Image Completion via Inference in Deep Generative Models

William Harvey [1] **Saeid Naderiparizi** [1] **Frank Wood** [1 2]

Abstract

We consider image completion from the perspective of amortized inference in an image generative model. We leverage recent state of the art variational auto-encoder architectures that have been shown to produce photo-realistic natural images at non-trivial resolutions. Through amortized inference in such a model we can train neural artifacts that produce diverse, realistic image completions even when the vast majority of an image is missing. We demonstrate superior sample quality and diversity compared to prior art on the CIFAR-10 and FFHQ-256 datasets. We conclude by describing and demonstrating an application that requires an in-painting model with the capabilities ours exhibits: the use of Bayesian optimal experimental design to select the most informative sequence of small field of view x-rays for chest pathology detection.

1. Introduction

Given an image with missing pixel values, the image completion task is to infer, or inpaint, these missing values. Research on this subject has often been motivated by applications such as image restoration (where damaged parts of a photograph may be removed and inpainted) and object removal (where a user deletes and then inpaints part of an image containing an unwanted object). Image completion without uncertainty is as an ill-posed problem, since it is impossible to exactly recover the missing pixel values. Exact recovery is often not even desired: an object removal tool would be of little use if it draws back objects deleted by the user. Instead, traditional image completion techniques (Bertalmio et al., 2001; Barnes et al., 2009) aim to inpaint missing regions to be similar to the rest of the image and sometimes explicitly penalise completed regions with features that "stand out" (Barnes et al., 2009). These techniques work well for many applications, but break down if more than a small portion of the image is missing, or if multiple diverse samples are required.

We view image completion from a Bayesian perspective. This requires us to first specify a prior distribution over images $p(\mathcal{I})$, which we model with a variational auto-encoder (VAE). We then consider how pixels can be removed from $\mathcal{I}$ by modeling the known pixel values $\hat{\mathcal{I}}$ as being drawn from the conditional distribution $p(\hat{\mathcal{I}}|\mathcal{I})$. This leads to a posterior distribution over image completions via Bayes' rule:

$$p(\mathcal{I}|\hat{\mathcal{I}}) = \frac{p(\mathcal{I})p(\hat{\mathcal{I}}|\mathcal{I})}{\int p(\mathcal{I})p(\hat{\mathcal{I}}|\mathcal{I})\mathrm{d}\mathcal{I}}. \tag{1}$$

The image completion task is then to draw approximate samples from this posterior. Basing our approach on this perspective allows us to inherit two major advantages of Bayesian inference. First, we can naturally incorporate uncertainty, meaning that our sampled image completions are diverse without the need for ad-hoc loss terms. Second, our approach can work well with few observations; i.e., provided we have a good generative model, we can produce plausible completions for images with arbitrarily many missing pixels.

Our approach to sampling from $p(\hat{\mathcal{I}}|\mathcal{I})$ draws heavily from the VAE framework (Kingma & Welling, 2013). VAEs can generate images by sampling latent variables $z \sim p_{\text{model}}(\cdot)$ and then the image $\mathcal{I} \sim p_{\text{model}}(\cdot|z)$. They can also reconstruct images if z is instead sampled from an encoder $q(z|\mathcal{I})$. Image completion can be viewed as an interpolation between these two extremes; we wish to reconstruct observed pixels and sample the values of the remainder. This motivates us to tackle image completion via the learning of an additional component: a partial encoder $\hat{q}(z|\hat{\mathcal{I}})$ which maps from a corrupted image to a distribution over z. In the extreme case where all pixels are missing from $\hat{\mathcal{I}}$, the optimal $\hat{q}(z|\hat{\mathcal{I}})$ is the prior $p_{\text{model}}(z)$. Similarly if no pixels are missing, $\hat{q}(z|\hat{\mathcal{I}})$ should match the encoder. We show that, in between these two extremes, $\hat{q}(z|\hat{\mathcal{I}})$ can learn distributions over z which lead to highly diverse and plausible image completions. An additional benefit to basing our architecture on a VAE is that, when a pretrained VAE is available, we can use it to speed up training.

In Section 5 we demonstrate an application of an uncertainty-preserving image completion system. In particular, we show

[1]Department of Computer Science, University of British Columbia, Canada [2]MILA. Correspondence to: William Harvey <wsgh@cs.ubc.ca>.

Preprint.

Figure 1. **Left column:** Images with most pixels masked out. **Rest:** Completions from our method. We show 5 handpicked from a batch of 10. They exhibit greater diversity and image quality than competing methods.

how it can be used as part of a Bayesian optimal experimental design (BOED) (Chaloner & Verdinelli, 1995) pipeline to control a sensor. The sensor can observe only a small portion of a scene at a time. After observing each portion, we sample scene completions based on the observed part. Further processing of these completions can be used to determine where in the scene will be most informative to observe next. This is useful if there is a high cost to observing the scene. We consider a scenario where the BOED pipeline is used for x-ray imaging. In this case, it is desirable to use as few observations as possible so as to limit the patient's exposure to radiation.

To summarise, we present an approach to convert a VAE into an image completion model. We then derive and compare two training objectives and show that one of them consistently outperforms all baselines. Finally, we show results indicating that this method could prove useful in the field of medical imaging.

2. Background

2.1. Image completion

Early diffusion-based approaches to image completion (Bertalmio et al., 2000; 2001; Ballester et al., 2001; Levin et al., 2003) could complete small gaps in images, such as those made by creases in an old photograph. A later strand of work filled larger gaps by copying and aligning patches from the rest of the image (Criminisi et al., 2003; Barnes et al., 2009). These methods produce convincing continuations of the background but no new structure.

Breakthroughs in deep learning enabled a plethora of advances in image completion. Köhler et al. (2014); Ren et al. (2015) trained convolutional neural networks (CNNs) to map directly from masked images to estimates of the missing pixel values using, e.g., a mean squared error (MSE) reconstruction loss. While this works well with small missing patches, the networks are deterministic and give blurry completions when the missing regions are large. Attempts to make more realistic completions have led to increasingly complex methods being used. Image completions are now often produced in several stages, with intermediate outputs such as segmentation masks (Song et al., 2018) or coarse-grained images (Yu et al., 2018; 2019) allowing additional losses to be imposed. In addition to reconstruction losses, it is also now common to use an adversarial loss (Pathak et al., 2016), or more often several adversarial losses from different discriminators (Iizuka et al., 2017). Domain-specific losses have also been considered: Li et al. (2017) crafted a loss term specifically for the completion of images of faces. Although adversarial losses were introduced as a tool to learn distributions (Goodfellow et al., 2014), most of these works are incapable of generating diverse completions from a single input. Creating diverse completions has therefore been the focus of several recent studies (Zheng et al., 2019; Zhao et al., 2020). However, these impose diversity using ad-hoc loss terms weighted by hyperparameters, and it is not clear how well they match the desired posterior $p(\mathcal{I}|\hat{\mathcal{I}})$. They also typically require many image pixels to be observed to generate compelling completions.

Prior work has used VAEs in conditional generative models, often for image completion (Sohn et al., 2015; Zheng et al., 2019; Ivanov et al., 2018; Ma et al., 2018). The decoders in these methods receive the conditioning input as well as sampled latent variables. A major difference in our method is that the conditioning happens purely through the sampled latent variables, and so we can use architectures and weights designed/trained for unconditional VAEs. Another related approach is that of Song et al. (2020), who present a stochastic differential equation-based image model. This

can be conditioned on a subset of image pixels to perform image completion, but is slow to sample from.

2.2. Variational auto-encoders

We describe VAEs in terms of three components. A decoder with parameters θ maps from latent variables z to data $\mathcal{I}$ (e.g. an image). The prior over latent variables, $p_{\text{model}}(z)$, may have learnable parameters, which we consider to be part of θ. Together, the prior and decoder define a joint distribution, $p_{\text{model}}(z, \mathcal{I})$. Finally, an encoder with parameters ϕ maps from data to an approximate posterior distribution over latent variables, $q(z|\mathcal{I}) \approx p_{\text{model}}(z|\mathcal{I})$. Ideally, these networks would be trained to maximise the data likelihood $p_{\text{model}}(\mathcal{I}) = \int p_{\text{model}}(z)p_{\text{model}}(\mathcal{I}|z)\mathrm{d}z$, averaged over some dataset. Since this is intractable, θ and ϕ are instead trained jointly to maximise an average of the evidence lower-bound (ELBO) over the data distribution $p_{\text{data}}(\mathcal{I})$:

$$\begin{aligned}
&\mathbb{E}_{p_{\text{data}}(\mathcal{I})}\left[\text{ELBO}(\theta, \phi, \mathcal{I})\right] \\
&\quad= \mathbb{E}_{p_{\text{data}}(\mathcal{I})q(z|\mathcal{I})}\left[\log \frac{p_{\text{model}}(z)p_{\text{model}}(\mathcal{I}|z)}{q(z|\mathcal{I})}\right] \quad (2) \\
&\quad= \mathcal{H}\left[p_{\text{data}}(\mathcal{I})\right] - \text{KL}\big(r(z, \mathcal{I}) \,\big\|\, p_{\text{model}}(z, \mathcal{I})\big) \quad (3)
\end{aligned}$$

where we define $r(z, \mathcal{I}) = p_{\text{data}}(\mathcal{I})q(z|\mathcal{I})$ to be the joint distribution defined by the data distribution and learned encoder. As Eq. (3) shows, maximising this objective will drive $p_{\text{model}}(z, \mathcal{I})$ towards $r(z, \mathcal{I})$, and so drive the marginal $p_{\text{model}}(\mathcal{I})$ towards $p_{\text{data}}(\mathcal{I})$.

A major weakness of VAEs in the past has been difficulty scaling to produce convincing high-resolution images. However, recent hierarchical architectures have made significant progress towards this end (Vahdat & Kautz, 2020; Child, 2020). Hierarchical VAEs differ from traditional VAEs by partitioning their latent variables z into L disjoint groups, $z_1, \ldots, z_L$ (Gregor et al., 2015; Kingma et al., 2016; Sønderby et al., 2016; Klushyn et al., 2019). We show a typical hierarchical VAE architecture in Fig. 2a. The decoder (black) maintains a hidden state h_l, which is initialized with a fixed h_0 and then incorporates each group of latent variables as they are sampled. By conditioning the prior for each z_l on the previous hidden state, the prior for z is factorised as

$$p_{\text{model}}(z) = \prod_{l=1}^{L} p_{\text{model}}(z_l|z_{<l}) \quad (4)$$

where $z_{<l}$ is the null set for $l = 1$ and $\{z_1, \ldots, z_{l-1}\}$ otherwise. The distribution produced by the encoder (shown in orange) for each z_l also depends on the previous hidden state h_{l-1} and therefore factorises as $q(z|\mathcal{I}) = \prod_{l=1}^{L} q(z_l|z_{<l}, \mathcal{I})$. This dependence on the decoder's hidden state also means that $q(z|\mathcal{I})$ depends on θ as well as the encoder parameters ϕ. Both $p_{\text{model}}(z_l|z_{<l})$ and $q(z_l|z_{<l}, \mathcal{I})$ usually assume independence between dimensions of z_l, and are often diagonal Gaussian distributions (Sønderby et al., 2016; Vahdat & Kautz, 2020; Child, 2020).

