved using maximum-likelihood inference:

$$\underset{\mathbf{x} \in \mathbb{R}^{d \times d}}{\text{minimize}} \sum_{\psi \in \Psi} \|\mathcal{A}_\psi(\mathbf{x}) - \mathbf{y}_\psi\|_2^2. \quad (2)$$

A few important challenges arise. First, the CT problem is typically too high-dimensional to be solved in closed form, and hence one has to resort to iterative or gradient descent based schemes. This means that uncertainty quantification methods that require the second moment of the posterior distribution are not computationally feasible without further approximation. In particular, many active learning algorithms require an uncertainty estimate or samples from the posterior distribution, and are therefore challenging to implement in the CT setting. Second, in the sparse reconstruction regime ($n \times l < d^2$), the problem is underdetermined. The most common remedy is to add a regularizer (e.g., L2 or Total-Variation loss), or other means of adding an inductive bias (e.g., using non-linear representations and pre-training).

Turning now to the sequential experimental design setting, we consider the measurement space to be a uniform set of rotation angles $\Phi = \{i \cdot \Delta\phi\}_{i=0}^{N-1}$ out of which a subset ψ of size $n < N$ are going to be used for the reconstruction. In our setting, we use $N = 180$ angles with $\Delta\psi = 1^\circ$ increment. The goal is to iteratively select those angles that are most informative in the sense that they yield the lowest reconstruction error. For selecting the angles, we proceed by sequentially selecting one angle at a time until the budget of n angles is exhausted (see Algorithm 1).

3 DIFFUSION ACTIVE LEARNING

We now describe *Diffusion Active Learning* (DAL), a novel approach for data-driven, ‘smart’ angle selection in computed tomography. In a pre-training step, we train an unconditional Denoising Diffusion Probabilist Model (DDPM) on a training set consisting of tomographic slices from objects in

the domain of interest. This training is entirely independent of the inference (reconstruction) given the X-ray measurements, and requires only samples of tomographic slices to learn the image distribution. Our implementation uses the classical DDPM training (Ho et al., 2020). During the active learning loop, we use the trained diffusion model to approximate the posterior distribution conditioned on the current set of measurements (X-ray projections). We use a variant of the techniques proposed in Song et al. (2023) to generate these conditional samples. Given the relatively small size of the images in our study, we opt for a pixel-space diffusion model rather than a latent-space diffusion model. More generally, we emphasize that any generative posterior can be used for the active learning procedure described below. A more detailed comparison of different diffusion models is, however, beyond the scope of this work.

In the active learning loop, we use conditional samples from the diffusion model to approximate the uncertainty of the estimation. Using the forward model, we map the diffusion samples to the measurement space, to obtain the posterior distributions of the projections. Finally, we choose the angle that has the largest uncertainty to take the next measurement and repeat until our measurement budget is depleted. The active learning step is described in more detail in Section 3.2. The complete diffusion active learning framework is outlined in Algorithm 1.

From a Bayesian perspective, the Diffusion model corresponds to a learned prior $p(\mathbf{x})$ over images $\mathbf{x} \in \mathbb{R}^{d \times d}$. Diffusion posterior inference approximates the posterior, $p(\mathbf{x}|\mathcal{Y}, \psi) \propto p(\mathbf{x})p(\mathcal{Y}|\mathbf{x}, \psi)$ for measurements \mathcal{Y} . The likelihood term is specified by the forward model in Eq. (1) and the noise distribution. The posterior (mean) distribution of a new measurement $\mathbf{y}_{\psi_{new}} \in \mathbb{R}^d$ at angle ψ_{new} is specified again by the forward model, marginalized over the posterior, $p(\mathbf{y}_{\psi_{new}}|\mathcal{Y}, \psi, \psi_{new}) = \int \mathcal{A}_{\psi_{new}}(\mathbf{x})p(\mathbf{x}|\mathcal{Y}, \psi)d\mathbf{x}$. Equivalently, we can sample from the predictive posterior by sampling first from the posterior $\tilde{\mathbf{x}} \sim p(\mathbf{x}|\mathcal{Y}, \psi)$ and applying the forward model to obtain $\tilde{\mathbf{y}}_{\psi_{new}} = \mathcal{A}_{\psi_{new}}(\tilde{\mathbf{x}})$. We choose the angle ψ_{new} to maximize the total posterior variance, $\text{tr}(\text{Cov}[\tilde{\mathbf{y}}_{\psi_{new}}])$, also known as uncertainty sampling (Settles, 2009).

3.1 SCORE-BASED DIFFUSION MODELS AND CONDITIONAL SAMPLING

For a data distribution $p_0(\mathbf{x}_0) = p(\mathbf{x})$, a family of distributions $p_t(\mathbf{x}_t)$ can be defined by injecting i.i.d. Gaussian noise to data samples, such that $\mathbf{x}_t = \mathbf{x}_0 + \sigma_t \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, I)$ and σ_t monotonically increasing with respect to time $t \in [0, T]$. The score function $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$ (i.e., gradient of log-probability) can be learned using a neural network via a denoising score matching objective $\mathcal{L}(\theta) = \mathbb{E}_{t, \mathbf{x}_0, \mathbf{x}_t} [\|\mathbf{s}_\theta(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t|\mathbf{x}_0)\|_2^2]$ (Ho et al., 2020).

So far, the trained diffusion model is independent of measurements obtained during the active learning loop. At inference time, our objective is to sample from the posterior distribution by conditioning the diffusion model on X-ray projection measurements; this is done without any additional training. Our setup is similar to that of Chung et al. (2022) and Song et al. (2023). Given the set of current measurements \mathcal{Y} for angles ψ , the goal is to sample from the posterior distribution $p(\mathbf{x}|\mathcal{Y}, \psi)$. The conditional score at time t can be obtained via Bayes’ rule, where the second term still needs to be approximated $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t|\mathcal{Y}, \psi) = \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p_t(\mathcal{Y}, \psi|\mathbf{x}_t)$. To this end, we use a variant of the *Hard Data Consistency* approach proposed by Song et al. (2023). At step t of the reverse diffusion process, we start from our current noisy estimate \mathbf{x}_t , and use Tweedie’s formula $\hat{\mathbf{x}}_0(\mathbf{x}_t) = \mathbf{x}_t + \sigma_t^2 \mathbf{s}_\theta(\mathbf{x}_t, t)$ to get a noiseless estimate $\hat{\mathbf{x}}_0(\mathbf{x}_t)$. We then take several gradient steps solving the minimization problem (2) initialized with $\mathbf{x} = \hat{\mathbf{x}}_0(\mathbf{x}_t)$.

However, instead of fully solving the problem as done in *Hard Data Consistency* and then combining the result linearly with $\hat{\mathbf{x}}_0(\mathbf{x}_t)$, as proposed by Song et al. (2023) (for latent diffusion models), we use early stopping—that is, we perform a predefined, limited number of gradient steps to solve (2). We refer to this approach as *Soft Data Consistency*. Since the minimization is initialized with $\hat{\mathbf{x}}_0(\mathbf{x}_t)$, the resulting estimate $\mathbf{x}_0^*(\mathbf{x}_t)$ retains features of $\hat{\mathbf{x}}_0(\mathbf{x}_t)$, similar to the linear combination in Song et al. (2023), while promoting consistency with the current measurements. This approach avoids convergence to the exact solution, reducing computation time while maintaining solution quality. Moreover, since we use diffusion models in pixel space, we do not require complex scheduling to avoid the overhead of back-propagation through the latent diffusion model decoder, making our method simpler and faster for pixel-space diffusion models while achieving comparable performance to *Hard Data Consistency*.

270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323

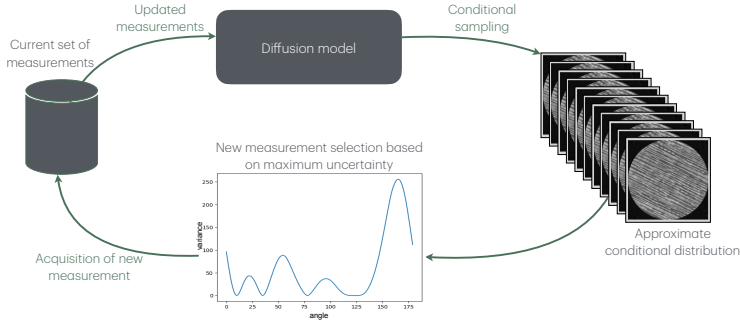


Figure 2: An illustration of the active learning loop using diffusion models to approximate the posterior distribution conditioned on current measurements.

Finally, $\mathbf{x}_0^*(\mathbf{x}_t)$ must be mapped back to the manifold defined by the noisy samples at time t , to go further the reverse diffusion process. To this end, we use the fact that $p(\mathbf{x}_t|\mathbf{x}_0^*(\mathbf{x}_t), \mathcal{Y}, \psi)$ is a tractable Gaussian distribution whose mean is a scaling of $\mathbf{x}_0^*(\mathbf{x}_t)$. For further details and a comparison with Hard Data Consistency see Appendix A.

3.2 SAMPLING-BASED ACTIVE LEARNING

The selection process in active learning usually involves an information criteria (sometimes called acquisition function) which scores the informativeness of each possible measurement. To this end, most methods make use of uncertainty quantification, which, at the same time, poses one of the major challenges in the deep learning setting. The literature proposes a plethora of different acquisition functions (c.f. Settles, 2009; Gal et al., 2017; Ren et al., 2021), depending on the setting (e.g., regression, classification), the statistical model (e.g., linear, deep neural networks), and the learning target (e.g., model identification, PAC). A common approach is to use a Bayesian model and (approximate) information theoretic measures such as mutual information. However, the pure Bayesian approach is often intractable beyond linear models, requiring further approximations.


We formulate the acquisition process in our sampling-based framework. At time step t of the active learning loop, we sample k images $\mathbf{x}_t^1, \dots, \mathbf{x}_t^k \in \mathbb{R}^{d \times d}$ from the (approximate) posterior distribution. We use the conditional diffusion model to obtain the samples, but the formulation is general and works with any generative model. In this light, we can also view the sampled Laplace approximation of Barbano et al. (2022a) in the same framework. Intuitively, the samples $\mathbf{x}_t^1, \dots, \mathbf{x}_t^k$ are consistent with the observation data, but differ in places where there is not enough data to constrain the posterior distribution. Our goal is to acquire additional measurements that differentiate the samples $\mathbf{x}_t^1, \dots, \mathbf{x}_t^k$. We choose the angle ψ_{t+1} that maximizes the posterior total variance,

$$\psi_{t+1} = \arg \max_{\psi \in \Phi} \frac{1}{k} \sum_{i=1}^k \|\mathcal{A}_\psi(\mathbf{x}_t^i) - \mathcal{A}_\psi(\bar{\mathbf{x}}_t)\|^2, \quad \text{where } \bar{\mathbf{x}}_t = \frac{1}{k} \sum_{i=1}^k \mathbf{x}_t^i. \quad (3)$$

The optimization is over the discrete set of angles Φ , therefore requires to apply the forward model $k \cdot |\Phi|$ times. In the tomographic setup, this computation can be batched efficiently on a GPU. Selecting the measurement angle that displays the largest sample variance introduces additional inference constraints in the consecutive round, in a way that effectively reduces the remaining variance in the posterior distribution. This is known as *uncertainty sampling* (Lewis, 1995) and has been analyzed formally in various settings (see, e.g. Settles, 2009; Liu & Li, 2023). We discuss alternative acquisition strategies in Appendix D, however we remark already now that we found no significant difference among the variants we tested. Therefore, based on our evaluation, we recommend Eq. (3) as a simple and yet effective choice.

