

---

# Parallel and Flexible Sampling from Autoregressive Models via Langevin Dynamics

---

Vivek Jayaram<sup>\*1</sup> John Thickstun<sup>\*1</sup>

## Abstract

This paper introduces an alternative approach to sampling from autoregressive models. Autoregressive models are typically sampled sequentially, according to the transition dynamics defined by the model. Instead, we propose a sampling procedure that initializes a sequence with white noise and follows a Markov chain defined by Langevin dynamics on the global log-likelihood of the sequence. This approach parallelizes the sampling process and generalizes to conditional sampling. Using an autoregressive model as a Bayesian prior, we can steer the output of a generative model using a conditional likelihood or constraints. We apply these techniques to autoregressive models in the visual and audio domains, with competitive results for audio source separation, super-resolution, and inpainting.

## 1. Introduction

Neural autoregressive models (Larochelle & Murray, 2011) are a popular family of generative models, with wide-ranging applications in a variety of domains including audio (van den Oord et al., 2016a; Dhariwal et al., 2020), images (van den Oord et al., 2016b; Salimans et al., 2017; Parmar et al., 2018; Razavi et al., 2019), and text (Radford et al., 2019; Brown et al., 2020). These models parameterize the conditional distribution over a token in an ordered sequence, given previous tokens in the sequence. The standard approach to sampling from an autoregressive model iteratively generates tokens, according to a conditional distribution over tokens defined by the model, conditioned on the partial sequence of previously generated tokens. We will refer to this approach to sampling as the ancestral sampler.

There are two major drawbacks to ancestral sampling that

---

<sup>\*</sup>Equal contribution <sup>1</sup>Department of Computer Science, University of Washington. Correspondence to: Vivek Jayaram <vjayaram@cs.washington.edu>, John Thickstun <thickstn@cs.washington.edu>.

limit the usefulness of autoregressive models in practical settings. First, ancestral sampling has time complexity that scales linearly in the length of the generated sequence. For data such as high-resolution images or audio, ancestral sampling from an autoregressive model (where the tokens are pixels or sound pressure readings respectively) can be impractically slow. Second, ancestral sampling is frustratingly inflexible. It is easy to sample the second half of a sequence conditioned on the first, but filling in the first half a sequence conditioned on the second naively requires training a new model that reverses the ordering of tokens in the autoregressive factorization. Conditioning on arbitrary subsets of tokens for tasks such as inpainting or super-resolution seems beyond reach of autoregressive modeling.

This paper introduces an alternative, parallel and flexible (PnF) sampler for autoregressive models that can be parallelized and steered using conditioning information or constraints.<sup>1</sup> Instead of sampling tokens sequentially, the PnF sampler initializes a complete sequence (with random tokens) and proceeds to increase the log-likelihood of this sequence by following a Markov chain defined by Langevin dynamics (Neal et al., 2011) on a smoothed log-likelihood. The smoothing temperature is cooled over time according to an annealing schedule informed by Song & Ermon (2019; 2020). Convergence time of this annealed Langevin dynamics is empirically independent of the sequence length and, generalizing Jayaram & Thickstun (2020), the PnF sampler can be applied to posterior log-likelihoods to incorporate conditional information into the sampling process.

The primary technical contribution of this paper is the development of the PnF sampler for discretized autoregressive models (Section 3.1). Our interest in these models is motivated by their success as unconditional models of audio waves (van den Oord et al., 2016a; Mehri et al., 2017; Dhariwal et al., 2020). Defined over a discrete lattice within a continuous space, these models occupy a middle ground between continuous and discrete models. For continuous models such as RNADE (Uria et al., 2013), PnF sampling can be directly applied as in Song & Ermon (2019); Jayaram

---

<sup>1</sup>Code and examples of PnF sampling are available at: <https://grail.cs.washington.edu/projects/pnf-sampling/>.