We will present a method to complete images through inference in the model defined by a VAE with a subset of image pixel values observed. For this inference to be well-posed, it is important for the observed pixels to have non-zero probability under the model; the denominator in Bayes' rule (Eq. (1)) would otherwise be zero. Hierarchical VAEs are well-suited to satisfying this requirement for two reasons. First, the VAE objective optimises θ with a mass-covering KL divergence (Eq. (3)), which harshly penalises θ for assigning low probability density to any images in the data distribution. Second, recent hierarchical architectures use latent variables with extremely high dimensionality: the VD-VAE (Child, 2020) uses 412 000 latent variables to model the CIFAR-10 dataset of 3072-dimensional images. This "over-parameterisation" ensures that the images produced are not constrained to a manifold. We confirm that these architectures can produce extremely mass-covering distributions; we find that a VD-VAE trained on images of faces can near-perfectly reconstruct wildly different data, including images of text and chest x-rays (see appendix). This property prevents pitfalls similar to those observed by Yeh et al. (2017) whose method, which completes images by estimating latent variables in a GAN, would often fail to reconstruct the known pixel values.

3. Method

To convert an unconditional VAE architecture to a conditional architecture, we introduce a *partial encoder* with parameters $\hat{\phi}$. This is fed an observation $\hat{\mathcal{I}}$ (e.g. a corrupted image) and defines an approximate posterior over the latent variables, $\hat{q}(z|\hat{\mathcal{I}})$. We can then approximate the posterior over image completions,

$$p_{\text{model}}(\mathcal{I}|\hat{\mathcal{I}}) = \int p_{\text{model}}(\mathcal{I}|z)p_{\text{model}}(z|\hat{\mathcal{I}})\mathrm{d}z \quad (5)$$

$$\approx p_{\text{inpaint}}(\mathcal{I}|\hat{\mathcal{I}}) = \int p_{\text{model}}(\mathcal{I}|z)\hat{q}(z|\hat{\mathcal{I}})\mathrm{d}z. \quad (6)$$

We can sample from $p_{\text{inpaint}}(\mathcal{I}|\hat{\mathcal{I}})$ by simply sampling $z \sim \hat{q}(\cdot|\hat{\mathcal{I}})$ and then $\mathcal{I} \sim p_{\text{model}}(\cdot|z)$.

Our method requires that we have a model of how $\hat{\mathcal{I}}$ is generated from $\mathcal{I}$. More precisely, we need to define the conditional distribution $p(\hat{\mathcal{I}}|\mathcal{I})$. We do so by defining a distribution over image-sized binary masks. Then to sample from $p(\hat{\mathcal{I}}|\mathcal{I})$, we sample a mask m before returning $\hat{\mathcal{I}} = \texttt{concatenate}(\mathcal{I} \odot m, m)$. Here, $\odot$ is a pixel-wise multiplication operation which removes information from the missing pixels. The concatenation is performed along the channel dimension and makes it possible to distinguish between unobserved pixels and zero-valued pixels.

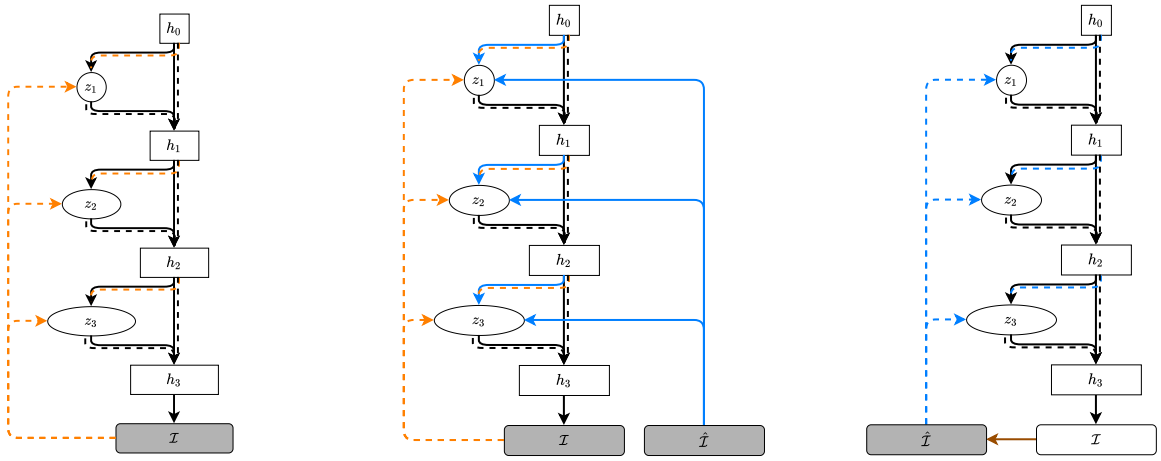

(a) A standard hierarchical VAE. (b) Forward KL training. (c) Reverse KL training.

Figure 2. A hierarchical VAE architecture with $L = 3$ groups of latent variables. The blocks $h_0, \ldots, h_L$ represent the deterministic hidden state of the decoder. Part (a) shows the standard VAE setting with an encoder (orange) and decoder (black). Parts (b) and (c) show two ways to extend this with a partial encoder (blue). Each is trained with with a different ELBO which lower bounds a different quantity. The dashed lines are the dependencies for distributions that are sampled from when the corresponding ELBO is computed. The solid lines are the dependencies of the 'model' distribution on the numerator of the corresponding ELBO.

3.1. Training objectives

Recall that a VAE defines two joint distributions over z and $\mathcal{I}$: the distribution of samples from the generator, $p_{\text{model}}(z, \mathcal{I})$; and the distribution resulting from sampling data $\mathcal{I}$ and encoding it, which we will now call $r(z, \mathcal{I}) = p_{\text{data}}(\mathcal{I})q(z|\mathcal{I})$. For notational convenience, we define the following joint distributions:

$$p_{\text{model}}(z, \mathcal{I}, \hat{\mathcal{I}}) = p_{\text{model}}(z)p_{\text{model}}(\mathcal{I}|z)p(\hat{\mathcal{I}}|\mathcal{I}) \quad (7)$$

$$p_{\text{data}}(\mathcal{I}, \hat{\mathcal{I}}) = p_{\text{data}}(\mathcal{I})p(\hat{\mathcal{I}}|\mathcal{I}) \quad (8)$$

$$r(z, \mathcal{I}, \hat{\mathcal{I}}) = p_{\text{data}}(\mathcal{I}, \hat{\mathcal{I}})q(z|\mathcal{I}) \quad (9)$$

We will use p_{model}, p_{data}, and r in this paper to refer to marginals and conditionals of these joint distributions, with the specific marginal or conditional made clear from context. With this notation, we now derive and compare two objectives for training the conditional generative model, with particular focus on the properties of the learned $\hat{q}(z|\hat{\mathcal{I}})$.

3.1.1. Forward KL

Consider maximising $\log p_{\text{model}}(\mathcal{I}|\hat{\mathcal{I}})$ on expectation over $\mathcal{I} \sim p_{\text{data}}(\cdot)$ and $\hat{\mathcal{I}} \sim p(\cdot|\mathcal{I})$. We can lower bound this as

$$\mathcal{O}_{\text{for}}(\theta, \phi, \hat{\phi}) = \mathbb{E}_{q(z|\mathcal{I})p_{\text{data}}(\mathcal{I}, \hat{\mathcal{I}})}\left[\log \frac{p_{\text{model}}(\mathcal{I}|z)\hat{q}(z|\hat{\mathcal{I}})}{q(z|\mathcal{I})}\right] \quad (10)$$

$$\leq \mathbb{E}_{p_{\text{data}}(\mathcal{I}, \hat{\mathcal{I}})}\left[\log p_{\text{inpaint}}(\mathcal{I}|\hat{\mathcal{I}})\right] \quad (11)$$

That is, we can lower-bound $\log p_{\text{inpaint}}(\mathcal{I}|\hat{\mathcal{I}})$ very similarly to how the ELBO in an unconditional VAE lower-bounds $\log p_{\text{model}}(\mathcal{I})$ (Eq. (2)). The only difference is that the prior, $p_{\text{model}}(z)$, is replaced by $\hat{q}(z|\hat{\mathcal{I}})$. This is reflected in Fig. 2b, where each z_l is now conditioned on $\hat{\mathcal{I}}$ via the partial encoder (blue).

We now consider the effect this objective has on the learned distribution $\hat{q}(z|\hat{\mathcal{I}})$. We can decompose Eq. (10) as follows. Starting by multiplying both sides of the fraction by the intractable conditional distribution $r(z|\hat{\mathcal{I}})$,

$$\mathcal{O}_{\text{for}}(\theta, \phi, \hat{\phi})$$

$$= \mathbb{E}_{r(z, \mathcal{I}, \hat{\mathcal{I}})}\left[\log \frac{p_{\text{model}}(\mathcal{I}|z)\hat{q}(z|\hat{\mathcal{I}})r(z|\hat{\mathcal{I}})}{q(z|\mathcal{I})r(z|\hat{\mathcal{I}})}\right] \quad (12)$$

$$= \mathbb{E}_{r(z, \mathcal{I}, \hat{\mathcal{I}})}\left[\log \frac{p_{\text{model}}(\mathcal{I}|z)r(z|\hat{\mathcal{I}})}{q(z|\mathcal{I})}\right] - \mathbb{E}_{r(z, \hat{\mathcal{I}})}\left[\log \frac{r(z|\hat{\mathcal{I}})}{\hat{q}(z|\hat{\mathcal{I}})}\right] \quad (13)$$

$$= C - \mathbb{E}_{p_{\text{data}}(\hat{\mathcal{I}})}\left[\text{KL}\big(r(z|\hat{\mathcal{I}}) \,\|\, \hat{q}(z|\hat{\mathcal{I}})\big)\right] \quad (14)$$

where C does not depend on the partial encoder's parameters $\hat{\phi}$. This shows that, when maximising $\mathcal{O}_{\text{for}}(\theta, \phi, \hat{\phi})$, the partial encoder will be trained to minimise the *forward* KL divergence to $r(z|\hat{\mathcal{I}})$. Minimising the forward KL leads to mass-covering behaviour (Bishop, 2006) and so this objective could be expected to lead to diverse samples of the latent variables $z \sim \hat{q}(\cdot|\hat{\mathcal{I}})$, and subsequently diverse image completions. We name the method of training with this objective AIPO (amortized inference with partial observations). We also note that this objective is the same as

Table 1. Results summary. The best performance on each metric is shown in **bold**, and the second best is underlined.