4 EXPERIMENTS

In our experiments, we closely follow the setup introduced in Section 2. Diffusion active learning is implemented as described in Section 3. We evaluate the proposed approach and several baselines



332
333
334
335

Figure 3: Two cropped and rescaled slices of size 128×128 from each of the test datasets.

336
337
338

on three real-world tomography image datasets as described below. The code is provided in the supplementary and will be released as open-source upon publication.

339 4.1 DATASETS

340
341
342
343
344
345
346
347

Chip Data. Our first dataset is an integrated circuit measured with ptychographic X-ray laminography (PyXL) (Holler et al., 2019b;a). PyXL uses ptychography to scan the sample with a coherent X-ray beam, acquiring diffraction patterns for projection reconstruction at multiple rotation angles. The sample rotation axis is tilted relative to the detector, allowing high-resolution, non-destructive imaging of planar samples across various scales. The dataset features a 3D volume of an integrated circuit, with large metal interconnects in the upper layers and progressively smaller features towards the transistor layer at the bottom.

348
349
350
351
352
353
354
355

Composite Materials. As a second dataset, we use tomography data of a composite material (Auenhammer et al., 2020a;b). This dataset contains 3D tomographic reconstructions of non-crimp fabric reinforced composites, which captures the arrangement and orientation of the fiber bundles and the matrix in which they are embedded. Such composites are extensively used in wind turbine blades and consist of fiber bundles aligned in one direction with stitching yarns to aid in manufacturing and handling. Hence, this dataset was used for the creation of precise finite element models to simulate and analyze the mechanical behavior of the materials, particularly their stiffness and response to fatigue.

356
357
358
359
360
361

Lung Data. The LIDC/IDRI dataset (Armato III et al., 2011) is designed specifically for training and comparing deep learning-based methods for low-dose CT reconstruction, and consists of helical thoracic CT scans. We chose 40,000 CT scan slices from LIDC/IDRI, data from approximately 800 patients, to define the dataset used in this paper. We use the same 40,000 examples chosen by Leuschner et al. (2019).

362
363
364
365
366
367
368

Due to computational constraints, we worked with two image sizes namely, 128×128 and 512×512 pixels. To produce smaller images we used cropping and rescaling of the slices of the reconstructed 3D objects in the datasets. For 128×128 images, we took 256×256 crops from the chip and composite material slices, and then rescaled them to size 128×128 using bilinear interpolation. For the Lung dataset, we rescaled directly each of the slices in the dataset to 128×128 . For the 512×512 images, we simply took crops of size 512×512 from the CT reconstruction datasets described above. All these crops are independent with no overlap. The final images are then split into disjoint train and test datasets.

369
370
371
372
373

In all cases, we use the Radon transform to generate the projection data synthetically using parallel beam geometry. For training the diffusion models, we use data augmentation on the train data by random rotation and scaling (1x to 1.3x). The evaluation on the test set was run on P100 compute nodes (one node per instance). For pre-training the diffusion models, we used a single A100 node.

374 4.2 METHODS

375
376
377

SWAG. Maddox et al. (2019) propose SWA-Gaussian (SWAG), a simple method to enhance uncertainty representation and calibration in deep learning models. SWAG extends Stochastic Weight Averaging (SWA) by capturing the first moment (mean) of the weights using SWA (Izmailov et al.,

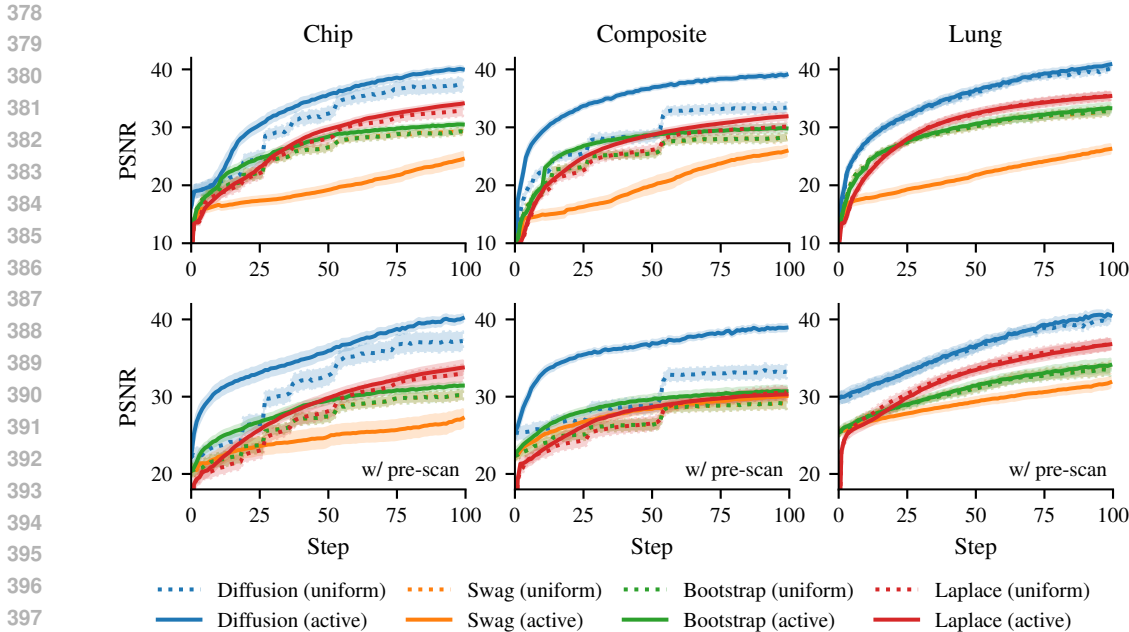


Figure 4: Benchmark for 128×128 images, averaged over 30 data points. Confidence bands show two times standard error. Top (bottom) row correspond to results without (with) low-res pre-scan.

2018), and then modeling the second moment (covariance) using a low-rank plus diagonal approach. This forms an approximate Bayesian posterior distribution over the neural network weights. The approach utilizes the stochastic gradient descent (SGD) trajectory to estimate the posterior’s mean and covariance efficiently, leveraging the observed behavior that SGD iterates approximate a Gaussian distribution in the parameter space of deep networks. As our active learning approach is itself based on samples from the posterior distribution, we directly subsample the SGD trajectory around the mode instead of fitting a Gaussian distribution and resampling, leading to no significant differences in the evaluation.

Bootstrap. A second way of using the inherent randomness of the SGD trajectory is to train multiple instances of the model. Lakshminarayanan et al. (2017) present a scalable method for predictive uncertainty estimation in deep neural networks, referred to as *Deep Ensembles* (or *Bootstrap*). This approach avoids the complexities and computational burdens of Bayesian neural networks by instead employing ensembles of neural networks to approximate uncertainty. Each network in the ensemble is trained independently on the same dataset, leveraging random initializations to induce diversity among the models. Predictive uncertainty is then quantified by the empirical distribution of the ensemble members.

Laplace Approximation and Deep Image Prior. While closed-form solutions of the linear model Eq. (2) are computationally intractable due to the high-dimensional input and observation spaces, Antoran et al. (2023) have proposed a kernelized and optimize-based approximation applicable specifically in the CT setting. This approach was used for Bayesian experimental design (Barbano et al., 2022a) in combination with a linearized deep image prior (DIP) network (Ulyanov et al., 2018). The method introduces a Gaussian surrogate for the total variation (TV) regularizer to preserve Gaussian-linear conjugacy, enabling efficient computation of uncertainty measures. However, the experimental evaluation of Barbano et al. (2022a) is still restricted to a synthetic toy example. Expanding on this framework, Antoran et al. (2022) introduce a scalable sampling-based approach to uncertainty quantification in large-scale linear models. They extend this approach to non-linear parametric models using a linearized Laplace approximation. To calibrate the posterior hyperparameters, they propose an efficient Expectation Maximization scheme for marginal likelihood optimization. Our evaluation uses the reference implementation of Antoran et al. (2023).

Dataset	Diffusion	Swag	Bootstrap	Laplace
Chip	27 [-2, +3]	> 100	80 [-8, +14]	53 [-4, +4]
Composite	15 [-3, +1]	> 100	> 100	65 [-6, +8]
Lung	18 [-2, +3]	> 100	44 [-4, +7]	34 [-3, +4]

Table 1: Number of measurements required to achieve an average score of PSNR of 30 dB using active learning acquisition for the different models and datasets studied in this paper. Ranges are computed using two times standard error. Diffusion active learning achieves up to a $4.3\times$ improvement compared to the Laplace model on the Composite data.

4.3 RESULTS

We compared the performance of several baselines including four different generative models and non-adaptive uniform acquisition (computed sequentially by halving the remaining angle space, *e.g.*, $0^\circ, 90^\circ, 45^\circ, 135^\circ, \dots$). As generative models, we consider diffusion models, SWAG, Bootstrap, and Laplace, and for acquisition we consider a non-adaptive, uniform allocation, and active learning based on (3). We consider other acquisition function in Appendix D, but we found no significant difference among the tested variants. In addition, we consider two settings: with and without *pre-scan*. In the pre-scan setting, there is a previously computed low-resolution scan of the object that can be used as a prior in the reconstruction process. Pre-scanning is a common procedure at synchrotron X-ray beamlines to get a quick estimate of the sample using much lower radiation. All models are conditioned on the pre-scan data in the same way as for the other measurements. Our experiments show that diffusion models with active learning consistently outperforms other methods in terms of Peak Signal-to-Noise Ratio (PSNR) and computational efficiency. We showcase other metrics in Appendix C, which also follow a similar pattern.