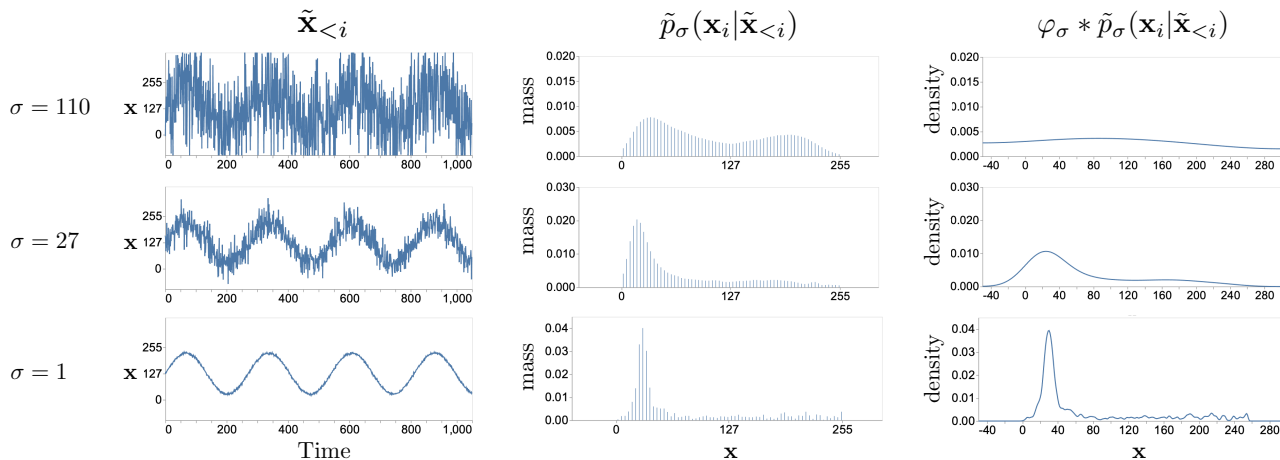


Figure 1. A visual summary of discretized autoregressive smoothing. Given a noisy history  $\tilde{\mathbf{x}}_{<i>} = \mathbf{x}_{<i>} + \boldsymbol{\varepsilon}_{<i>}$  (left column) where  $\boldsymbol{\varepsilon} \sim \mathcal{N}(0, \sigma^2 I)$ , we train a model to predict the un-noised distribution over  $\mathbf{x}_i \in \mathbb{R}$  (middle column). This distribution is discrete and non-differentiable in  $\tilde{\mathbf{x}}$ ; we convolve with a Gaussian  $\varphi_\sigma(t) = \mathcal{N}(t; 0, \sigma^2)$  to produce a continuous estimate of  $\tilde{\mathbf{x}}_i$  (right column). We can run Langevin dynamics on the continuous distribution, and gradually anneal the smoothing to approximate the target distribution.

& Thickstun (2020). We defer the development of the PnF sampler for fully discrete models to future work.

In Section 3.2, we present a stochastic variant of the PnF sampler based on stochastic gradient Langevin dynamics (Welling & Teh, 2011). This is an embarrassingly parallel, asynchronous distributed algorithm for autoregressive sampling. Using a WaveNet model, we show in Section 4.3 that stochastic PnF sampling approximates the quality of ancestral sampling to arbitrary accuracy, with compute time that is inverse proportional to the number of computing devices. This allows PnF sampling to take full advantage of modern, massively parallel computing infrastructure.

We will see in Section 3.3 how the PnF sampler can find solutions to general posterior sampling problems, using an unconditional generative model as a prior. In Section 4 we present applications of the PnF sampler to a variety of Bayesian image and audio inverse problems. We focus on linear inverse problems, using PixelCNN++ (Salimans et al., 2017) and WaveNet (van den Oord et al., 2016a) models as priors. Sections 4.4, 4.5, and 4.6 demonstrate PnF conditional sampling for source separation, super-resolution, inpainting respectively. PnF sampling results correlate strongly with the strength of the generative model used as a prior; as better autoregressive models are developed, they can be used with PnF sampling to improve performance on conditional generation tasks. We refer the reader to the project website for demonstrations of audio PnF sampling.