	CIFAR-10			FFHQ-256		
Method	FID	IS	FID (holes)	FID	LPIPS	FID (holes)
AIPO (ours)	**27.7**	**3.41**	**2.08**	**17.4**	**0.325**	**1.62**
AIPO-R (ours)	49.2	2.61	2.44	55.2	0.204	2.05
PIC	75.3	2.01	4.00	24.6	0.286	3.01
ANP	80.8	1.86	3.85	168.5	0.225	125.4
CE	117.6	1.00	3.45	107.6	0.000	3.93

that used by neural processes with latent variables (Garnelo et al., 2018).

3.1.2. Reverse KL

An alternative is to maximise a lower bound on $\log p_{\text{model}}(\hat{\mathcal{I}})$, averaged over $p_{\text{data}}(\hat{\mathcal{I}})$:

$$\mathcal{O}_{\text{rev}}(\theta, \phi, \hat{\phi}) = \mathbb{E}_{p_{\text{data}}(\hat{\mathcal{I}})\hat{q}(z|\hat{\mathcal{I}})}\left[\log \frac{p_{\text{model}}(z, \hat{\mathcal{I}})}{\hat{q}(z|\hat{\mathcal{I}})}\right]. \quad (15)$$

This can be computed with the architecture shown in Fig. 2c. This architecture has two differences from the unconditional VAE (Fig. 2a): the partial encoder with input $\hat{\mathcal{I}}$ is used in place of the full encoder with input $\mathcal{I}$; and the likelihood $p_{\text{model}}(\mathcal{I}|z)$ is replaced by $p_{\text{model}}(\hat{\mathcal{I}}|z)$. We note that computing $p_{\text{model}}(\hat{\mathcal{I}}|z) = \int p_{\text{model}}(\mathcal{I}|z)p(\hat{\mathcal{I}}|\mathcal{I})\mathrm{d}\mathcal{I}$ involves a marginalisation over $\mathcal{I}$. In the setting we focus on, where $p_{\text{model}}(\mathcal{I}|z)$ consists of independent likelihoods for each pixel and $\hat{\mathcal{I}}$ consists of a subset of the pixels of $\mathcal{I}$, this is simple. We can compute $p_{\text{model}}(\hat{\mathcal{I}}|z)$ by only summing up the likelihoods for pixels in $\hat{\mathcal{I}}$ (rather than for all pixels in $\mathcal{I}$, as to compute $p_{\text{model}}(\mathcal{I}|z)$).

To reveal the properties of the learned partial encoder, we can decompose the objective as

$$\mathcal{O}_{\text{rev}}(\theta, \phi, \hat{\phi}) = \mathbb{E}_{p_{\text{data}}(\hat{\mathcal{I}})}\Big[\log p_{\text{model}}(\hat{\mathcal{I}}) - \text{KL}\big(\hat{q}(z|\hat{\mathcal{I}}) \,\big\|\, p_{\text{model}}(z|\hat{\mathcal{I}})\big)\Big]. \quad (16)$$

Maximising this objective will therefore minimise the *reverse* KL, and so $\hat{q}(z|\hat{\mathcal{I}})$ may exhibit mode-seeking behaviour. Networks trained to maximise $\mathcal{O}_{\text{rev}}$ may therefore produce less diverse samples than those with $\mathcal{O}_{\text{for}}$. This is counterbalanced by the potential that the samples produced could be of higher quality. We denote the method of training with this objective AIPO-R (amortized inference with partial observations and the reverse KL).

Note that these two KL divergences are with respect to different conditional distributions: using $\mathcal{O}_{\text{for}}$ minimises the divergence to $r(z|\hat{\mathcal{I}})$, whereas using $\mathcal{O}_{\text{rev}}$ will minimise the divergence to $p_{\text{model}}(z|\hat{\mathcal{I}})$. However, these two conditional distributions are likely to be close since they are matched by optimising the ELBO (as Eq. (3) shows).

3.2. Faster training with a pretrained VAE

We consider the two objectives we propose and the standard unconditional VAE objective to be consistent in the sense that they can all be simultaneously maximised. Specifically, all three objectives will be maximised if we have **(1)** a generator and full encoder network learned such that $\text{KL}(r(z, \mathcal{I}) \,\|\, p_{\text{model}}(z, \mathcal{I})) = 0$; and **(2)** a partial encoder such that $\hat{q}(z|\hat{\mathcal{I}}) = r(z|\hat{\mathcal{I}})$. Training an unconditional VAE with the standard ELBO will lead to a generator and full encoder which satisfy the first condition (assuming sufficient expressivity and convergence to a global optimum). We can therefore speed up training of the image completion model by using a pretrained generator and full encoder, and only training the partial encoder. This leads to faster convergence, as well as faster training iterations since we only need to compute gradients for, and perform update steps on, the partial encoder's parameters $\hat{\phi}$. Another possibility is to use the pretrained full encoder's weights as an initialisation for the partial encoder's weights. In our experiments we do this with AIPO-R, for which we found it made training more stable, but not with AIPO, for which it made training less stable. For experiments on CIFAR-10 and FFHQ-256, we made use of the pretrained models released by Child (2020). This allowed us to leverage models trained for about 2 GPU-weeks and 1 GPU-year, respectively. Training each complementary partial encoder in our experiments took no more than 1 GPU-week.

4. Experiments

We create AIPO and AIPO-R image completion models based on the VD-VAE unconditional architecture (Child, 2020). We evaluate them on two datasets: CIFAR-10 (Krizhevsky et al., 2009) and FFHQ-256 (Karras et al., 2019) and compare against three baselines: Pluralistic Image Completion (PIC) (Zheng et al., 2019), Context Encoders (CE) (Pathak et al., 2016) and Attentive Neural Processes (ANP) (Kim et al., 2019).

Our focus is on completing images when the majority of pixel values are unknown. Therefore, where not otherwise specified, all results we show are with networks trained and evaluated using the type of mask seen in the left column

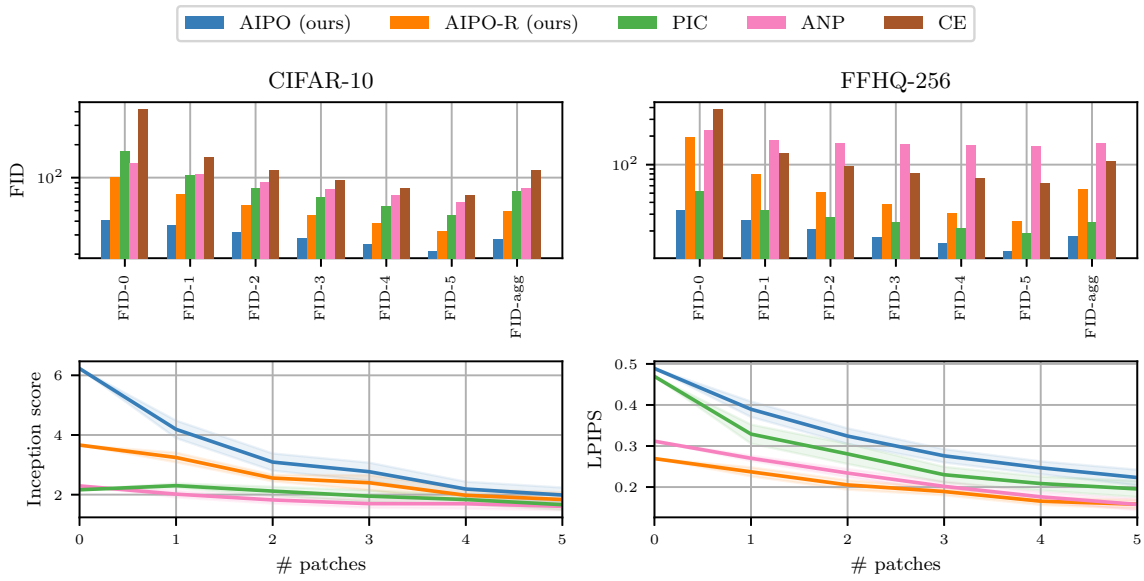

Figure 3. **Top row**: FID scores (lower is better) for completions on CIFAR-10 and FFHQ-256. FID-n is the FID between the test set and a set consisting of one completion from n patches of each test image. FID-agg is the FID between the test set and the union of all completions generated for each FID-n. **Bottom row:** Inception scores for CIFAR-10 completions (left) and LPIPS scores for FFHQ-256 completions (right). When these are high, it indicates that the method can generate diverse completions for a single $\hat{\mathcal{I}}$. Each line shows an aggregation of results on 10 test images. AIPO achieves both the lowest FID scores and the highest diversity scores.

of Fig. 1. In these, all pixels are masked out except for those within one of n squares (where $n \in \{0, 1, 2, 3, 4, 5\}$ is sampled uniformly for each image). The position of each square is randomly sampled and the side length is always 35% of that of the image. We now describe the metrics used to evaluate each method, for which results are summarised in Table 1.

4.1. Quantitative metrics

FID for completed images The FID score (Heusel et al., 2017) provides a measure of distance between image distributions. We use it to quantify the distance between the data distribution, $p_{\text{data}}(\mathcal{I})$, and the distribution of completed images. By this, we mean the distribution resulting from sampling a dataset image, masking out some pixels, and replacing them by performing inpainting. These distributions match if the distribution over completed images matches $p_{\text{data}}(\mathcal{I}|\hat{\mathcal{I}})$, since $\int p_{\text{data}}(\mathcal{I}|\hat{\mathcal{I}})p_{\text{data}}(\hat{\mathcal{I}})\mathrm{d}\hat{\mathcal{I}} = p_{\text{data}}(\mathcal{I})$. Although this metric is zero when an image completion model perfectly captures $p_{\text{data}}(\mathcal{I}|\hat{\mathcal{I}})$ it does not explicitly consider multiple completions of the same observations. We therefore view it as a measure of image completion quality more than diversity. Figure 3 shows FID scores for completions by each method with varying numbers of observed image patches. AIPO achieves significantly lower FID scores than the nearest method (which is AIPO-R on CIFAR-10 and PIC on FFHQ-256). The FID scores aggregated over each number of patches is reported in Table 1. This also reports the FID scores (described as "holes") for a different mask distribution where all pixels are observed other than those covered by a patch; essentially the inverse of the masks from our standard distribution. This is closer to the standard inpainting setting where only small patches are missing. For this evaluation, we used the AIPO and AIPO-R models trained on the standard mask distribution. However, to ensure a fair comparison rather than a test of the baselines' generalisation, we used baseline models trained on masks sampled from the "holes" distribution. Despite this, both AIPO and AIPO-R achieve significantly better FID scores than the baselines for this task.