Figure 4 summarizes the performance on our three datasets for images of size 128×128 . If no pre-scan is used, diffusion models outperform all other generative models in terms of PSNR with up to $4.3\times$ reduction in the number of measurements needed and, while the gap is dataset-dependent, there is always a clear advantage in using diffusion models to generate tomographic reconstructions. In terms of acquisition functions, we benchmark against uniform acquisition in angular space. While being data-distribution independent, uniform remains on par with active learning acquisition functions for the Lung dataset. This can be attributed to the fact that there are not clear directions of preference for lung images and their features appear isotropic in all directions. For chip and composite materials samples however, active learning acquisitions based in (3) are able to identify the most relevant directions, and clearly outperform the uniform strategy. Table 1 summarizes the gains of diffusion active learning in terms of number of measurements for a target PSNR of 30 dB.

When using pre-scan, the different generative models can make out the rough shape of the object from the beginning of the experiment, which leads into higher quality reconstructions over the whole range of measurements. This in turn provides more information to decide which measurement to take next. This is particularly noticeable for SWAG which becomes competitive with Bootstrap and Laplace. Diffusion models still provide better reconstructions. However, after sufficient number of measurements (*e.g.*, 50 for Lung data and 100 for chip data) there is no clear advantage in using the learned prior.

Figure 5 corroborates our findings with larger images of 512×512 pixels tested on the chip dataset. The results are consistent with the 128×128 case, showcasing that the advantages in the use of both diffusion models and active learning acquisitions are not restricted to low resolution scans. Figure 5 further shows computation times for the different methods. At 512×512 resolution, DAL requires less than two minutes per step; significantly faster than the Laplace approach by Barbano et al. (2022a). In the context of long image acquisition times of up to several days (Aidukas et al., 2024), the proposed method can provide a significant advantage by reducing the data requirements.

5 CONCLUSION

We introduced *Diffusion Active Learning* (DAL), a novel framework that combines generative diffusion models with active learning for CT reconstruction. Our experimental evaluation showcases

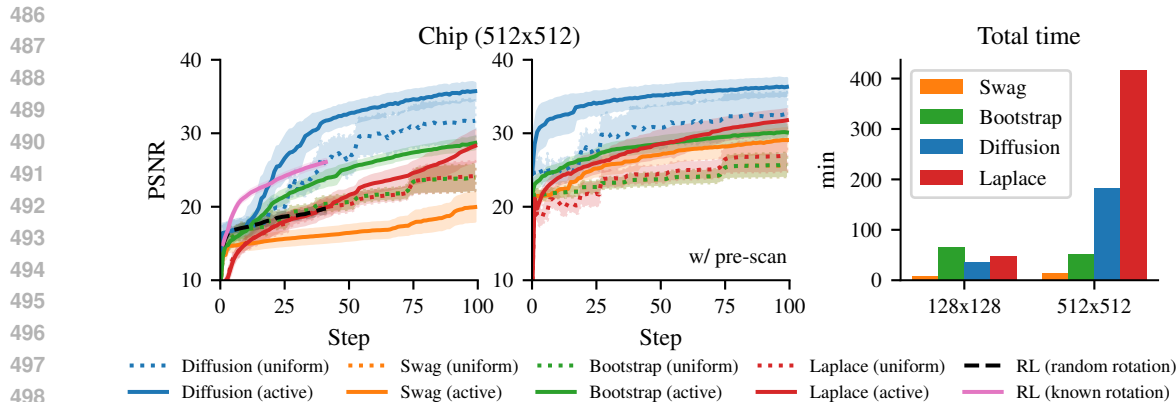


Figure 5: Left: Benchmark for 512×512 images averaged over 10 data points. Confidence bands show two times standard error. Right: Average running time for 100 steps of active learning using different models.

significant gains in improving the reconstruction quality with fewer samples compared to static uniform allocation and several baselines from prior works. The reduction in X-ray dose and cost savings achieved by diffusion active learning have significant implications for practical CT applications. These improvements can lead to wider adoption in scientific imaging and material sciences at synchrotron facilities, due to shorter X-ray beam-time allocation, faster experiments and enhanced quality. As noticed in our results, the gains provided by our method are dataset dependent, showing improvements in particular with highly structured images. Lastly, we remark that the DAL framework applies to any (differentiable) forward process, and as such can be applied to other setups such as MRI or ptychographic reconstruction methods.

As often, the achieved gains come with trade-offs. First, training the diffusion model requires a sufficient training data in the domain of interest. This is a reasonable assumption for many applications where prior reconstruction data is readily available. A possible way forward is to pre-train large foundation models on a variety of tomographic images, or even use pre-trained models such as Stable Diffusion (Podell et al., 2023) for fine-tuning on much smaller datasets (Hu et al., 2021).

Second, while diffusion models are computationally more costly than iterative reconstruction or filtered back-projection, our methods run on an NVIDIA P100 with 16GB of memory for resolutions up to 1024×1024 , displaying reasonable scaling properties with many possibilities for further improvements. In addition, image acquisition time in micro- and nano-tomography can take up to several days (Aidukas et al., 2024), providing a clear use-case where the additional computational effort can be justified for improving sample efficiency.

Lastly, the key advantage of a learned prior comes at the cost of introducing reconstruction bias in cases where the measured sample is not contained in the training distribution; this is in particular relevant when the goal is to detect small deviations or defects in the samples. The same reason might prohibit the use of the proposed approach in high-stakes setting such as medical applications, or at least not without enhanced safety measures. The work of Barbano et al. (2023) takes a first steps towards working with out-of-distribution samples. More generally, in a real-world tomography experimental setup, many additional challenges arise, including sample alignment and measurement noise. Despite our effort of making the evaluation more realistic, there remains a sim-to-real gap to be addressed in future works, *e.g.*, by considering distribution shifts in the testing distribution.

REFERENCES

Tomas Aidukas, Nicholas W Phillips, Ana Diaz, Emiliya Poghosyan, Elisabeth Müller, AFJ Levi, Gabriel Aeppli, Manuel Guizar-Sicairos, and Mirko Holler. High-performance 4-nm-resolution x-ray tomography using burst ptychography. *Nature*, 632(8023):81–88, 2024.

- 540 Javier Antoran, Shreyas Padhy, Riccardo Barbano, Eric Nalisnick, David Janz, and José Miguel
541 Hernández-Lobato. Sampling-based inference for large linear models, with application to lin-
542 earised laplace. In *The Eleventh International Conference on Learning Representations*, 2022.
543
- 544 Javier Antoran, Riccardo Barbano, Johannes Leuschner, José Miguel Hernández-Lobato, and Bangti
545 Jin. Uncertainty estimation for computed tomography with a linearised deep image prior. *Trans-*
546 *actions on Machine Learning Research*, 2023.
- 547 Samuel G Armato III, Geoffrey McLennan, Luc Bidaut, Michael F McNitt-Gray, Charles R Meyer,
548 Anthony P Reeves, Binsheng Zhao, Denise R Aberle, Claudia I Henschke, Eric A Hoffman,
549 et al. The lung image database consortium (lidc) and image database resource initiative (idri): a
550 completed reference database of lung nodules on ct scans. *Medical physics*, 38(2):915–931, 2011.
551
- 552 Robert M Auenhammer, Lars P Mikkelsen, Leif E Asp, and Brina J Blinzler. Dataset of non-crimp
553 fabric reinforced composites for an x-ray computer tomography aided engineering process. *Data*
554 *in Brief*, 33:106518, 2020a.
- 555 Robert M. Auenhammer, Lars P. Mikkelsen, Leif E. Asp, and J. Blinzler, Brina. Automated X-
556 ray computer tomography segmentation method for finite element analysis of non-crimp fabrics
557 reinforced composites, September 2020b. URL [https://doi.org/10.5281/zenodo.](https://doi.org/10.5281/zenodo.3830790)
558 3830790.
- 559 Daniel Otero Bager, Johannes Leuschner, and Maximilian Schmidt. Computed tomography recon-
560 struction using deep image prior and learned reconstruction methods. *Inverse Problems*, 36(9):
561 094004, 2020.
562
- 563 Tim Bakker, Herke van Hoof, and Max Welling. Experimental design for mri by greedy policy
564 search. *Advances in Neural Information Processing Systems*, 33:18954–18966, 2020.
565
- 566 Maria-Florina Balcan, Steve Hanneke, and Jennifer Wortman Vaughan. The true sample complexity
567 of active learning. *Machine learning*, 80:111–139, 2010.
- 568 Riccardo Barbano, Johannes Leuschner, Javier Antorán, Bangti Jin, and José Miguel Hernández-
569 Lobato. Bayesian experimental design for computed tomography with the linearised deep image
570 prior. *arXiv preprint arXiv:2207.05714*, 2022a.
571
- 572 Riccardo Barbano, Johannes Leuschner, Maximilian Schmidt, Alexander Denker, Andreas Haupt-
573 mann, Peter Maass, and Bangti Jin. An educated warm start for deep image prior-based micro ct
574 reconstruction. *IEEE Transactions on Computational Imaging*, 8:1210–1222, 2022b.
- 575 Riccardo Barbano, Alexander Denker, Hyungjin Chung, Tae Hoon Roh, Simon Arridge, Peter
576 Maass, Bangti Jin, and Jong Chul Ye. Steerable conditional diffusion for out-of-distribution adap-
577 tation in imaging inverse problems. *arXiv preprint arXiv:2308.14409*, 2023.
578
- 579 Samuel Budd, Emma C Robinson, and Bernhard Kainz. A survey on active learning and human-in-
580 the-loop deep learning for medical image analysis. *Medical image analysis*, 71:102062, 2021.
581
- 582 Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical*
583 *science*, pp. 273–304, 1995.
- 584 Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul
585 Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh Interna-*
586 *tional Conference on Learning Representations*, 2022.
587
- 588 S.R. Deans. *The Radon Transform and Some of Its Applications*. A Wiley-Interscience publication.
589 Wiley, 1983. ISBN 9780471898047.
- 590 Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data.
591 In *International conference on machine learning*, pp. 1183–1192. PMLR, 2017.
592
- 593 Hideitsu Hino and Shinto Eguchi. Active learning by query by committee with robust divergences.
Information Geometry, 6(1):81–106, 2023.