## 2. Related Work

The PnF autoregressive sampler is based on the annealed Langevin dynamics introduced in Song & Ermon (2019), which accelerates standard Langevin dynamics (Neal et al., 2011; Du & Mordatch, 2019) using a smoothing procedure

in the spirit of simulated annealing (Kirkpatrick et al., 1983) and graduated optimization (Blake & Zisserman, 1987). The extension of annealed Langevin dynamics to conditional sampling problems was discussed in Jayaram & Thickstun (2020) for source separation and image coloring problems, and developed further in Song et al. (2021) for general posterior sampling problems. The present work extends these methods to discretized autoregressive models, for which the smoothing procedures described in previous work are not directly applicable (Frank & Ilse, 2020). Markov-chain Monte Carlo posterior samplers based on Gibbs sampling rather than Langevin dynamics are proposed in Theis & Bethge (2015) and Hadjeres et al. (2017) as solutions for inpainting tasks.

The slow speed of ancestral sampling is a persistent obstacle to the adoption and deployment of autoregressive models. This has inspired algorithms that seek to parallelize the sampling process. Parallel WaveNet (van den Oord et al., 2018) and ClariNet (Ping et al., 2019) train generative flow models to mimic the behavior of an autoregressive model. Sampling a flow model requires only one pass through a feed-forward network and can be distributed across multiple devices. Wiggers & Hooeboom (2020) and Song et al. (2020) propose fixed-point algorithms that, like PnF sampling, iteratively refine an initial sample from a simple distribution into a sample from the target distribution. But none of these methods are easily adaptable to source separation (Section 4.4) or more general conditional sampling tasks.

Like anytime sampling (Xu et al., 2021), PnF sampling offers a tradeoff between sample quality and computational budget. The algorithm’s iterates gradually mix to the target distribution and, by stopping early, we can approximate samples from this distribution using less computation. We explore the empirical tradeoff between sample quality and

computation for PnF sampling from autoregressive models in Section 4.2. The anytime sampler proposed in Xu et al. (2021) requires a specific model architecture based on the VQ-VAE (van den Oord et al., 2017; Razavi et al., 2019). In contrast, the PnF sampler can be used with any likelihood-based model. Unlike an anytime sampler, the computational budget for PnF sampling must be specified in advance: halting prior to completing the annealing schedule will result in noisy samples.

Bayesian inverse problems are explored extensively in theoretical settings, where the prior is given by a simple analytical distribution (Tropp & Wright, 2010; Knapik et al., 2011; Wang et al., 2017). These problems have also been studied using learned priors given by GAN’s, with a focus on linear inverse problems (Rick Chang et al., 2017; Bora et al., 2017; Raj et al., 2019). These GAN-based approaches are tailored to the latent variable architecture of the model, performing latent space optimizations to find codes that correspond to desired outputs. There is no obvious extension of these latent variable approaches to autoregressive models.

While we focus on autoregressive models, due to their strong empirical performance as unconditional models of audio, PnF sampling could be applied more generally with other likelihood-based models. In the audio space, this includes recent diffusion models (Kong et al., 2021; Chen et al., 2021). Note however that audio vocoder models (Prenger et al., 2019; Kim et al., 2018; Ping et al., 2020), which rely on spectrogram conditioning, cannot be adapted as priors for the source separation, super-resolution, and inpainting experiments presented in Section 4. In addition, GAN based models (Donahue et al., 2019; Kumar et al., 2019), which are not likelihood based, cannot be sampled using PnF.