Inception score The inception score (IS) (Salimans et al., 2016) is a widely-adopted heuristic for assessing the quality of image generative models. It takes into account both how diverse and how realistic the generated images are. We compute inception scores of the distribution over image completions for a fixed input $\hat{\mathcal{I}}$. We then report the average of this inception score over multiple inputs. In Fig. 3 (bottom left), we show inception scores for CIFAR-10 with varying numbers of observed patches. As recommended by Barratt & Sharma (2018), these were computed with a classifier trained on CIFAR-10 (details in the appendix). With no, or few, observed patches, AIPO and AIPO-R both significantly outperform the baselines. Of these two, AIPO achieves the higher score; this may reflect greater diversity

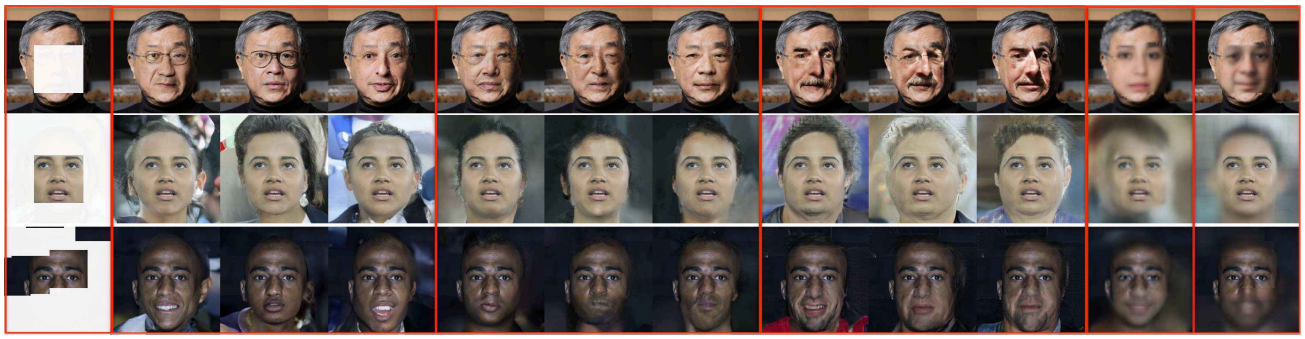

(a) $\hat{\mathcal{I}}$ (b) AIPO (c) AIPO-R (d) PIC (e) ANP (f) CE

Figure 4. Uncurated qualitative results of different models. Best viewed zoomed in. Each row shows an image completion task where the leftmost column is the input and the rest are outputs by different moels. The lighter regions of (a) are the masked out pixels.

due to the mass-covering property of its objective. As more patches are observed, the inception scores for all methods decrease towards their minimum of 1. This is to be expected since there is little uncertainty in the completion when most of the image is observed. FFHQ-256 has no labels for training a classifer, so we report an LPIPS distance (described in the following paragraph) instead.

LPIPS Following Zhu et al. (2017); Zheng et al. (2019), we use a measure of image completion diversity based on the LPIPS metric (Zhang et al., 2018), a measure of distance between two images. we assess the diversity of an image completion distribution for observation $\hat{\mathcal{I}}$ by drawing 50k pairs of images and averaging the LPIPS for each pair. We average the result over 10 different observations. A higher LPIPS score means more dissimilar pairs of images are generated, implying a higher diversity among the drawn images. It is important to note that, in contrast to the inception score, LPIPS measures only diversity and does not take into account the quality or coherency of the images. Figure 3 (bottom right) plots the LPIPS on FFHQ-256 for image completions from various numbers of patches. We see that AIPO creates the most diverse completions, closely followed by PIC. AIPO-R creates the least diverse completions.

4.2. Qualitative results

Each row of Fig. 4 shows an example image completion task. The images are randomly sampled from the test set. The observation masks for the first two tasks are manually selected to cover most of the face parts. The last observation mask is sampled from the mask distribution $p(\hat{\mathcal{I}}|\mathcal{I})$ used for training the models. While both AIPO and AIPO-R produce more plausible images than the baselines, the samples from AIPO are more diverse. PIC achieves performance close to that of our models but struggles when most of the image is missing; for an example, see the hair in the second task. Samples from ANP and CE, although consistent with the general context of the patches, are blurry. See the appendix for more examples.

5. Application to medical imaging

In this section, we investigate a possible application for image completion in medical imaging. We do not claim that our system, as it is now, is suitable for use in a clinical setting but believe this is a worthwhile avenue to explore. In particular, we consider whether it is possible to automatically target a chest x-ray at areas most likely to reveal abnormalities. This could avoid the need to scan the entire chest and so bring benefits including reducing the patient's radiation exposure. Specifically, we consider performing a series of x-ray scans, each targeted at only a small portion of the area of interest. We can select the coordinates $c_t = (x_t, y_t)$ of the scan at each step t, and this selection can be informed by what was observed in the previous scans. The task we consider is how to select c_t to be maximally informative. In particular, assume we wish to infer a variable v (representing, e.g., whether the patient has a particular illness). Bayesian optimal experimental design (BOED) (Chaloner & Verdinelli, 1995) provides a framework to select a value of c_t that is maximally informative about v. It involves taking a Bayesian perspective on the problem of estimating v given a series of scans. We have one posterior distribution over v after taking scans at $c_1, \ldots, c_{t-1}$ and another (typically lower entropy) distribution after conditioning on a scan at c_t as well. The *expected information gain*, or EIG, quantifies the utility of the choice of c_t as the expected difference in entropy between these two distributions. Using BOED involves estimating the EIG and selecting c_t to minimise it.

We use an estimator for the EIG similar to that of Harvey et al. (2019). It requires two components:

- A CNN trained to approximate the posterior over v given a series of scans at coordinates $c_1, \ldots, c_t$. We call this approximate posterior $g(v|f_{c_1,\ldots,c_t}(\mathcal{I}))$, where

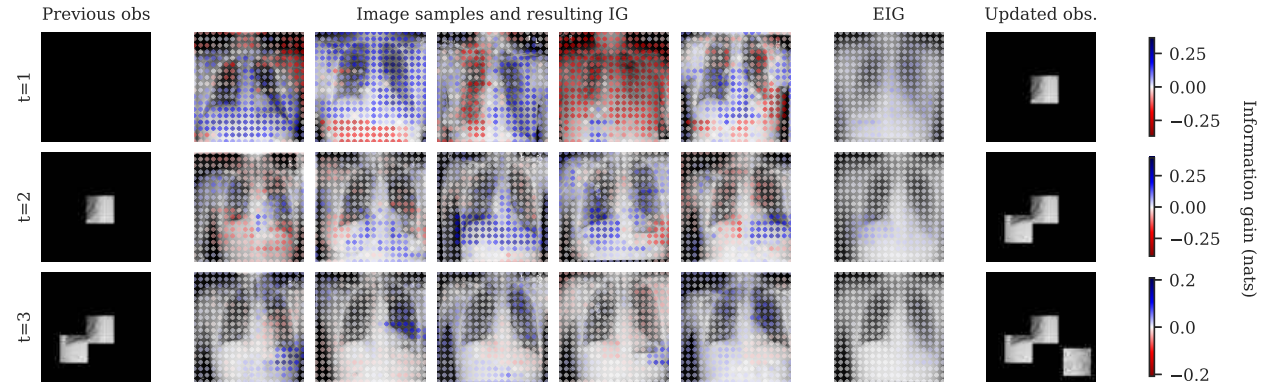

Figure 5. Visualisation of the use of BOED to select three scan locations for diagnosing 'Effusion'. The left column shows the observations made prior to each time step. The next five columns show samples from the image completion model (or dataset when $t = 1$), overlaid with a visualisation of the information gain for each possible value of c_t. The second column from the right shows the expected information gain (EIG), which is the average of the information gain (IG) over all sampled images. The pixel-space average of the sampled images is shown underneath as a visual aid. The rightmost column shows updated observations after scanning the area with the highest EIG.

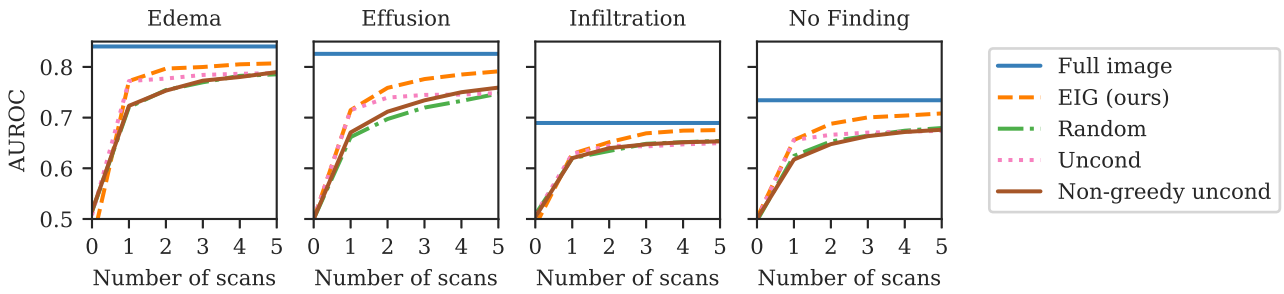

Figure 6. AUROC classification score plotted against the number of scans. We show three tasks of diagnosing particular illnesses and another, labelled 'No Finding', based on identifying the absence of any signs of illness. Using AIPO to estimate and maximise the EIG achieves performance closest to the upper bound given by processing the full image. We compare against baselines which select scan coordinates without AIPO.

$f_{c_1,\ldots,c_t}$ is a function mapping from an image to the values of the pixels observed by scans at $c_1, \ldots, c_t$.

- A method for sampling images conditioned on some observed pixel values $f_{c_1,\ldots,c_{t-1}}(\mathcal{I})$. We use AIPO.

Let the pixel values observed so far be $\hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}} = f_{c_1,\ldots,c_{t-1}}(\mathcal{I})$ for a latent image $\mathcal{I}$. Given these, we estimate the EIG of coordinate c_t as

$$\text{EIG}(c_t;\hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}}) \approx \overbrace{\mathcal{H}\left[\frac{1}{N}\sum_{n=1}^{N} g(\cdot|f_{c_1,\ldots,c_t}(\mathcal{I}^{(n)}))\right]}^{\text{entropy after } t-1 \text{ scans}} - \underbrace{\frac{1}{N}\sum_{n=1}^{N}\mathcal{H}\left[g(\cdot|f_{c_1,\ldots,c_t}(\mathcal{I}^{(n)}))\right]}_{\text{expected entropy after } t \text{ scans}}. \quad (17)$$

where $\mathcal{I}^{(1)}, \ldots, \mathcal{I}^{(N)}$ are sampled image completions from AIPO given $\hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}}$. In Appendix A we report hyper-parameters, provide further details of our EIG estimator, and compare it to the estimators used in related work. To select c_t, we simply estimate $\text{EIG}(c_t;\hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}})$ for many different values of c_t and select the value which maximises it. This process of selecting c_t and then taking a scan is repeated for each $t = 1, \ldots, T$.