- 594 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*
595 *neural information processing systems*, 33:6840–6851, 2020.
- 596
- 597 Mirko Holler, Michal Odstrcil, and Manuel Guizar-Sicairos. Nanolaminography dataset: Three-
598 dimensional imaging of integrated circuits with a macro to nanoscale zoom, August 2019a. URL
599 <https://doi.org/10.5281/zenodo.2657340>.
- 600 Mirko Holler, Michal Odstrcil, Manuel Guizar-Sicairos, Maxime Lebugle, Elisabeth Müller, Simone
601 Finizio, Gemma Tinti, Christian David, Joshua Zusman, Walter Unglaub, et al. Three-dimensional
602 imaging of integrated circuits with macro-to nanoscale zoom. *Nature Electronics*, 2(10):464–470,
603 2019b.
- 604 Malcolm R Howells, Tobias Beetz, Henry N Chapman, C Cui, JM Holton, CJ Jacobsen, J Kirz, Enju
605 Lima, Stefano Marchesini, Huijie Miao, et al. An assessment of the resolution limitation due to
606 radiation-damage in x-ray diffraction microscopy. *Journal of electron spectroscopy and related*
607 *phenomena*, 170(1-3):4–12, 2009.
- 608 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
609 and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint*
610 *arXiv:2106.09685*, 2021.
- 611
- 612 Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wil-
613 son. Averaging weights leads to wider optima and better generalization. *arXiv preprint*
614 *arXiv:1803.05407*, 2018.
- 615 Kyong Hwan Jin, Michael Unser, and Kwang Moo Yi. Self-supervised deep active accelerated mri.
616 *arXiv preprint arXiv:1901.04547*, 2019.
- 617
- 618 Avinash C Kak and Malcolm Slaney. *Principles of computerized tomographic imaging*. SIAM,
619 2001.
- 620 Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive
621 uncertainty estimation using deep ensembles. *Advances in neural information processing systems*,
622 30, 2017.
- 623 Johannes Leuschner, Maximilian Schmidt, and Daniel Otero Baguer. Lodopab-ct dataset, October
624 2019. URL <https://doi.org/10.5281/zenodo.3384092>.
- 625
- 626 David D Lewis. A sequential algorithm for training text classifiers: Corrigendum and additional
627 data. In *Acm Sigir Forum*, volume 29, pp. 13–19. ACM New York, NY, USA, 1995.
- 628 Disen Li, Limin Ma, Jining Li, Shouliang Qi, Yudong Yao, and Yueyang Teng. A comprehensive
629 survey on deep learning techniques in ct image quality improvement. *Medical & Biological*
630 *Engineering & Computing*, 60(10):2757–2770, 2022.
- 631 Shang Liu and Xiaocheng Li. Understanding uncertainty sampling. *arXiv preprint*
632 *arXiv:2307.02719*, 2023.
- 633
- 634 David Lowell, Zachary C Lipton, and Byron C Wallace. Practical obstacles to deploying active
635 learning. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language*
636 *Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-*
637 *IJCNLP)*, pp. 21–30, 2019.
- 638 Wesley J Maddox, Pavel Izmailov, Timur Garipov, Dmitry P Vetrov, and Andrew Gordon Wilson.
639 A simple baseline for bayesian uncertainty in deep learning. *Advances in neural information*
640 *processing systems*, 32, 2019.
- 641
- 642 Arshi Parvaiz, Muhammad Anwaar Khalid, Rukhsana Zafar, Huma Ameer, Muhammad Ali, and
643 Muhammad Moazam Fraz. Vision transformers in medical computer vision—a contemplative
644 retrospection. *Engineering Applications of Artificial Intelligence*, 122:106126, 2023.
- 645 Luis Pineda, Sumana Basu, Adriana Romero, Roberto Calandra, and Michal Drozdal. Active
646 mr k-space sampling with reinforcement learning. In *Medical Image Computing and Computer*
647 *Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8,*
2020, Proceedings, Part II 23, pp. 23–33. Springer, 2020.

- 648 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe
649 Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image
650 synthesis. *arXiv preprint arXiv:2307.01952*, 2023.
- 651
- 652 Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zihui Li, Brij B Gupta, Xiaojiang Chen,
653 and Xin Wang. A survey of deep active learning. *ACM computing surveys (CSUR)*, 54(9):1–40,
654 2021.
- 655
- 656 Carlos Riquelme, Mohammad Ghavamzadeh, and Alessandro Lazaric. Active learning for accurate
657 estimation of linear models. In *International Conference on Machine Learning*, pp. 2931–2939.
658 PMLR, 2017.
- 659
- 660
- 661 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomed-
662 ical image segmentation. In *Medical image computing and computer-assisted intervention–*
663 *MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceed-*
664 *ings, part III 18*, pp. 234–241. Springer, 2015.
- 665
- 666 Burr Settles. Active learning literature survey. 2009.
- 667
- 668 H Sebastian Seung, Manfred Opper, and Haim Sompolinsky. Query by committee. In *Proceedings*
669 *of the fifth annual workshop on Computational learning theory*, pp. 287–294, 1992.
- 670
- 671
- 672 Ziju Shen, Yufei Wang, Dufan Wu, Xu Yang, and Bin Dong. Learning to scan: A deep reinforcement
673 learning approach for personalized scanning in ct imaging. *Inverse Problems & Imaging*, 16(1),
674 2022.
- 675
- 676 Bowen Song, Soo Min Kwon, Zecheng Zhang, Xinyu Hu, Qing Qu, and Liyue Shen. Solv-
677 ing inverse problems with latent diffusion models via hard data consistency. *arXiv preprint*
678 *arXiv:2307.08123*, 2023.
- 679
- 680 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv*
681 *preprint arXiv:2010.02502*, 2020.
- 682
- 683
- 684 Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging
685 with score-based generative models. In *International Conference on Learning Representations*,
686 2021.
- 687
- 688 Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the*
689 *IEEE conference on computer vision and pattern recognition*, pp. 9446–9454, 2018.
- 690
- 691
- 692 Tianyuan Wang, Felix Lucka, and Tristan van Leeuwen. Sequential experimental design for x-ray
693 ct using deep reinforcement learning. *arXiv preprint arXiv:2307.06343*, 2023.
- 694
- 695 Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zeyd Huang, Matthew J Muckley,
696 Aaron Defazio, Rachel Stern, Philip Johnson, Michael Bruno, et al. fastmri: An open dataset and
697 benchmarks for accelerated mri. *arXiv preprint arXiv:1811.08839*, 2018.
- 698
- 699 Zizhao Zhang, Adriana Romero, Matthew J Muckley, Pascal Vincent, Lin Yang, and Michal
700 Drozdal. Reducing uncertainty in undersampled mri reconstruction with active acquisition.
701 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
2049–2058, 2019.




Figure 6: Ablation of soft (Square) vs hard data consistency (disk) on both Lung and Chip datasets. We evaluate the performance with different sparsity. While the quality is comparable, our approach requires a fraction of the time. We use SGD consistency with a fixed number of steps, each with a fixed batch size, our running time is then predictable and stable, without any loss in quality.

A BACKGROUND ON SCORE-BASED DIFFUSION MODELS AND CONDITIONAL SAMPLING

We succinctly review the fundamentals of diffusion models, namely the formulation of denoising diffusion probabilistic model (DDPM) (Ho et al., 2020). The forward diffusion process incrementally adds Gaussian noise to the data. Let $\mathbf{x}_0 \sim p(\mathbf{x})$ denote the initial data sample, and \mathbf{x}_t the data at time step t . The forward noising process can be described by the following stochastic differential equation (SDE):

$$d\mathbf{x}_t = -\frac{\beta_t}{2}\mathbf{x}_t dt + \sqrt{\beta_t}d\mathbf{w}, \quad (4)$$

where β_t is the noise scheduler, and \mathbf{w} represents the standard Wiener process.

The reverse process aims to recover the original data from the noised version by reversing the SDE:

$$d\mathbf{x}_t = \left[\frac{\beta_t}{2}\mathbf{x}_t - \beta_t \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) \right] dt + \sqrt{\beta_t}d\mathbf{w}_t, \quad (5)$$

Here, $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$ is the score function, which is the gradient of the log probability density at time t . We train a diffusion model $\mathbf{s}_\theta(\mathbf{x}_t, t)$ to approximate the true score function. The training objective is typically:

$$\mathcal{L}(\theta) = \mathbb{E}_{t, \mathbf{x}_0, \mathbf{x}_t} \left[\|\mathbf{s}_\theta(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t | \mathbf{x}_0)\|_2^2 \right]. \quad (6)$$

Once the score function is learned, data samples can be generated by solving the reverse SDE using numerical methods such as Euler-Maruyama or more sophisticated solvers.

In order to improve the sampling speed of DDPM, Song et al. (2020) proposed the denoising diffusion implicit model (DDIM) which defines the diffusion process as a non-Markovian process. A crucial step to achieve this is to notice that one can predict a noiseless variant of $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ from \mathbf{x} using Tweedie’s formula

$$\hat{\mathbf{x}}_0(\mathbf{x}_t) = \frac{1}{\alpha_t} (\mathbf{x}_t + \sqrt{1 - \alpha_t} \mathbf{s}_\theta(\mathbf{x}_t, t)), \quad (7)$$

where $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$.

Solving Inverse Problems using Score-Based Diffusion Models. Scientific inverse problems like CT, where we only have access to partial information due to the limited number of observations, are inherently ill-posed and hence, no unique reconstruction of \mathbf{x} is possible.

To address this problem, we learn a prior $p(\mathbf{x})$ from the training set, and sample from the posterior distribution $p_t(\mathbf{x} | \mathcal{Y}, \psi)$. To achieve this goal, Chung et al. (2022) introduced Diffusion Posterior Sampling. Using a pre-trained score-based model as a prior, they modified (5) and obtain a reverse diffusion process to sample from the posterior distribution:

$$d\mathbf{x}_t = \left[\frac{\beta_t}{2}\mathbf{x}_t - \beta_t (\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p_t(\mathcal{Y}, \psi | \mathbf{x}_t)) \right] dt + \sqrt{\beta_t}d\mathbf{w}_t, \quad (8)$$

756 where they use the fact that

$$757 \quad \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t|\mathcal{Y}, \psi) = \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p_t(\mathcal{Y}, \psi|\mathbf{x}_t). \quad (9)$$

760 Using (7) allows to use the forward operator of the inverse problem on $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ to obtain an estimate
761 $\hat{\mathcal{Y}}_t$ of the measurements. By defining a loss between $\hat{\mathcal{Y}}_t$ and the true measurements \mathcal{Y} , we can
762 back-propagate the error to \mathbf{x}_t and obtain an approximation of $\nabla_{\mathbf{x}_t} \log p_t(\mathcal{Y}, \psi|\mathbf{x}_t)$.