### 3. Parallel and Flexible Sampling

We want to sample from an autoregressive generative model over some indexed sequence of values  $\mathbf{x} \in \mathcal{X}^n$  where

$$p(\mathbf{x}) = \prod_{i=1}^n p(\mathbf{x}_i | \mathbf{x}_{<i}). \quad (1)$$

We are particularly interested in developing a sampler for discretized autoregressive models, where  $\mathcal{X} = \mathbb{R}$  and each conditional  $p(\mathbf{x}_i | \mathbf{x}_{<i})$  has support on a finite set of scalar values  $\mathcal{D} = \{e_1, \dots, e_d\} \subset \mathbb{R}$ . The set  $\mathcal{D}^n$  could represent, for example, an 8-bit encoding of an image or audio wave.

We propose to sample from  $p(\mathbf{x})$  via Langevin dynamics. Let  $\mathbf{x}_0 \sim \text{Uniform}(\mathbb{R}^n)$ ,  $\varepsilon_t \sim \mathcal{N}(0, I_n)$ , and define a Markov chain

$$\mathbf{x}^{(t+1)} \equiv \mathbf{x}^{(t)} + \eta \nabla_{\mathbf{x}} \log p(\mathbf{x}^{(t)}) + \sqrt{2\eta} \varepsilon_t \quad (2)$$

If  $p(\mathbf{x})$  were a smooth density then, for sufficiently small  $\eta$ , the Markov chain mixes and  $\mathbf{x}^{(t)}$  converges in distribution

---

#### Algorithm 1 Parallel and Flexible Sampling

---

**Input:**  $\mathbf{y}, \{\sigma_i\}_{i=1}^L, \delta, T$   
 Sample  $\mathbf{x}^{(0)} \sim \mathcal{N}(0, \sigma_1^2 I_n)$   
**for**  $i \leftarrow 1$  **to**  $L$  **do**  
      $\eta_i \leftarrow \delta \cdot \sigma_i^2 / \sigma_L^2$   
     **for**  $t = 1$  **to**  $T$  **do**  
         Sample  $\varepsilon_t \sim \mathcal{N}(0, I_n)$   
          $\mathbf{g}^{(t)} \leftarrow \nabla_{\mathbf{x}} \log p_{\sigma_i}(\mathbf{x}^{(t)}) + \nabla_{\mathbf{x}} \log p_{\sigma_i}(\mathbf{y} | \mathbf{x}^{(t)})$   
          $\mathbf{x}^{(t+1)} \leftarrow \mathbf{x}^{(t)} + \eta_i \mathbf{g}^{(t)} + \sqrt{2\eta} \varepsilon_t$   
     **end for**  
**end for**

---

to  $p$  as  $t \rightarrow \infty$ . But a discretized probability distribution defined over  $\mathcal{D}^n$  is not smooth; the gradient  $\nabla_{\mathbf{x}} \log p(\mathbf{x}^{(t)})$  is not even well-defined. In Section 3.1 we propose a smoothing of the discrete model  $p(\mathbf{x})$ , creating a differentiable density on which the Markov chain (2) can mix.

To support conditional generation, we then turn our attention to sampling from the posterior of a joint distribution  $p(\mathbf{x}, \mathbf{y}) = p(\mathbf{y} | \mathbf{x}) p(\mathbf{x})$ , where  $p(\mathbf{x})$  is an autoregressive model over  $\mathbf{x} \in \mathbb{R}^n$  and  $p(\mathbf{y} | \mathbf{x})$  is a conditional likelihood. Langevin dynamics for sampling from the posterior  $p(\mathbf{x} | \mathbf{y})$  are given by

$$\begin{aligned} \mathbf{x}^{(t+1)} &\equiv \mathbf{x}^{(t)} + \eta \nabla_{\mathbf{x}} \log p(\mathbf{x}^{(t)} | \mathbf{y}) + \sqrt{2\eta} \varepsilon_t \\ &= \mathbf{x}^{(t)} + \eta \nabla_{\mathbf{x}} \left( \log p(\mathbf{x}^{(t)}) + \log p(\mathbf{y} | \mathbf{x}^{(t)}) \right) + \sqrt{2\eta} \varepsilon_t. \end{aligned} \quad (3)$$

This is a convenient Markov chain for posterior sampling, because the partition function  $p(\mathbf{y})$  vanishes when we take the gradient. However, like  $p(\mathbf{x})$ , the posterior  $p(\mathbf{x} | \mathbf{y})$  is not smooth. In Section 3.3 we propose a smoothing of the joint distribution  $p(\mathbf{x}, \mathbf{y})$ , for which the posterior  $p(\mathbf{x} | \mathbf{y})$  is differentiable and the Markov chain (3) can mix.