We experiment on the NIH Chest X-ray 14 dataset (Wang et al., 2017) with images downsampled to 256×256. We simulate a scanner which returns a 64×64 pixel patch from this image. BOED can be used to select scan locations for an x-ray without necessarily performing an automated diagnosis. However, to quantify the informativeness of the chosen locations, Fig. 6 shows the results of using g to perform a diagnosis, or classification, based on the chosen scan locations. We show results for diagnosing four different illnesses. Since the conditional distribution g (used to estimate the EIG) depends on which illness we are classifying, the choice of scan locations is different in each case. We

compare against baselines which choose scan coordinates without AIPO; see Appendix A for details. Our method (denoted EIG) leads to performance much closer to the upper bound given by a CNN with access to the entire image. Another benefit is that the choice of scan locations is highly interpretable; we can see why a particular coordinate was chosen with visualisations similar to Fig. 5. This shows several quantities from Eq. (17): the sampled images $\mathcal{I}^{(n)}$; the information gain given each $\mathcal{I}^{(n)}$ for all values of c_t; and finally the EIG for each value of c_t.

6. Discussion and conclusion

We have presented a method for generating diverse and high-quality image completions through amortized inference in the latent space of a VAE. We derived two training objectives, one of which significantly outperformed all baselines on every metric we computed. We demonstrated an application for such a model by using it to perform BOED. There are numerous directions for future research. One is to scale AIPO to larger images: Child (2020) scale the VD-VAE architecture to 1024×1024, albeit with increasing computational cost. Another is to further improve the image quality, for example by using a partial encoder with more expressive distributions. Preliminary experiments revealed that normalizing flows could help the partial encoder better match the posterior, but with little impact on FID scores.

Acknowledgements

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canada CIFAR AI Chairs Program, and the Intel Parallel Computing Centers program. This material is based upon work supported by the United States Air Force Research Laboratory (AFRL) under the Defense Advanced Research Projects Agency (DARPA) Data Driven Discovery Models (D3M) program (Contract No. FA8750-19-2-0222) and Learning with Less Labels (LwLL) program (Contract No.FA8750-19-C-0515). Additional support was provided by UBC's Composites Research Network (CRN), Data Science Institute (DSI) and Support for Teams to Advance Interdisciplinary Research (STAIR) Grants. This research was enabled in part by technical support and computational resources provided by WestGrid (https://www.westgrid.ca/) and Compute Canada (www.computecanada.ca). William Harvey acknowledges support by the University of British Columbia's Four Year Doctoral Fellowship (4YF) program.

References

Ballester, C., Bertalmio, M., Caselles, V., Sapiro, G., and Verdera, J. Filling-in by joint interpolation of vector fields and gray levels. *IEEE transactions on image processing*, 10(8):1200–1211, 2001.

Barnes, C., Shechtman, E., Finkelstein, A., and Goldman, D. B. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3):24, 2009.

Barratt, S. and Sharma, R. A note on the inception score. *arXiv preprint arXiv:1801.01973*, 2018.

Bertalmio, M., Sapiro, G., Caselles, V., and Ballester, C. Image inpainting. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pp. 417–424, 2000.

Bertalmio, M., Bertozzi, A. L., and Sapiro, G. Navier-stokes, fluid dynamics, and image and video inpainting. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, volume 1, pp. I–I. IEEE, 2001.

Biewald, L. Experiment tracking with weights and biases, 2020. URL https://www.wandb.com/. Software available from wandb.com.

Bishop, C. M. *Pattern recognition and machine learning*. springer, 2006.

Chaloner, K. and Verdinelli, I. Bayesian experimental design: A review. *Statistical Science*, pp. 273–304, 1995.

Child, R. Very deep vaes generalize autoregressive models and can outperform them on images. *arXiv preprint arXiv:2011.10650*, 2020.

Criminisi, A., Perez, P., and Toyama, K. Object removal by exemplar-based inpainting. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, pp. II–II. IEEE, 2003.

Garnelo, M., Schwarz, J., Rosenbaum, D., Viola, F., Rezende, D. J., Eslami, S., and Teh, Y. W. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018.

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial networks. *arXiv preprint arXiv:1406.2661*, 2014.

Gregor, K., Danihelka, I., Graves, A., Rezende, D., and Wierstra, D. Draw: A recurrent neural network for image generation. In *International Conference on Machine Learning*, pp. 1462–1471. PMLR, 2015.

Harvey, W., Teng, M., and Wood, F. Near-optimal glimpse sequences for improved hard attention neural network training. *arXiv preprint arXiv:1906.05462*, 2019.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *arXiv preprint arXiv:1706.08500*, 2017.

Iizuka, S., Simo-Serra, E., and Ishikawa, H. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)*, 36(4):1–14, 2017.

Ivanov, O., Figurnov, M., and Vetrov, D. Variational autoencoder with arbitrary conditioning. *arXiv preprint arXiv:1806.02382*, 2018.

Karras, T., Laine, S., and Aila, T. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4401–4410, 2019.

Kim, H., Mnih, A., Schwarz, J., Garnelo, M., Eslami, A., Rosenbaum, D., Vinyals, O., and Teh, Y. W. Attentive neural processes. *arXiv preprint arXiv:1901.05761*, 2019.

Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL http://arxiv.org/abs/1412.6980.

Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improving variational inference with inverse autoregressive flow. *arXiv preprint arXiv:1606.04934*, 2016.

Klushyn, A., Chen, N., Kurle, R., Cseke, B., and van der Smagt, P. Learning hierarchical priors in vaes. *arXiv preprint arXiv:1905.04982*, 2019.

Köhler, R., Schuler, C., Schölkopf, B., and Harmeling, S. Mask-specific inpainting with deep neural networks. In *German conference on pattern recognition*, pp. 523–534. Springer, 2014.

Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.

Levin, A., Zomet, A., and Weiss, Y. Learning how to inpaint from global image statistics. In *ICCV*, volume 1, pp. 305–312, 2003.

Li, Y., Liu, S., Yang, J., and Yang, M.-H. Generative face completion. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3911–3919, 2017.

Loshchilov, I. and Hutter, F. SGDR: stochastic gradient descent with warm restarts. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL https://openreview.net/forum?id=Skq89Scxx.

Ma, C., Tschiatschek, S., Palla, K., Hernández-Lobato, J. M., Nowozin, S., and Zhang, C. Eddi: Efficient dynamic discovery of high-value information with partial vae. *arXiv preprint arXiv:1809.11142*, 2018.

Minka, T. et al. Divergence measures and message passing. Technical report, Microsoft Research, 2005.

Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., and Efros, A. A. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2536–2544, 2016.

Ren, J. S., Xu, L., Yan, Q., and Sun, W. Shepard convolutional neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 1*, pp. 901–909, 2015.

Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X., and Chen, X. Improved techniques for training gans. In *Advances in Neural Information Processing Systems 29*, pp. 2234–2242. Curran Associates, Inc., 2016.

Sohn, K., Lee, H., and Yan, X. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28: 3483–3491, 2015.

Sønderby, C. K., Raiko, T., Maaløe, L., Sønderby, S. K., and Winther, O. Ladder variational autoencoders. *arXiv preprint arXiv:1602.02282*, 2016.

Song, Y., Yang, C., Shen, Y., Wang, P., Huang, Q., and Kuo, C.-C. J. Spg-net: Segmentation prediction and guidance network for image inpainting. *arXiv preprint arXiv:1805.03356*, 2018.

Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.

Vahdat, A. and Kautz, J. Nvae: A deep hierarchical variational autoencoder. *arXiv preprint arXiv:2007.03898*, 2020.

Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., and Summers, R. M. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2097–2106, 2017.

Yeh, R. A., Chen, C., Yian Lim, T., Schwing, A. G., Hasegawa-Johnson, M., and Do, M. N. Semantic image inpainting with deep generative models. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5485–5493, 2017.

Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., and Huang, T. S. Generative image inpainting with contextual attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5505–5514, 2018.

Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., and Huang, T. S. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4471–4480, 2019.

Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 586–595, 2018.

Zhao, L., Mo, Q., Lin, S., Wang, Z., Zuo, Z., Chen, H., Xing, W., and Lu, D. Uctgan: Diverse image inpainting based on unsupervised cross-space translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5741–5750, 2020.

Zheng, C., Cham, T.-J., and Cai, J. Pluralistic image completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1438–1447, 2019.

Zhu, J.-Y., Zhang, R., Pathak, D., Darrell, T., Efros, A. A., Wang, O., and Shechtman, E. Toward multimodal image-to-image translation. *arXiv preprint arXiv:1711.11586*, 2017.

A. Bayesian optimal experimental design

A.1. EIG estimators

In this section, we expand on how our estimator for the expected information gain differs from prior work (Harvey et al., 2019). Repeating some relevant definitions from the main text, $f_{c_1,\ldots,c_t}$ is a function which maps from a full image to a sequence of observations of the image at coordinates $c_1, \ldots, c_t$. When estimating the EIG at time step t, we will have already observed a sequence of observations denoted $\hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}}$, extracted from some latent image $\mathcal{I}$ as $f_{c_1,\ldots,c_{t-1}}(\mathcal{I})$. We use a CNN which outputs $g(v|f_{c_1,\ldots,c_t}(\mathcal{I}))$, an approximation of the posterior over the latent variable of interest given a sequence of observations. Following Harvey et al. (2019), we will from now on refer to this as the AVP-CNN (attentional variational posterior CNN). Both our method and that of Harvey et al. (2019) sample image completions, which we denote $\mathcal{I}^{(1)}, \ldots, \mathcal{I}^{(N)}$, conditioned on $\hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}}$. Harvey et al. (2019) do so by retrieving images which roughly match the observed pixel values from a database. Although completions from this are diverse, they can match the observed values poorly. We therefore replace this stage with an image completion network trained using AIPO.

Given these components, Harvey et al. (2019) approximate the expected information gain with

$$\mathrm{EIG}(c_t;\hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}}) \approx \overbrace{\mathcal{H}\left[g(\cdot|\hat{\mathcal{I}}_{1,\ldots,c-1})\right]}^{\text{entropy after } t-1 \text{ scans}} - \underbrace{\frac{1}{N}\sum_{n=1}^{N}\mathcal{H}\left[g(\cdot|f_{c_1,\ldots,c_t}(\mathcal{I}^{(n)}))\right]}_{\text{expected entropy after } t \text{ scans}}. \tag{18}$$

Since the prior entropy term of Eq. (18) does not depend on c_t, it can be neglected when choosing c_t with BOED. As only the expected posterior entropy then needs to be estimated and compared for various c_t, we will from now on refer to this estimator with the acronym EPE (for 'expected posterior entropy').