763 Based on these ideas, Song et al. (2023) proposed the use latent diffusion models, instead of pixel-
764 space diffusion models, and introduced a variant of the conditional sampling. Additionally, they
765 proposed Hard Data Consistency, which consists in using (7) to obtain an estimate $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ and then
766 solve completely the following optimization problem initialized with $\hat{\mathbf{x}}_0(\mathbf{x}_t)$:

$$767 \quad \mathbf{x}_0^*(\mathcal{Y}, \psi) \in \arg \min_{\mathbf{x} \in \mathbb{R}^{d \times d}} \sum_{\psi \in \Psi} \|\mathcal{A}_\psi(\mathbf{x}) - \mathbf{y}_\psi\|_2^2 \quad (10)$$

770 where ψ is the current set of angles measured, and \mathcal{A}_ψ is the Radon transform with angle ψ . Notice
771 that if the inverse problem is ill-posed, the initialization and the optimization algorithm determine
772 the value of $\mathbf{x}_0^*(\mathcal{Y}, \psi)$. Finally, one needs to map $\mathbf{x}_0^*(\mathcal{Y}, \psi)$ back to the manifold defined by the
773 noisy samples at time t , to go further with the reverse diffusion process. To this end, Song et al.
774 (2023) proposed *Stochastic Encoding*, which uses the fact that $p(\mathbf{x}_t|\mathbf{x}_0^*(\mathcal{Y}, \psi), \mathcal{Y}, \psi)$ is a tractable
775 Gaussian distribution with mean being a scaling of $\mathbf{x}_0^*(\mathcal{Y}, \psi)$. However, since the measurements
776 might be noisy, and since we have an ill-posed problem, using solely $\mathbf{x}_0^*(\mathcal{Y}, \psi)$ leads to noisy image
777 reconstructions. Thus, they propose a variance reduction by taking a linear combination between
778 $\mathbf{x}_0^*(\mathcal{Y}, \psi)$ and $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ before mapping it back, a method which they coined *ReSample*.

780 **Soft Data Consistency and Early Stopping.** Our setup is similar to that of Chung et al. (2022)
781 and Song et al. (2023). We first pre-train a diffusion model on domain-specific tomogram data
782 that captures the underlying distribution of the desired application. We then condition on the set of
783 current measurements, and produce samples of the posterior distribution using an approach similar
784 to that of *Hard Data Consistency* (Song et al., 2023). More specifically, at step t of the reverse
785 diffusion process, we start from our current estimate \mathbf{x}_t , and use Tweedie’s formula (7) to get a
786 noiseless estimate $\hat{\mathbf{x}}_0(\mathbf{x}_t)$. We then take several *consistency* gradient steps solving the following
787 minimization problem initialized with $\mathbf{x} = \hat{\mathbf{x}}_0(\mathbf{x}_t)$,

$$788 \quad \underset{\mathbf{x}}{\text{minimize}} \sum_{\psi \in \Psi} \|\mathcal{A}_\psi(\mathbf{x}) - \mathbf{y}_\psi\|_2^2. \quad (11)$$

790 However, instead of solving the problem completely as in Hard Data Consistency, and then taking a
791 linear combination with $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ as done in ReSample proposed by Song et al. (2023), we instead take
792 a few steps of Stochastic Gradient Descent (SGD), effectively applying early stopping. Since the
793 minimization problem is initialized with $\hat{\mathbf{x}}_0(\mathbf{x}_t)$, this effectively ends with an estimate that retains
794 features of $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ while encouraging consistency with the current measurements. Moreover, it
795 reduces the computational cost and allows us to use data consistency in each step of the reversed
796 diffusion process. This is not the case with ReSample, where they have to pick a small subset of
797 time steps in which to perform the consistency optimization. We finally use Stochastic Encoding to
798 map back our estimate to the manifold defined by the noisy samples at time t and continue with the
799 reversed diffusion process.

800 Figure 6 shows a comparison of our method with Hard Data Consistency (Song et al., 2023) and
801 Diffusion Posterior Sampling Song et al. (2023). Since we use SGD for consistency steps with a
802 fixed batch-size, the running time of our approach is the same regardless of the number of angles
803 sampled. This is not the case for Hard Data Consistency, where they implement full batch updates
804 in each update, and where they run until a certain specified threshold has been reached with the loss
805 in equation (11). This takes longer and longer as more angles are sampled. Moreover, this training
806 has higher variance as can be seen in Figure 6. In contrast, we have a fixed number of consistency
807 gradient steps after each diffusion step, hence the running time of our algorithm is predictable and
808 constant.

809 In their implementation, Song et al. (2023) notices already that they could not apply hard consistency
over the entire diffusion process. They therefore partitioned the time in 3 equal sections. In the first,

810 they perform no conditioning, in the second, they perform the same update that we do; a fixed
 811 number of gradient steps per diffusion step. Finally, it is only in their third phase that they introduce
 812 the hard consistency. This is how we implemented it for our comparison. While they implemented it
 813 for latent diffusion models, some of their ideas can be repurposed to work with pixel-space diffusion
 814 models, but they needed a strong adaptation that we presented in this paper.

817 B DIFFUSION ACTIVE LEARNING FOR MRI

818
 819 MRI operates within the same framework described in Equation (1), with $\mathbf{y}_\psi = \mathcal{A}_\psi(\mathbf{x}) + \epsilon$, where
 820 x is a complex value object, $\mathcal{A}_\psi(\mathbf{x})$ is a forward model consisting of taking the Fourier transform,
 821 which produces the k -space of \mathbf{x} , followed by a masking that takes a single row (or column) specified
 822 by the index ψ . The set of possible measurements is then the set \mathcal{Y}_ψ of all possible rows of the k -
 823 space of \mathbf{x} . Assuming a Gaussian distribution of the noise ϵ , the reconstruction problem can be
 824 solved using maximum-likelihood inference given by Equation 2.

825 Given the similarity between this setup and the CT setup studied in the main body of this paper, we
 826 extend our framework to work with MRI reconstructions using the formulation of Diffusion Active
 827 Learning and Algorithm 1.

828
 829 **Datasets.** We benchmarked our algorithm with the FastMRI knee dataset (Zbontar et al., 2018),
 830 which is a large-scale open dataset designed to accelerate research in magnetic resonance imaging
 831 (MRI) reconstruction using machine learning. It consists of raw k -space data and fully-sampled
 832 ground truth images, enabling both supervised and unsupervised training. We focus on knee scans
 833 acquired using single coils. We trained our diffusion model with the provided train dataset and tested
 834 it on the disjoint test set.

835
 836 **Methods.** We compare against the same methods introduced in the main body of the paper for the
 837 CT setting, i.e., SWAG and Bootstrap. Laplace is not included as we were relying on the author’s
 838 implementation which does not support the MRI model at this time. Each of the baselines works
 839 with active acquisition using uncertainty sampling (maximizing the ‘variance’ score), and a non-
 840 adaptive baseline that selects from a uniform pool of angles (‘uniform’) or columns corresponding
 841 to frequencies in increasing order (‘low to high’).

842
 843 **Results.** Figure 7 summarizes our results for the FastMRI dataset. As generative models, we
 844 consider diffusion models, SWAG, and Bootstrap; and for acquisition, we consider two non-adaptive
 845 allocations, *uniform* and *low to high*, explained below, and the active learning allocation based on
 846 Equation (3). On each acquisition step, both non-adaptive allocations choose the non-measured
 847 row closest to the center from a pool of predefined available rows. For an AL process with k steps,
 848 the pool of rows for the uniform allocation consists of k equispaced rows, while for the low to
 849 high allocation, the pool consists of the k rows closest to the center, i.e., the ones with the lowest
 850 frequencies.

851
 852 As in the case of CT, the diffusion model outperforms all other generative models in terms of PSNR
 853 with up to $2\times$ reduction in the number of sampled rows needed to achieve the same PSNR.

854
 855 In MRI, similar to the pre-scan in CT, it is common to pre-scan the first few rows with lower fre-
 856 quencies. We examine two cases: 2 and 30 rows preselected closest to the center. As anticipated,
 857 the preselected rows with lower frequencies enable the generative model to obtain more accurate
 858 predictions through various sparsity levels. While diffusion models retain the lead, the other models
 catch up more quickly.

859
 860 While we believe the performance of Active Learning strategies is also data-distribution dependent,
 861 we can see that for FastMRI, DAL clearly outperforms the non-adaptive allocation, which was not
 862 the case in CT medical images of the Lung dataset, though the gap is narrow with low-to-high
 863 allocations. The uniform allocation, however, struggles to generate meaningful reconstructions with
 only two preselected rows, likely meaning that some important low frequencies were skipped. For
 30 preselected rows, uniform is more competitive, but is eventually clearly outperformed.

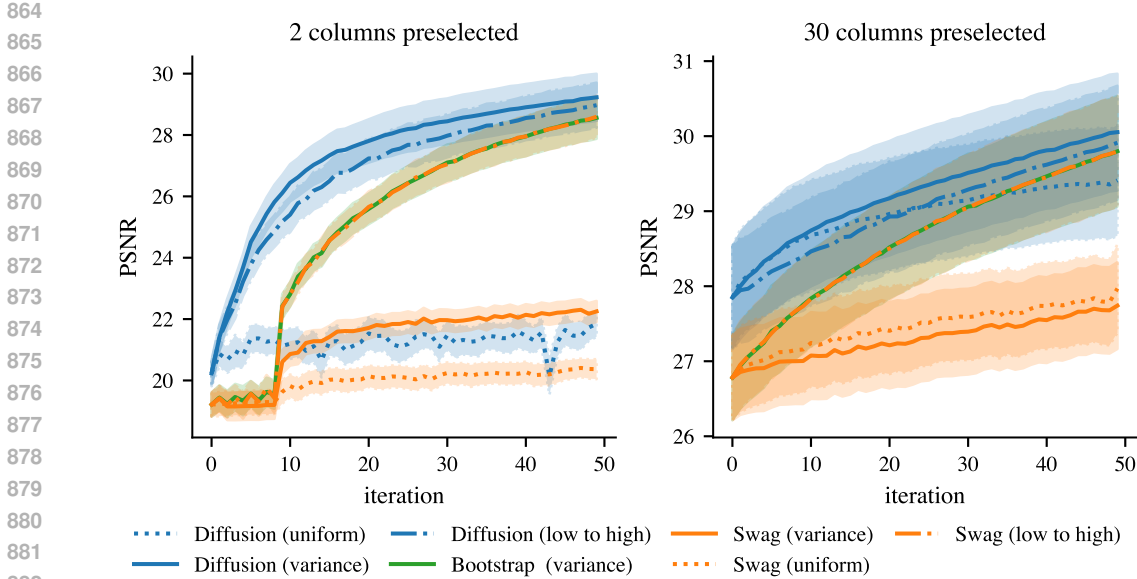


Figure 7: Benchmark over 26 test images of the FastMRI dataset. Confidence bands show two times standard error. The reconstructions of Bootstrap and Swag are equivalent to solving the inverse Fourier problem given the observed columns in k-space; for active learning selection (‘variance’), a separate copy of the model is used for random sampling. Sampling 50 rows corresponds to an acceleration of $6.38\times$ for 2 columns and $4.1\times$ for 30 columns.

Fig. 17 and Fig. 18 show qualitative results of the reconstructions. Here we can see that the diffusion model obtains a sharper image with as few as 10 adaptively chosen measurements (plus two pre-selected columns), while the other generative models struggle to obtain a meaningful reconstruction.

Accelerated MRI diffusion inference.