Given a smoothing procedure parameterized by a temperature parameter, we appeal to the simulated annealing heuristic developed in Song & Ermon (2019) to turn down the temperature as the Markov chain (2) or (3) mixes. In contrast to classical Markov chain sampling, for which samples  $\mathbf{x}^{(t)}$  converge in distribution to  $p$ , annealed Langevin dynamics converges asymptotically to a single point distributed approximately according to  $p$ . Algorithm 1 describes these annealed Langevin dynamics given a smoothed prior  $p_{\sigma}(\mathbf{x})$  and smoothed likelihood  $p_{\sigma}(\mathbf{y} | \mathbf{x})$ . The structure of this algorithm is the same as the annealed Langevin dynamics presented in Song & Ermon (2019) (the unconditional case where  $\nabla_{\mathbf{x}} p_{\sigma}(\mathbf{y} | \mathbf{x}) = 0$ ) and Jayaram & Thickstun (2020) (the conditional case); the novel contribution of this paper is the smoothing algorithm for evaluating  $\nabla_{\mathbf{x}} p_{\sigma}(\mathbf{x})$  given  $p(\mathbf{x})$  (Section 3.1) and  $\nabla_{\mathbf{x}} p_{\sigma}(\mathbf{y} | \mathbf{x})$  given  $p(\mathbf{y} | \mathbf{x})$  (Section 3.3).

Each step of the Langevin dynamics described in Equations (2) or (3) requires inference of  $\log p(\mathbf{x})$ , an  $O(n)$  operation.

Unlike sequential sampling, calculating  $\log p(\mathbf{x})$  for a given sequence  $\mathbf{x} \in \mathbb{R}^n$  is embarrassingly parallel; for moderate sequence lengths  $n$ , the cost of computing  $\log p(\mathbf{x})$  is essentially constant using a modern parallel computing device. When  $n$  is very large, e.g. for WaveNet models where just a minute of audio has  $n > 10^6$  samples, it can be convenient to distribute sampling across multiple computing devices. In Section 3.2 we describe a stochastic variant of PnF sampling based on stochastic gradient Langevin dynamics that is easily distributed across a cluster of devices.

### 3.1. Autoregressive Smoothing

We consider models that parameterize the conditional distribution  $p(\mathbf{x}_i | \mathbf{x}_{<i})$  with a categorical softmax distribution over  $d = |\mathcal{D}|$  discrete values; given functions  $f_i : \mathcal{X}^i \rightarrow \mathbb{R}^d$ , we define

$$p(\mathbf{x}_i = e_k | \mathbf{x}_{<i}) = \frac{\exp(f_{i,k}(\mathbf{x}_{<i}))}{\sum_{\ell=1}^d \exp(f_{i,\ell}(\mathbf{x}_{<i}))}. \quad (4)$$

The functions  $f_i$  are typically given by a neural network, with shared weights across the sequential indices  $i$ . Collectively, these conditional models define the joint distribution  $p(\mathbf{x})$  according to Equation (1).