We use a different estimator for the prior entropy. It becomes equivalent as $N \to \infty$ if g and the image completion method are perfect but, when this is not the case, produces an estimate of the prior entropy with a dependence on c_t. To repeat Eq. (17), this results in the following estimator for the EIG:

$$\mathrm{EIG}(c_t;\hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}}) \approx \overbrace{\mathcal{H}\left[\frac{1}{N}\sum_{n=1}^{N} g(\cdot|f_{c_1,\ldots,c_t}(\mathcal{I}^{(n)}))\right]}^{\text{entropy after } t-1 \text{ scans}} - \underbrace{\frac{1}{N}\sum_{n=1}^{N}\mathcal{H}\left[g(\cdot|f_{c_1,\ldots,c_t}(\mathcal{I}^{(n)}))\right]}_{\text{expected entropy after } t \text{ scans}}. \tag{19}$$

While it is not immediately clear that this estimator is better than that in Eq. (18), one useful property is that, like the true expected information gain, it is guaranteed to be non-negative. This can be seen as follows. First, note that the prior entropy term is the entropy of an expectation (over $n \sim \mathrm{Uniform}(1, N)$), and the expected posterior entropy term swaps the order of the entropy and expectation. Since the mapping from a distribution $g(\cdot)$ to its entropy $\mathcal{H}[g(\cdot)]$ is strictly concave, Jensen's inequality is applicable. Invoking Jensen's inequality shows that the prior entropy term must be greater than or equal to the expected posterior entropy, and so the estimate of the EIG will be non-negative. In Fig. 7, we compare the performance of this estimator (denoted EIG) and that in Eq. (18) (EPE). We find that the estimator denoted EIG leads to significantly better classification performance.

Another view of our EIG estimator is as a variant of the inception score (Salimans et al., 2016) for the sampled images $\mathcal{I}^{(1)}, \ldots, \mathcal{I}^{(N)}$. The standard inception score is computed using an image classifier trained on ImageNet (Salimans et al., 2016). Equation (17) is the inception score if this classifier is replaced with the AVP-CNN acting on observations of each image at coordinates $c_1, \ldots, c_t$.

A.2. BOED baselines

Random This baseline simply samples the scan coordinate independently at each time step from a uniform distribution over all valid locations.

Uncond This baseline ablates our estimator for the EIG (Eq. (17)) by sampling images $\mathcal{I}^{(1)}, \ldots, \mathcal{I}^{(N)}$ from the training dataset (without conditioning on $\hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}}$) instead of from AIPO (conditioned on $\hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}}$). That is, we use:

$$\mathrm{U}^{\mathrm{uncond}}(c_t; c_1, \ldots, c_{t-1}) = \mathcal{H}\left[\frac{1}{N}\sum_{n=1}^{N} g(\cdot|f_{c_1,\ldots,c_t}(\mathcal{I}^{(n)}))\right] - \frac{1}{N}\sum_{n=1}^{N}\mathcal{H}\left[g(\cdot|f_{c_1,\ldots,c_t}(\mathcal{I}^{(n)}))\right]. \tag{20}$$

with $\mathcal{I}^{(1)}, \ldots, \mathcal{I}^{(N)} \sim p(\mathcal{I})$. The resulting choice of scan coordinates c_t at each time step t has no dependence on previous observations $\hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}}$.

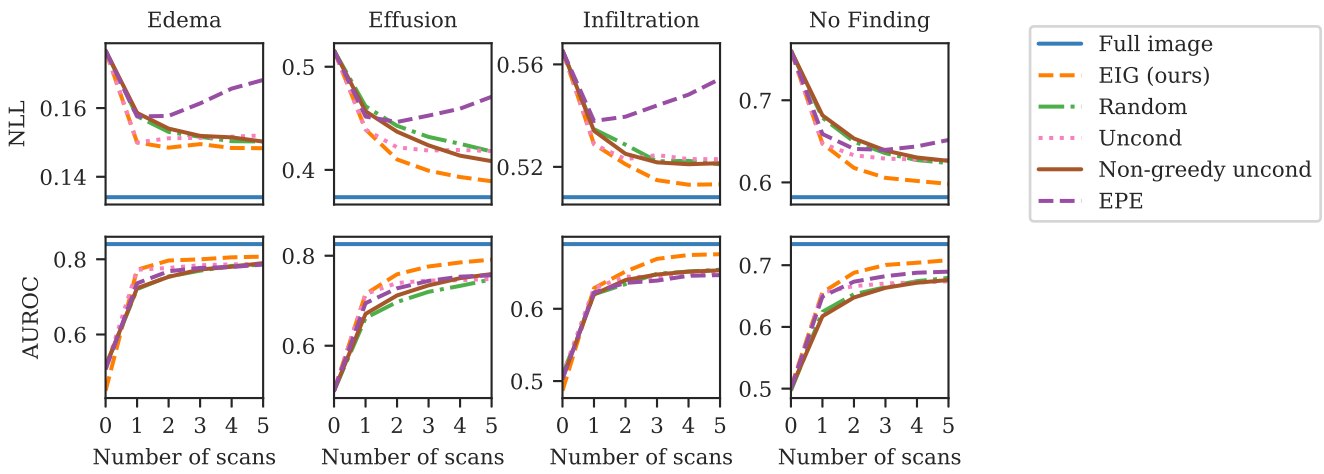

Figure 7. We extend on Fig. 6 by additionally reporting negative log-likelihoods (top row) for the classification tasks. We also include results from the 'EPE' estimator used by Harvey et al. (2019). We see that our BOED estimator provides significant improvements; resulting performance is always superior to the other baselines. This is only sometimes the case for the 'EPE' estimator.

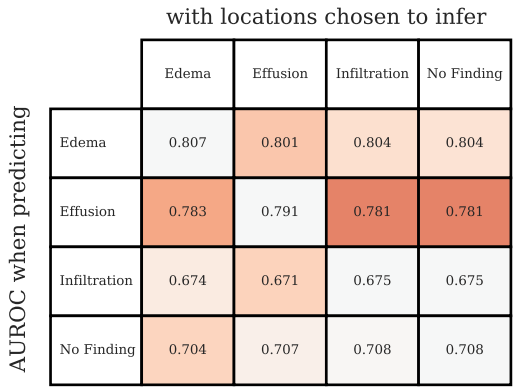

Figure 8. AUROC scores when performing one classification task with scan locations chosen for a different classification task. The intensity of the colour of each cell is proportional to the difference between its value and the greatest value in its row.

Non-greedy uncond When selecting scan locations that optimise the EIG, we select the scan location at each time step greedily to optimise the EIG from the next scan. To do otherwise (i.e. maximise the EIG from multiple next scans) is intractable because of the dependence of the EIG at each t on previous observations $\hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}}$. However, since the above 'Uncond' baseline selects each c_t independently of what is observed at previous scans, we can remove this greedy property. We do so by selecting all $c_1, \ldots, c_T$ simultaneously to maximize

$$\mathrm{U}^{\text{non-greedy}}(c_1, \ldots, c_T) = \mathcal{H}\left[\frac{1}{N}\sum_{n=1}^{N} g(\cdot|f_\emptyset(\mathcal{I}^{(n)}))\right] - \frac{1}{N}\sum_{n=1}^{N}\mathcal{H}\left[g(\cdot|f_{c_1,\ldots,c_T}(\mathcal{I}^{(n)}))\right]. \quad (21)$$

with $\mathcal{I}^{(1)}, \ldots, \mathcal{I}^{(N)} \sim p(\mathcal{I})$. Since the number of possible values of $c_1, \ldots, c_T$ increases exponentially with T, it is not feasible to search over all of them. We therefore optimize this coordinate sequence by randomly sampling a set of such sequences, and choosing the best from this set.

A.3. Additional BOED results

Figure 7 shows additional results from classification with sequences chosen by BOED and our baselines. In particular, we report negative log-likelihoods for the class labels as well as the AUROC scores. We see that the negative log-likelihood sometimes increases as more scans are taken, indicating that the classifier used is not perfectly trained. Interestingly, this only happens when glimpses are chosen with the EPE estimator of Harvey et al. (2019). This may be because the EPE estimator has a tendency to select coordinates which cause the classifier to be overconfident (and thus produce artificially low posterior entropies).

To illustrate that the coordinates chosen by BOED are task-dependent, Fig. 8 shows the results from using scan locations chosen for one classification task (i.e. diagnosing a particular disease) to perform a different classification task. Each row shows the results for a particular task, with each column showing results when scan coordinates are chosen

based on a particular task. As expected, the highest (or at least joint-highest) AUROC score in each row occurs on the diagonal, when the classification task and choice of scan locations are aligned. Away from the diagonal, the scores are only slightly lower, perhaps reflecting that there is a large overlap between the relevant areas for these tasks.

Figure 9 shows additional visualisations of the Bayesian experimental design process for different test images, adding to the one shown in the main paper. The 'information gains' shown are computed similarly to the EIG but the 'expected entropy after t scans' term is replaced by the entropy given a particular sampled $\mathcal{I}^*$. That is,

$$\mathrm{IG}(c_t; \hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}}) \approx \overbrace{\mathcal{H}\left[\frac{1}{N}\sum_{n=1}^{N} g(\cdot|f_{c_1,\ldots,c_t}(\mathcal{I}^{(n)}))\right]}^{\text{entropy after } t-1 \text{ scans}} - \underbrace{\mathcal{H}\left[g(\cdot|f_{c_1,\ldots,c_t}(\mathcal{I}^*))\right]}_{\text{entropy after } t \text{ scans given } \mathcal{I}^*}. \tag{22}$$

where $\mathcal{I}^*$ is a sampled conditioned on observations $\hat{\mathcal{I}}_{c_1,\ldots,c_{t-1}}$. Averaging information gains computed with $\mathcal{I}^* = \mathcal{I}^{(1)}$ to $\mathcal{I}^{(N)}$ gives the estimate of the EIG presented in Eq. (17).

A.4. BOED experimental details

The AVP-CNN, g, is trained to map from masked images, $\hat{\mathcal{I}}$, to distributions over the class labels. Each image in the Chest X-ray 14 dataset has labels indicating the presence or absence of each of 14 pathologies. We train g to produce 15 outputs: an estimated probability of the presence or absence of each of the 14 conditions individually; and an additional estimate of the probability that any (one or more) of these conditions is present. We train g to estimate these using a cross-entropy loss. Masked images are sampled using almost the same mask distribution as for training the image completion networks (described in Section 4); the only difference is that patches now have 25% rather than 35% of the image width, to match the observations we use in the experiments with BOED. We use an AVP-CNN pretrained on ImageNet and then trained on Chest X-ray 14 for 32 000 iterations with a batch size of 32 and learning rate 1×10^{-5}.

To use AIPO for the Chest X-ray dataset, we first trained an unconditional VAE. While AIPO allows us to train all components from scratch, training this unconditional model allowed us to speed up later experimentation with AIPO. The unconditional VAE was trained for 200 000 iterations (about 5 days) and the checkpoint with the best validation ELBO was used. The hyperparameters used were the same as those of Child (2020) on FFHQ-256, but with a larger skip threshold of 15 000[1] and a smaller batch size (of 8 rather than 32). We trained AIPO for 180 000 iterations (about 60 hours on 4 GPUs) using this unconditional model with a batch size of 8 and the original skip threshold of 100 used by Child (2020).