While diffusion inference can be done by optimizing Equation (2) as in the CT case using gradient descent steps, we developed an accelerated formulation for MRI described as follows. We recall first the optimization objective of Equation (2):

$$\text{minimize}_{\mathbf{x} \in \mathbb{R}^{d \times d}} \sum_{\psi \in \Psi} \|\mathcal{A}_\psi(\mathbf{x}) - \mathbf{y}_\psi\|_2^2.$$

Notice that a solution to the optimization is given by computing the k-space of \mathbf{x} and then replacing all the rows in \mathcal{Y}_ψ by the measurements $\{\mathbf{y}_\psi | \psi \in \Psi\}$. That is, we perform a *Fourier in-painting*, after which we apply the inverse Fourier transform to go back to pixel space and obtain our estimate $\hat{\mathbf{x}}$ being a minimizer of Equation (2). As shown in Figure 8, using this formulation leads to an up to $4\times$ improvement in inference time, and hence, also similar gains in the entire loop of DAL.

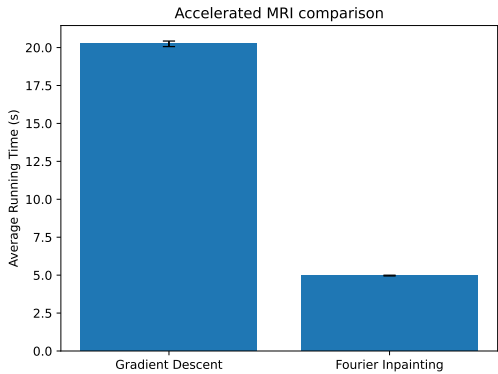



Figure 8: Compares the average running time of diffusion MRI inferences: gradient descent updates vs Fourier in-painting. Fourier in-painting achieves a $4\times$ acceleration. Results are averaged over 30 independent inferences, and bars show the standard error.



928 Figure 9: A diffusion reconstruction for a sample of the FastMRI dataset and its comparison with
929 the standard Fourier reconstruction.


931 C EXPERIMENTS: ADDITIONAL 932 DETAILS

934 C.1 COMPUTING PARAMETERS

935 To sample from our diffusion models, we use 50 steps of the reverse diffusion process out of 1000
936 using the DDIM scheduler. For the soft data consistency, we use 50 gradient steps of the loss in
937 equation (2). A single reconstruction in an A100 chip takes around 10 seconds, while 10 reconstruc-
938 tions take up to 25 seconds using less than 16GB of memory for images of size 512x512. Figure 5
939 (right) further shows the average running time of an entire active learning loop with 100 iterations
940 for image sizes 512x512 and 128x128 on a P100 GPU.
941

942 C.2 DATASET

943 A visualization of the three datasets introduced in Section 4 is given in Fig. 10 below.
944



958 Figure 10: A 1000x1000 pixels crop sample of each of our three datasets, each before rescaling.
959

960 C.3 ABLATION: RECONSTRUCTION METHOD AND ANGLE SELECTION

961 An interesting question is if the advantage of diffusion active learning is due to better angle selection
962 or due to better reconstruction, or a combination of both. To answer this question, we re-evaluated
963 the sequence of angles selected by the Bootstrap model using uncertainty sampling, and performed
964 the reconstruction using the diffusion model. As expected, this improves the PSNR of the recon-
965 struction. However, the resulting image quality as measured by PSNR is still significantly below
966 the quality achieved by diffusion active learning. This showcases that the best reconstruction is
967 achieved by the combination of both: The diffusion model captures the data distribution, and the
968 angles selected by diffusion active learning exploit the data distribution in a way that a distribution-
969 independent approach cannot; Bootstrap uses only information obtained from the current sample,
970 and therefore intuitively cannot “reason” about the posterior distribution as the diffusion model
971

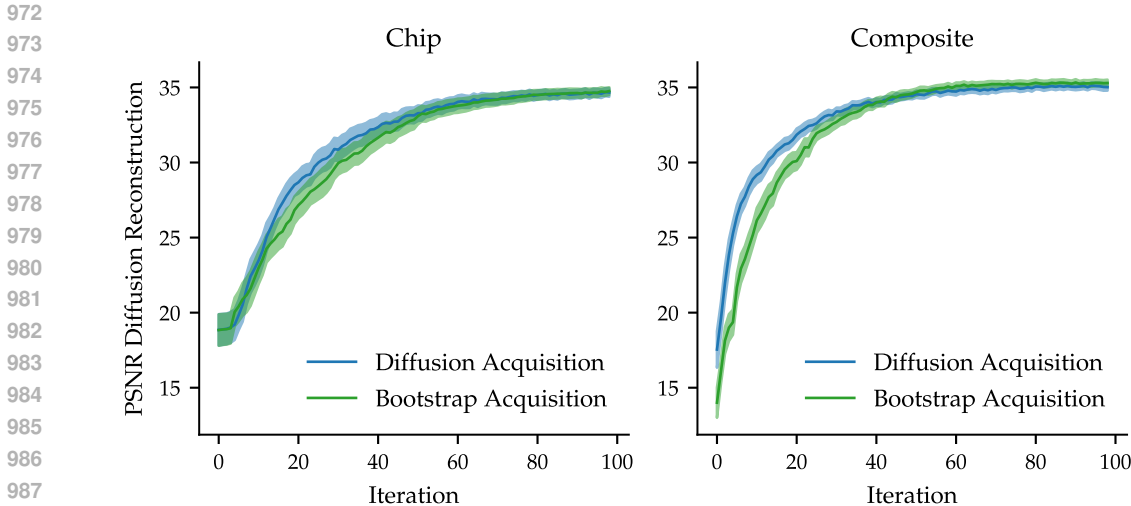


Figure 11: The plots show PSNR using Diffusion Posterior Sampling as a reconstruction method (using soft data consistency), with a sequence of angles selected either by Diffusion Active Learning or the Bootstrap uncertainty sampling. While eventually, both diffusion reconstructions reach the same PSNR score, for fewer angles there is a clear advantage of using Diffusion Active Learning for angle selection. This showcases that the selected sequence of angles is specific to the reconstruction method. In particular, the learned diffusion prior is not only used for better reconstruction, but is also used to choose a sequence of measurements that leads to better reconstructions.

does. Note also that the gains are not purely from better angle selection, as there is a significant gap between uniform and active selection on the composite and chip data.

C.4 ADDITIONAL EVALUATION METRICS

We provide plots for PSNR, RMSE and SSIM metrics in Fig. 12 below. Note that the trend for all metrics is the same, showing that active learning significantly outperforms uniform acquisition for the Chip and Composite data, and achieving similar construction quality for the Lung data set.

C.5 VISUALIZATION OF RECONSTRUCTIONS

To allow for a qualitative comparison, we provide the reconstructed images at steps 1,10,20 and 30 below in Figs. 13 to 16, for the active learning acquisition and without pre-scan. Note that visually the reconstruction quality of the diffusion model is already superior after 10 steps, and displaying intricate details of the 512×512 reconstruction with as few as 20 projections.

D SAMPLING-BASED ACTIVE LEARNING

D.1 ACQUISITION FUNCTIONS

We provide further details on the acquisition process of active learning, and discuss several acquisition function proposed in the literature (e.g., Settles, 2009) that can be applied in the tomographic reconstruction setting.

Entropy and Mutual Information Assume, for a moment, that at step t of the acquisition process with observations \mathcal{Y}_t for angles ψ_t , we have an exact and fully specified Bayesian model with a posterior over images, $p_t(\mathbf{x}) := p(\mathbf{x}|\mathcal{Y}_t, \psi_t) \propto p(\mathbf{x})p(\mathcal{Y}|\mathbf{x}, \psi_t)$. The observation likelihood $p(\mathcal{Y}|\mathbf{x}, \psi)$ is defined by Eq. (1) for taking a new measurement at angle ψ . As our goal is to reconstruct the image \mathbf{x} , a natural acquisition target is the conditional mutual information between the

reconstruction \mathbf{x} and the observation conditioned on the choice of angle $\psi_{t+1} = \psi$,

$$\psi_{t+1} = \arg \max_{\psi \in \Phi} \mathbb{I}_t(\mathbf{x}; \mathcal{Y}_{t+1} | \psi_{t+1} = \psi). \quad (12)$$

The subscript t indicates conditioning on the observed data filtration, $\mathbb{I}_t(\mathbf{x}; \mathcal{Y}_{t+1} | \psi_{t+1} = \psi) := \mathbb{I}(\mathbf{x}; \mathcal{Y}_{t+1} | \psi_{t+1} = \psi, \mathcal{Y}_t, \psi_t)$. Rewriting the mutual information using the entropy,

$$\mathbb{I}_t(\mathbf{x}; \mathcal{Y}_{t+1} | \psi_{t+1} = \psi) = \mathbb{H}_t(\mathcal{Y}_{t+1} | \psi_{t+1} = \psi) - \mathbb{H}_t(\mathcal{Y}_{t+1} | \mathbf{x}, \psi_{t+1} = \psi) \quad (13)$$

From this, we note that maximizing the mutual information Eq. (12) is the same as maximizing the posterior entropy of the observation in cases where the observation distribution is independent of the angle (e.g. homoscedastic Gaussian noise models). In practice, computing the mutual information or the entropy is computationally challenging except for in special cases such as a Gaussian conjugate model. Assuming that the posterior distribution is $\mathcal{N}(\mathbf{x}_t, \Sigma_t)$ with covariance $\Sigma_t \in \mathbb{R}^{(d \times d)^2}$ and the likelihood is also Gaussian, centered at $\mathcal{A}_\psi(\mathbf{x}_t)$ with variance σ , i.e. $\mathcal{N}(\mathbf{x}_t, \sigma \mathbf{1}_l)$ with the unit matrix $\mathbf{1}_l \in \mathbb{R}^{l \times l}$, the close form of Eq. (12) is

$$\log \det(\sigma \mathbf{1}_l + A_\psi \Sigma_t A_\psi^\top) + C, \quad (14)$$

where C is a constant that does not depend on the angle. We refer to Barbano et al. (2022a, Appendix A) for a derivation.

Departing from the exact Gaussian setting, we turn to approximating the acquisition functions using samples. More specifically, assume we are given image samples $\mathbf{x}_t^1, \dots, \mathbf{x}_t^k$ from the (approximate) posterior distribution, with mean prediction $\frac{1}{l} \sum_{i=1}^l \mathbf{x}_t^i$. To obtain a sampling based approximation of the acquisition functions Eqs. (3) and (14), note that in the Gaussian model, the term $A_\psi \Sigma_t A_\psi^\top \in \mathbb{R}^{l \times l}$ is the variance of the posterior mean observation. This term can be directly approximated from samples, i.e.

$$A_\psi \Sigma_t A_\psi^\top = \text{Cov}[A_\psi \mathbf{x}_t^1] \approx \sum_{i=1}^k (A_\psi \mathbf{x}_t^i - A_\psi \bar{\mathbf{x}}_t)(A_\psi \mathbf{x}_t^i - A_\psi \bar{\mathbf{x}}_t)^\top. \quad (15)$$

Uncertainty Sampling Uncertainty sampling aims at querying data points with the largest posterior total variance, i.e.

$$\psi_{t+1} = \arg \max_{\psi \in \Phi} \text{tr}(\text{Cov}[A_\psi \mathbf{x}_t^1]). \quad (16)$$

The variance can be analogously approximated from posterior samples using Eq. (15).

Query by Committee Lastly, committee based acquisition (Seung et al., 1992) aims at taking measurements that maximize the disagreement in the measurements for candidates $\mathbf{x}_t^1, \dots, \mathbf{x}_t^k$ towards reference prediction $\hat{\mathbf{x}}_t$. For example, taking the average squared Euclidean norm, we choose the angle maximizing the disagreement,

$$\psi_{t+1} = \arg \max_{\psi \in \Phi} \sum_{i=1}^k \|A_\psi \mathbf{x}_t^i - A_\psi \hat{\mathbf{x}}_t\|^2 \quad (17)$$

For a Gaussian likelihood, this corresponds to the average KL between the observation distribution induced by \mathbf{x}_t^i and $\hat{\mathbf{x}}_t$ respectively. Other variants such as worst-case disagreement or other divergence measures are also possible (Hino & Eguchi, 2023). We note that when taking the mean prediction as reference, $\hat{\mathbf{x}}_t = \bar{\mathbf{x}}_t$, the committee based approach reduces to maximizing the sample variance.

D.2 EVALUATION OF ACQUISITION FUNCTIONS

In our evaluation, we compare the three acquisition function (variance, log-determinant and committee based) to a static uniform design. While we observed significant gains using all of the active learning strategies, there was no significant difference among the different acquisition functions (although it is possible to construct examples where the acquisition functions differ). This leads us to recommend the simplest, variance acquisition strategy given our current evaluation. A detailed overview of the results is given in Fig. 19

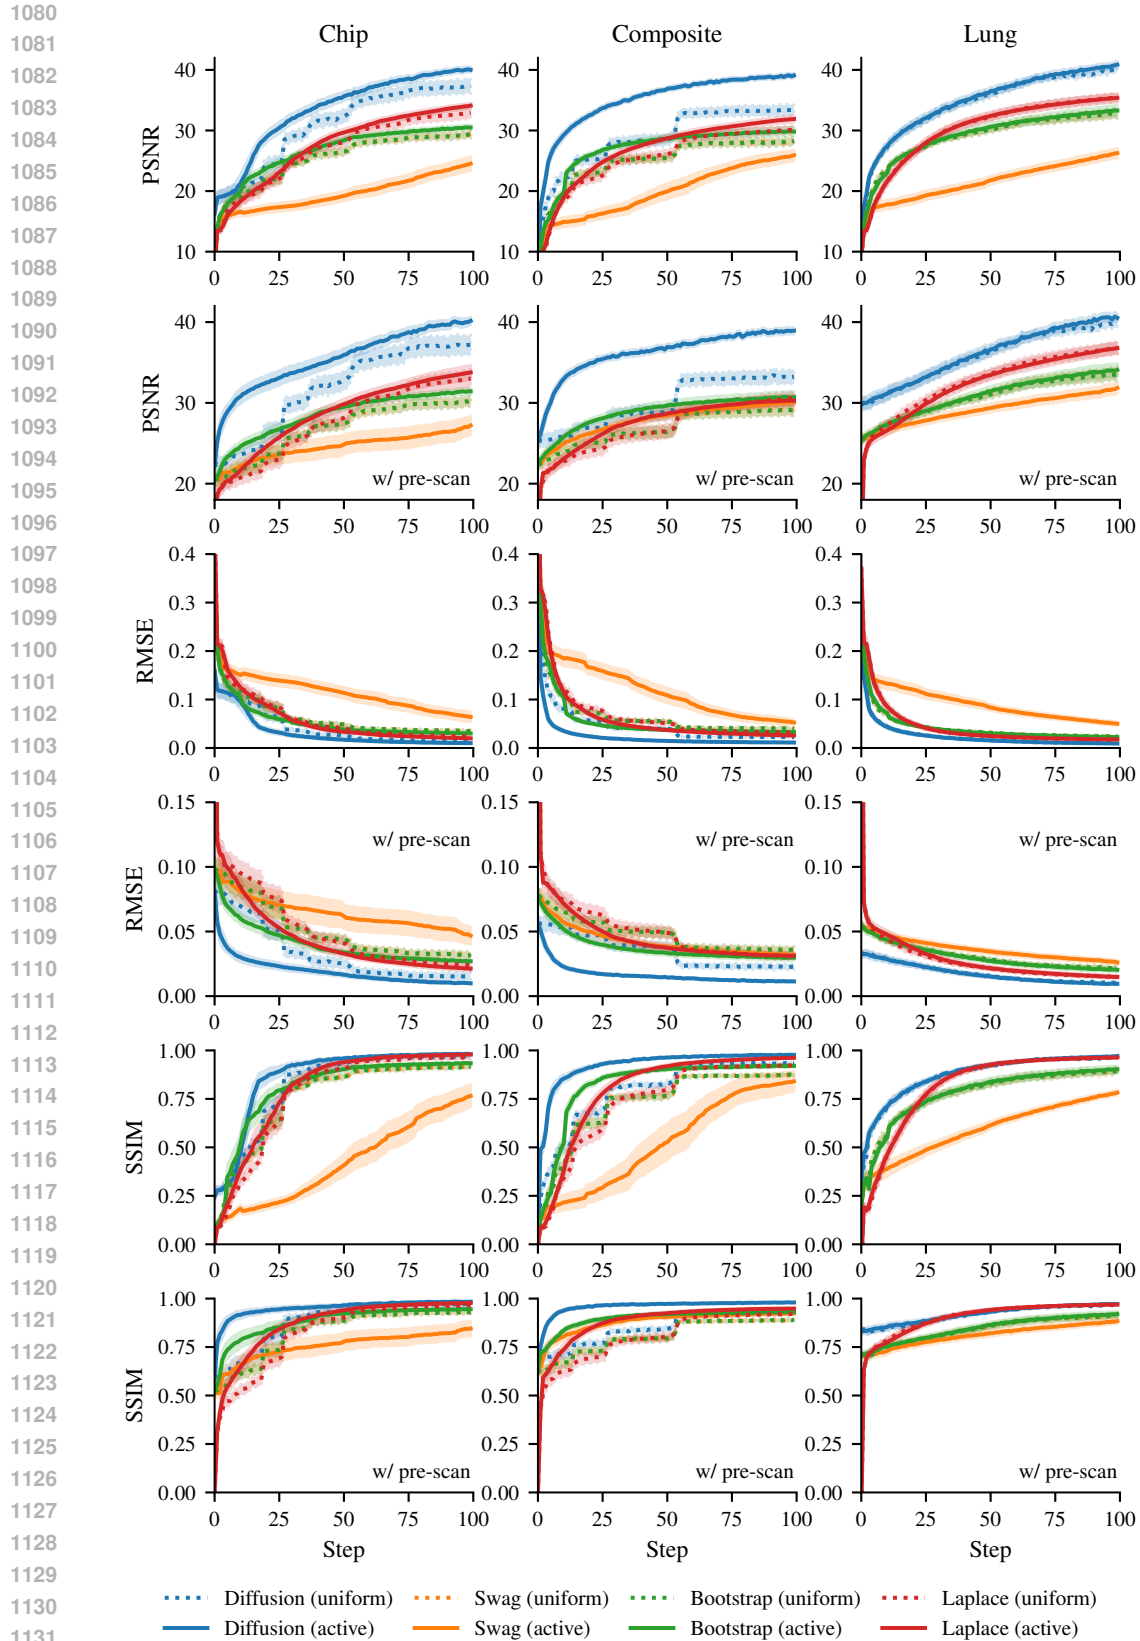


Figure 12: Results for different evaluation metrics (PSNR, RMSE, SSIM). The plots for PSNR are the same as in the main text.

1134
 1135
 1136
 1137
 1138
 1139
 1140
 1141
 1142
 1143
 1144
 1145
 1146
 1147
 1148
 1149
 1150
 1151
 1152
 1153
 1154
 1155
 1156
 1157
 1158
 1159
 1160
 1161
 1162
 1163
 1164
 1165
 1166
 1167
 1168
 1169
 1170
 1171
 1172
 1173
 1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187

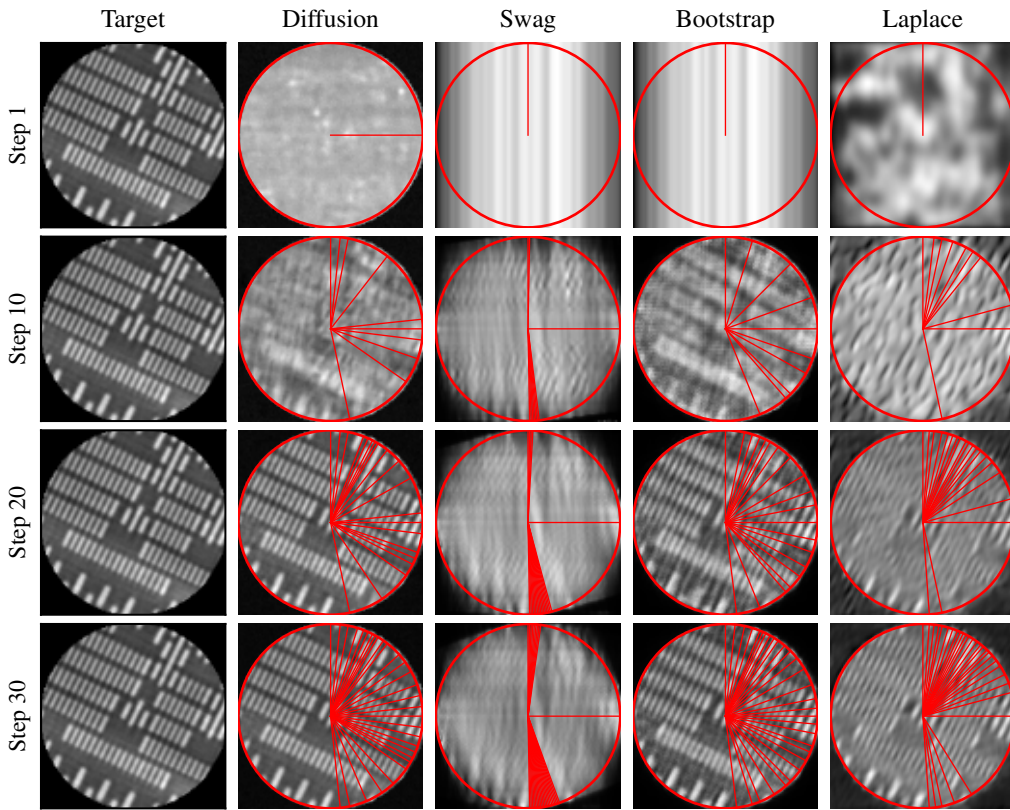


Figure 13: Qualitative results for the ‘Chip’ dataset (128×128). DAL eventually focuses its attention in the direction of the chip structures and its orthogonal direction; other algorithms struggle to pick both of them. .