Selecting each scan location was done by evaluating Eq. (17) at every point in an evenly-spaced 17×17 grid over the image, and choosing the maximum. We evaluate the EIG with $N = 10$ sampled image completions. This is repeated to select the scan location for each $t = 1, \ldots, 5$ (with the sampled images conditioned on observations up to $t-1$ at each stage). Where applicable, we use the same trained networks and hyperparameters for the baselines. For the 'non-greedy uncond' baseline, we select a coordinate sequence by choosing the best of 289 sampled sequences. This number was chosen to match the number of locations in the 17×17 grid that the other methods search over at each time step. The results reported in Figs. 6 and 7 were computed on a randomly-sampled 5000 image subset of the Chest X-ray 14 test set.

[1] Child (2020) improve training stability by skipping updates with gradients larger than a set threshold. Using a larger threshold, and so skipping fewer updates, was found to be necessary to train their architecture on the Chest X-ray dataset.

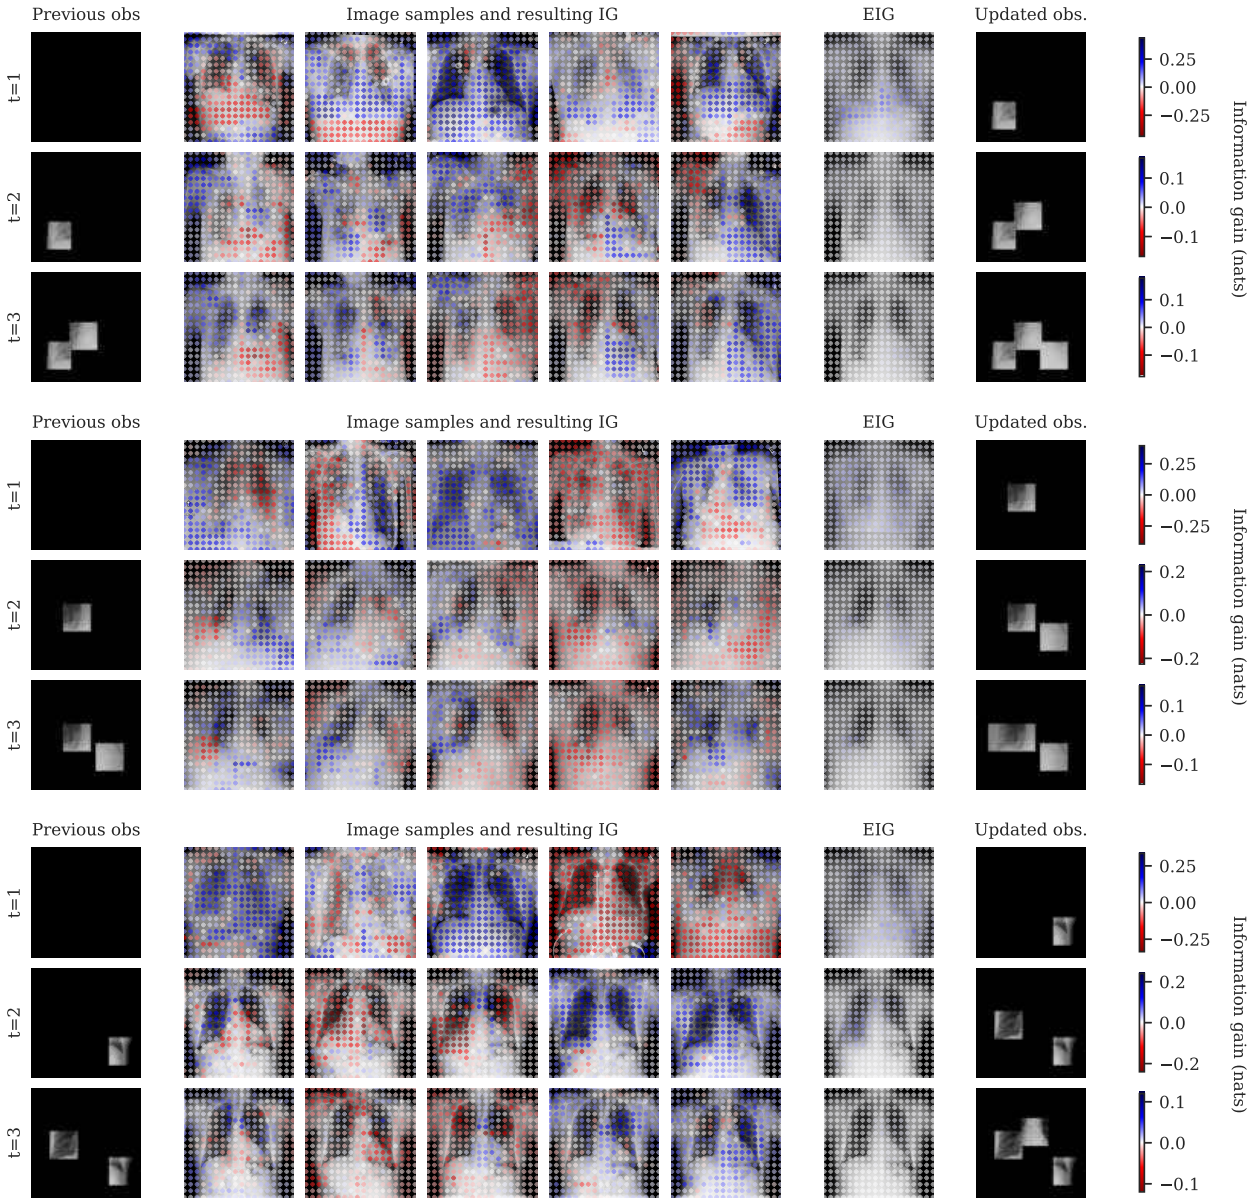

Figure 9. Visualisations in the style of Fig. 5 for three more test images. Each shows the experimental design process to choose three scan locations in order to detect whether 'Effusion' is present. As in Fig. 5, the left column shows the observations made before selecting each scan location. The next five columns show five of the $N = 10$ sampled image completions conditioned on these observations. Each is overlaid with the information gain predicted after scanning any location. The second column from the left shows the expected information gain at each location, given by averaging the information gains arising from each sampled image completion. The final column shows the updated observation after scanning the location which maximises the expected information gain.

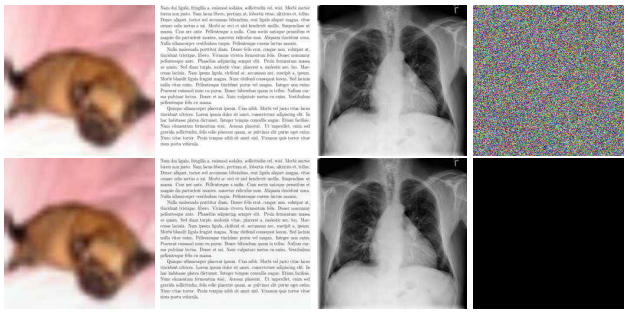

Figure 10. Top row: a CIFAR-10 image, an image of text, an image from Chest X-ray 14, and an image with all pixel values independently sampled from a uniform distribution, all at 256×256 resolution. Bottom row: reconstructions using the encoder and decoder of a VD-VAE trained on images of faces. All but the uniformly sampled image are faithfully reconstructed.

B. Reconstructions of o.o.d. images

Figure 10 shows images from various domains, all reconstructed by a VAE trained on a different domain (the FFHQ-256 dataset). The fact that the VD-VAE architecture can reconstruct most of these images alludes to its mass-covering property.

C. Experimental details

C.1. AIPO/AIPO-R

Architectures The AIPO encoder and decoder architectures we used for CIFAR-10 and FFHQ-256 were the same as those used by Child (2020) for the same datasets, with 45 and 62 groups of latent variables respectively. The encoder and decoder had 39 million parameters for CIFAR-10 and 115 million for FFHQ-256. We used partial encoders with structure identical to the encoders, other than an additional input channel to accept the concatenated mask. The partial encoders contained 18 million and 65 million parameters respectively for CIFAR-10 and FFHQ-256. The architecture used for Chest X-ray 14 was identical to that used for FFHQ-256. We use the same architecture for both AIPO and AIPO-R.

Training The AIPO and AIPO-R networks on CIFAR-10 were each trained for 160 000 iterations with a batch size of 14. This took 17 hours for AIPO and 21 hours for AIPO-R. The learning rates used were selected with a sweep over three values and were 2×10^{-4} for AIPO and 2×10^{-5} for AIPO-R. On FFHQ-256, AIPO and AIPO-R were trained for 240 000 and 160 000 iterations respectively, taking 141 and 110 hours. These used learning rates of 1.5×10^{-4} and 5×10^{-5}. As for CIFAR-10, these learning rates were selected with a search over three values. Other hyperparameters are the same as those used by Child (2020) for training VAEs on these datasets. We use the Weights & Biases (Biewald, 2020) experiment-tracking infrastructure.

C.2. Baselines

For all the baselines, we based our implementations on publicly available (official or unofficial) implementations. A link to these repositories is provided for each baseline. All the training procedures were modified to use the same distribution of partially observed images as for training AIPO (see Section 4).

PIC We adapted the official implementation of Zheng et al. (2019)[2] for the FFHQ-256 and CIFAR-10 datasets. We used the same architecture, hyperparameters and objective functions as their implementation. The networks for both CIFAR-10 and FFHQ-256 had 9,128,390 parameters. They were trained using Adam optimizer (Kingma & Ba, 2015) with a batch size of 20, a learning rate of 10^{-4} and β_1, β_2 parameters of 0.0 and 0.999 respectively. The model for FFHQ-256 was trained for 180 epochs and the model for CIFAR-10 was trained for 240 epochs. It took 118 hours to train each of these models on an RTX 2080 Ti GPU.

The PIC architecture is defined for 256×256 images. Therefore, in order to test it on the CIFAR-10 dataset, we resized the 32×32 CIFAR-10 images to 256×256 (via bilinear interpolation) before feeding them into the PIC model. We then down-sampled the inpainted images back to the original size of 32×32 for evaluation.

ANP Our ANP network architecture was based on that of Kim et al. (2019)[3], differing only in that we used hidden and latent dimensions of 512. Both the CIFAR-10 and FFHQ-256 models had 11,563,526 parameters and were trained for 220 epochs with a batch size of 16 using Adam optimizer. Learning rates of 5×10^{-5} and 10^{-4} was used for the CIFAR-10 and FFHQ-256 models respectively. Training took 29 hours for CIFAR-10 and 43 hours for FFHQ-256 on a Tesla V100 GPU.

In the original image inpainting experiments of (Kim et al., 2019) the images are 32×32. Since ANPs embed each pixel separately and their self-attention and cross-attention layers attend to every other pixel in the observed and target sets, it is computationally expensive to scale them to larger images. We therefore downsampled the FFHQ-256 images to 64×64, and upsampled the inpainted image back to 256×256 via bilinear interpolation. Additionally, at training time, we randomly dropped half of the observed pixels (a.k.a. context set in neural process literature) to reduce the computational

[2] https://github.com/lyndonzheng/Pluralistic-Inpainting

[3] Our implementation of ANP is based on https://github.com/EmilienDupont/neural-processes

cost. Finally, the target set at training time was half of the unobserved pixels. For CIFAR-10, on the other hand, no modfication was done to the image resolution or observation masks i.e. images were fed in at 32×32 resolution, the context sets were the set of all observed pixels in partial images $\hat{\mathcal{I}}$, and the target sets consisted of all unobserved pixels.