1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198
 1199
 1200
 1201
 1202
 1203
 1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241

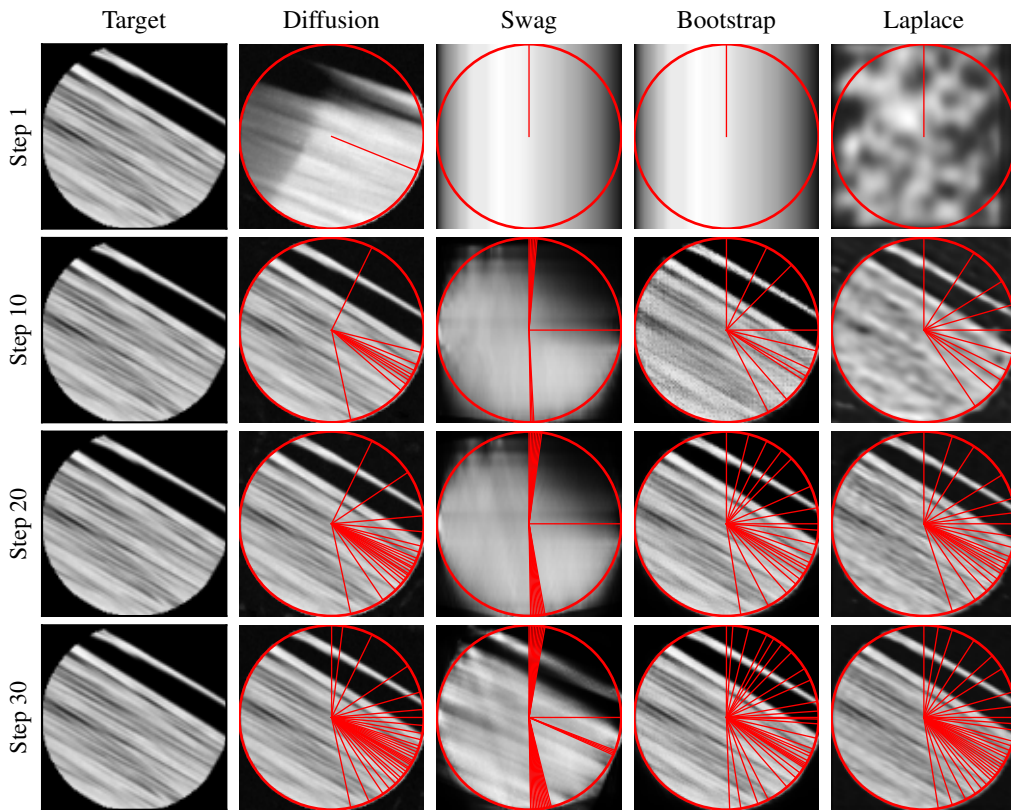


Figure 14: Qualitative results for the ‘Composite’ dataset (128×128). There exists one prominent direction containing most of the information, which is quickly picked up by DAL. Other algorithms eventually catch up, but take longer to converge to this direction.

1242
 1243
 1244
 1245
 1246
 1247
 1248
 1249
 1250
 1251
 1252
 1253
 1254
 1255
 1256
 1257
 1258
 1259
 1260
 1261
 1262
 1263
 1264
 1265
 1266
 1267
 1268
 1269
 1270
 1271
 1272
 1273
 1274
 1275
 1276
 1277
 1278
 1279
 1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295




Figure 15: Qualitative results for the ‘Lung’ dataset (128×128). While some structure exists in the sampling strategy with large sparsity, it quickly converges to a uniform distribution. This shows why uniform sampling is on par with Active Learning strategies for the Lung dataset.

1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313
 1314
 1315
 1316
 1317
 1318
 1319
 1320
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330
 1331
 1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349




Figure 16: Qualitative results for the ‘Chip’ dataset (512×512). DAL quickly focuses on the direction of the chip structures and its orthogonal direction; other algorithms struggle to pick both of them.

1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403




Figure 17: Qualitative results for the fastMRI dataset using active learning acquisition uncertainty sampling and 2 columns preselected.

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457




Figure 18: Qualitative results for the fastMRI dataset using active learning acquisition uncertainty sampling and 30 columns preselected.

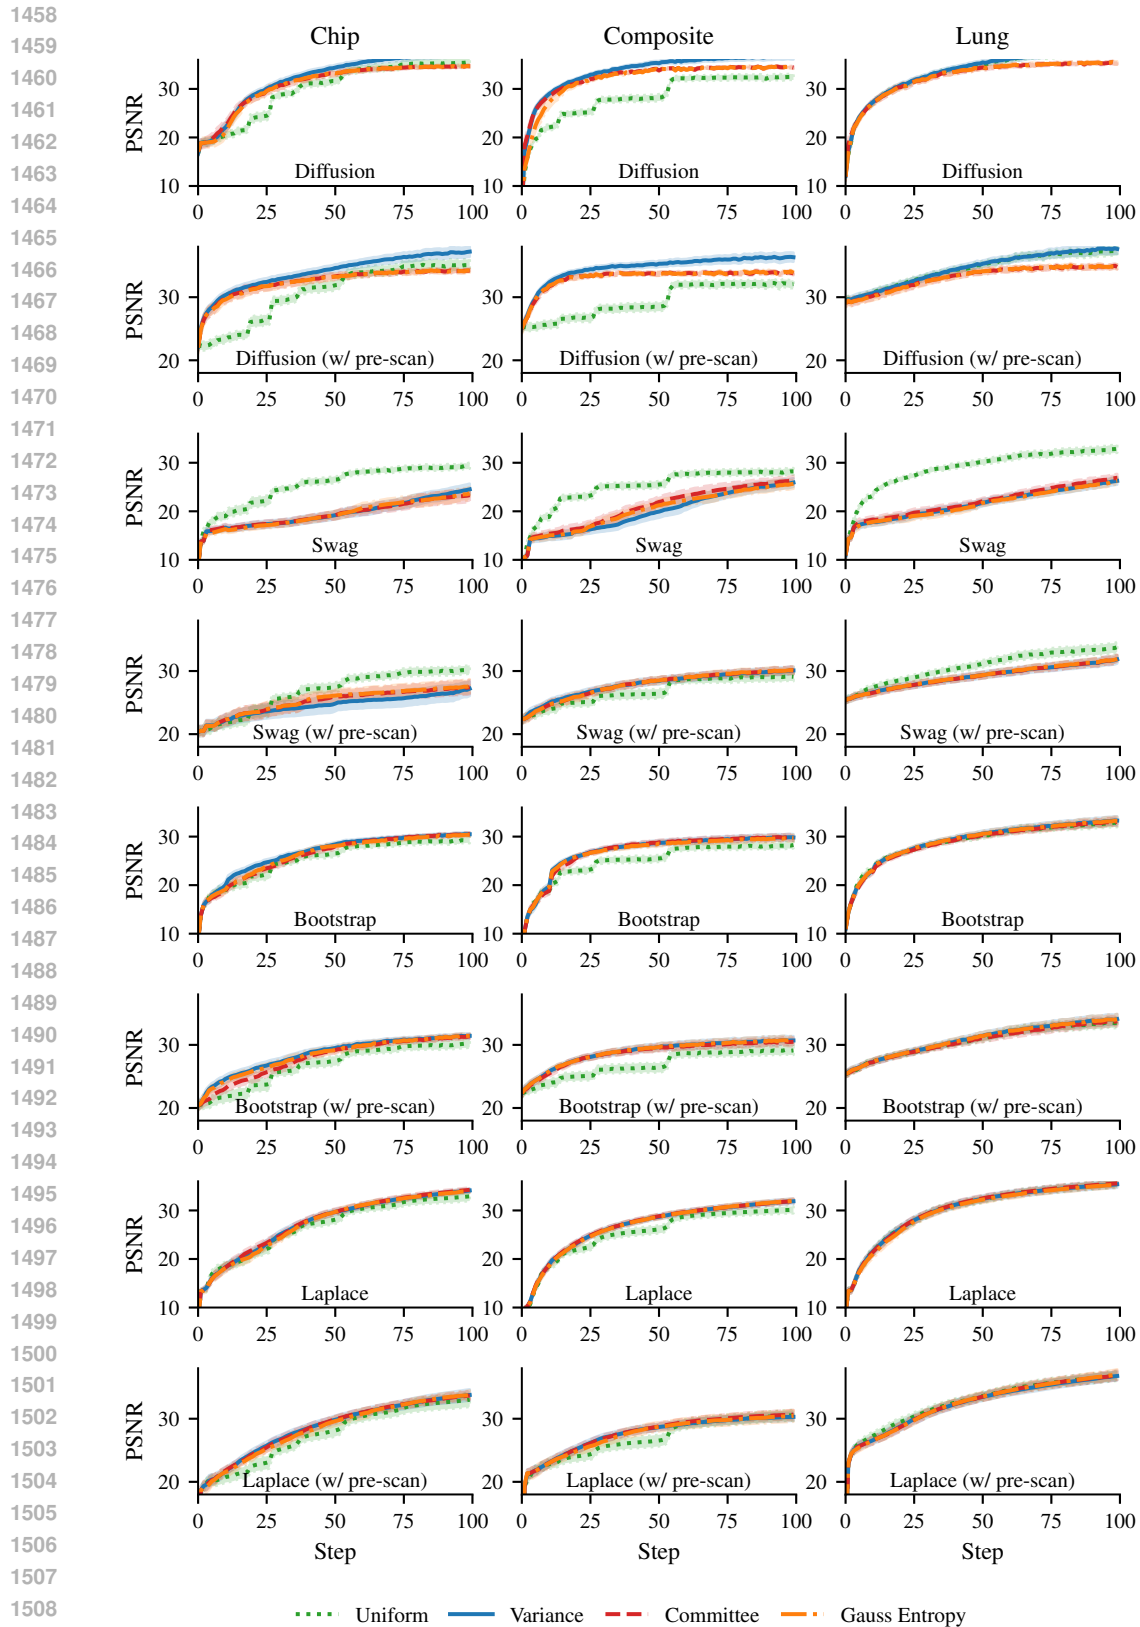


Figure 19: Results comparing different acquisition functions for each model. We show PSNR for three active acquisition strategies (Variance 3, Committee 17 and Gauss Entropy 14) and the uniform baseline. In almost all cases (except for SWAG), there is no visible difference among the different acquisition strategies.