CE We used the same architecture as reported by Pathak et al. (2016)[4]. The models for CIFAR-10 and FFHQ-256 had 34,094,528 and 71,321,536 parameters respectively. Both models were trained for 500 epochs with a batch size of 64. We used Adam optimizer with a learning rate of 2×10^{-4}, $\beta_1 = 0.5$, and $\beta_2 = 0.999$. It took 5 hours to train the CIFAR-10 model and 48 hours to train the FFHQ-256 model on a RTX 2080 Ti GPU.

C.3. Inception score

Definition The inception score is defined based on an image classifier. A high score requires the individual samples to be classifiable with high confidence and, at the same time, the marginal distribution of the samples' classes to be diverse. More formally, let G be a generator producing samples $\boldsymbol{x} \in \mathcal{X}$ and let the classifier output $p(y|\boldsymbol{x})$ where $y \in \mathcal{Y}$ denotes the classification labels specified by the classifier. Then the logarithm of the inception score is

$$\begin{aligned} \log \mathrm{IS} &= \mathbb{E}_{\boldsymbol{x}\sim G}\left[\mathrm{KL}\big(p(y|\boldsymbol{x}) \,\big\|\, p(y)\big)\right] \\ &= -\mathbb{E}_{\boldsymbol{x}\sim G}\left[H(p(y|\boldsymbol{x})) - H(p(y|\boldsymbol{x}), p(y))\right]. \end{aligned} \tag{23}$$

From an information-theoretic perspective, the inception score is the mutual information between class labels and generated images $I(\boldsymbol{x}; y)$.

Classifier network Results in the GAN literature suggest that the inception score is unreliable when applied to image domains other than ImageNet. We therefore replaced the standard, ImageNet-trained, Inception network with a classifier trained on CIFAR-10. In particular, we used a ResNet-18 (He et al., 2016) architecture, pretrained on ImageNet[5]. We fine-tuned the model on CIFAR-10 for 200 epochs using stochastic gradient descent with an initial learning rate of 0.1 and a cosine annealing learning rate scheduler (Loshchilov & Hutter, 2017). We evaluated the classifier after each training epoch and picked the one with the highest test accuracy as the final classifier. The best accuracy of 95.15% correct classification rate was achieved after 198 epochs of training.

[4]Our implementation of CE is based on https://github.com/BoyuanJiang/context_encoder_pytorch, with some modifications according to https://github.com/pathak22/context-encoder to support larger image sizes and non-centered observation masks.

[5]Our implementation of the CIFAR-10 classifier is based on github.com/kuangliu/pytorch-cifar.

D. Image samples

Chest X-ray 14 Figure 11 shows sampled completions of x-ray images from AIPO.

FFHQ-256 Figure 12 on page 19 shows sampled image completions from each method (AIPO, AIPO-R, and all baselines) for an FFHQ-256 image. The long form of the appendix[6] contains more completions for different FFHQ-256 images. For each image, completions from six different observation masks (with zero to five observed patches) are shown. For each mask, we plot five sampled completions from each stochastic method and the single deterministic completion from CE.

CIFAR-10 Figure 13 (on page 20) shows sampled completions for various observations from a CIFAR-10 image, formatted in the same way as Figure 12. See the long form of the appendix for completions of different CIFAR-10 images.

As is apparent from the image samples, AIPO-R fails to generate meaningful completions when the number of observed pixels is small. It also has less sample diversity than AIPO, which can be explained by the mode-seeking behaviour of the reverse-KL and subsequent under-estimation of posterior uncertainty (Minka et al., 2005).

Finally, although AIPO generally produces more realistic images than the others, it is still far from perfect. In particular, a lack of bilateral symmetry is visible in the FFHQ-256 samples, especially with smaller numbers of observed pixels. This issue can be seen in unconditional samples from the underlying VAEs as well (see the samples in Child (2020)). Therefore, it is likely that any future advances in image modeling with VAEs could be integrated to improve this aspect of the results.

Free-form observation masks We also investigated how the models generalise to observation masks which are outside the training distribution (as defined by $p(\hat{\mathcal{I}}|\mathcal{I})$). To do so, we created masks manually. The results can be seen in Figure 14 (on page 20). More available in the long form of the appendix. All methods generalise reasonably well to these masks, although better results may be obtained by changing the mask distribution use in training.

[6]https://www.cs.ubc.ca/~wsgh/icidgm-appendix.pdf

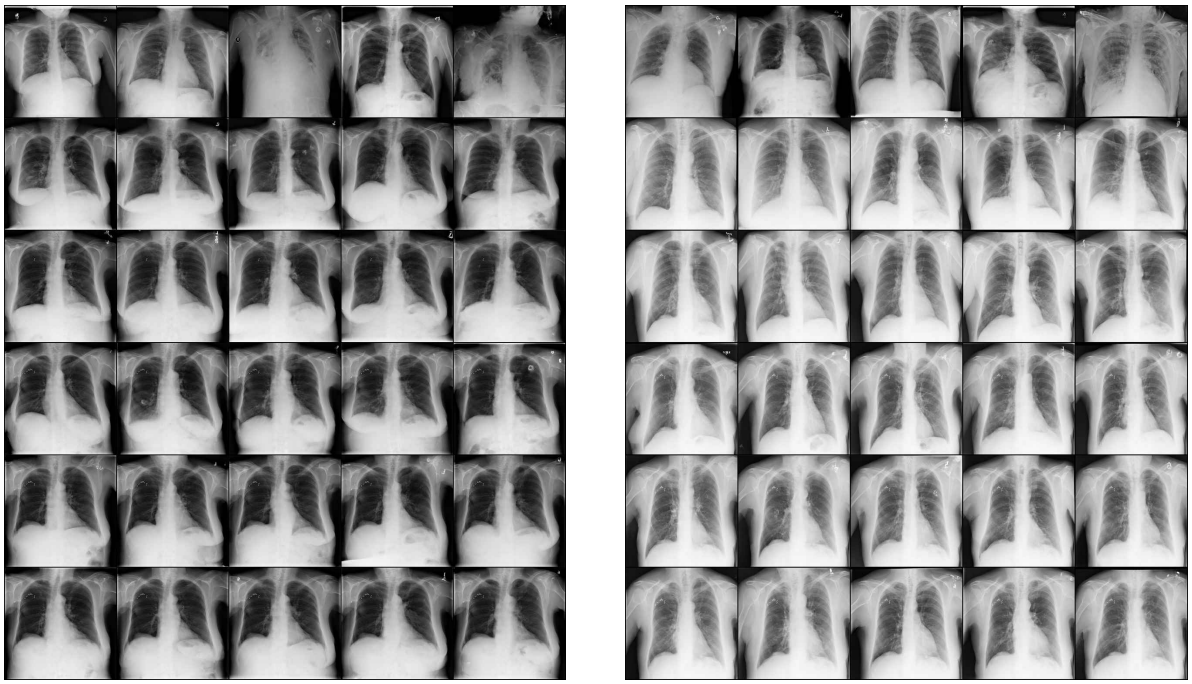

Figure 11. Two panels of sampled image completions from AIPO for two different ground truth images from the Chest X-ray 14 dataset. Each row shows completions based on a different number of observed patches, from zero (top) to five (bottom). The left column of each panel shows the ground truth image. The second column from the left shows the observations, and the remaining columns show five independent samples from AIPO.

E. Code

We have attached code for training AIPO and AIPO-R. We will later release code for using this with Bayesian experimental design. This will all be made publicly available, along with weights for networks trained with AIPO. Additionally, we will make the weights for the baseline models and the classifier used for inception score computation publicly available.

(a) $\mathcal{I}, \hat{\mathcal{I}}$ (b) AIPO (c) AIPO-R

(d) $\mathcal{I}, \hat{\mathcal{I}}$ (e) PIC (f) ANP (g) CE

Figure 12. Sampled completions from each method for an FFHQ-256 image. Panels (a) and (d) both show the true image, and the observation mask on which the samples in each row are conditioned. Panels (b) and (c) show completions from AIPO and AIPO-R. The remaining panels show completions from the baselines: (e) and (f) show five completions for the stochastic baselines PIC and ANP. Panel (g) shows the single deterministic completion produced by CE.

(a) $\mathcal{I}, \hat{\mathcal{I}}$ (b) AIPO (c) AIPO-R

(d) $\mathcal{I}, \hat{\mathcal{I}}$ (e) PIC (f) ANP (g) CE

Figure 13. Sampled completions from each method for a CIFAR-10 image. Panels (a) and (d) both show the true image, and the observation mask on which the samples in each row are conditioned. Panels (b) and (c) show completions from AIPO and AIPO-R. The remaining panels show completions from the baselines: (e) and (f) show five completions for the stochastic baselines PIC and ANP. Panel (g) shows the single deterministic completion produced by CE.

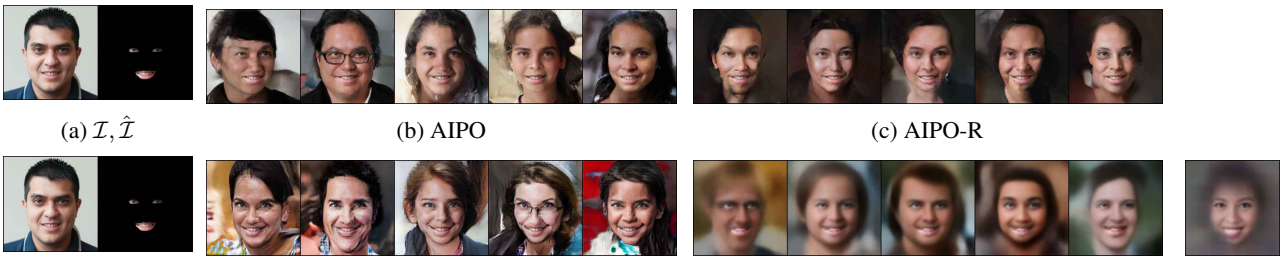

(a) $\mathcal{I}, \hat{\mathcal{I}}$ (b) AIPO (c) AIPO-R

(d) $\mathcal{I}, \hat{\mathcal{I}}$ (e) PIC (f) ANP (g) CE

Figure 14. Sampled completions from each method for an FFHQ-256 image with a free-form observation mask. Panels (a) and (d) both show the true image, and the observation mask on which the samples in each row are conditioned. Panels (b) and (c) show completions from AIPO and AIPO-R. The remaining panels show completions from the baselines: (e) and (f) show five completions for the stochastic baselines PIC and ANP. Panel (g) shows the single deterministic completion produced by CE.