We cannot directly compute gradients of the distribution  $p(\mathbf{x})$  defined by discrete conditionals  $p(\mathbf{x}_i | \mathbf{x}_{<i})$ . Instead, we smooth  $p(\mathbf{x})$  by convolving it with a spherical Gaussian  $\mathcal{N}(0, \sigma^2 I_n)$ . This smoothing relies on the fact that  $e_k \in \mathcal{D}$  represent scalar values on the real line, and therefore the discrete distribution  $p(\mathbf{x}_i | \mathbf{x}_{<i})$  can be viewed as a linear combination of weighted Dirac spikes on  $\mathcal{D} \subset \mathbb{R}$ . If  $\phi_\sigma(\mathbf{x})$  denotes the density of a spherical Gaussian  $\mathcal{N}(0, \sigma^2 I_n)$  on  $\mathbb{R}^n$  then we define a density  $p_\sigma(\tilde{\mathbf{x}})$  on  $\mathbb{R}^n$  given by

$$p_\sigma(\tilde{\mathbf{x}}) = (\phi_\sigma * p)(\tilde{\mathbf{x}}) = \int \phi_\sigma(\tilde{\mathbf{x}} - \mathbf{x}) p(\mathbf{x}) d\mathbf{x}. \quad (5)$$

This distribution has well-defined gradients  $\nabla_{\tilde{\mathbf{x}}} \log p_\sigma(\tilde{\mathbf{x}})$  and  $p_\sigma(\tilde{\mathbf{x}}) \rightarrow p(\mathbf{x})$  in total variation as  $\sigma^2 \rightarrow 0$ .

When  $p(\mathbf{x})$  is a deep autoregressive model, the convolution (5) is difficult to calculate directly. Previous work proposed training a smoothed model  $p_\sigma(\tilde{\mathbf{x}})$  by fine-tuning a model  $p(\mathbf{x})$  on noisy data  $\tilde{\mathbf{x}} = \mathbf{x} + \varepsilon_\sigma$  where  $\mathbf{x} \sim p$  and  $\varepsilon_\sigma \sim \mathcal{N}(0, \sigma^2 I_n)$  (Jayaram & Thickstun, 2020). This approach cannot be directly applied to discrete autoregressive models as defined by Equation (4). The obstruction is that noisy samples  $\tilde{\mathbf{x}}_i \in \mathbb{R}$  are not supported by the discretization  $\mathcal{D}$ . One way to address the problem is to replace the discrete model  $p(\mathbf{x})$  with a continuous autoregressive model of  $p_\sigma(\tilde{\mathbf{x}})$ , e.g. RNADE (Uribe et al., 2013). We avoid this approach because fine-tuning  $p(\mathbf{x})$  to  $p_\sigma(\tilde{\mathbf{x}})$  becomes complicated when these models have different architectures.

Instead of directly fine-tuning  $p(\mathbf{x})$  to a model  $p_\sigma(\tilde{\mathbf{x}})$ , we combine an analytic calculation with an auxiliary model

learned via fine-tuning. Let  $\tilde{p}_\sigma(\mathbf{x}_i | \tilde{\mathbf{x}}_{<i})$  denote a (discrete) conditional model trained to predict  $\mathbf{x}_i$  given noisy covariates  $\tilde{\mathbf{x}}_{<i} = \mathbf{x}_{<i} + \varepsilon_{\sigma, <i}$ . If  $\varphi_\sigma$  denotes the density of  $\mathcal{N}(0, \sigma^2)$  then we can re-write the factored density  $p_\sigma(\tilde{\mathbf{x}})$  as

$$p_\sigma(\tilde{\mathbf{x}}) = \prod_{i=1}^n p_\sigma(\tilde{\mathbf{x}}_i | \tilde{\mathbf{x}}_{<i}) = \prod_{i=1}^n (\varphi_\sigma * \tilde{p}_\sigma(\cdot | \tilde{\mathbf{x}}_{<i}))(\tilde{\mathbf{x}}_i). \quad (6)$$

On the right-hand side, we decompose the smoothed conditional densities  $p_\sigma(\tilde{\mathbf{x}}_i | \tilde{\mathbf{x}}_{<i})$  into Gaussian convolutions of discrete conditionals  $\tilde{p}_\sigma(\cdot | \tilde{\mathbf{x}}_{<i})$  evaluated at  $\tilde{\mathbf{x}}_i$ . This suggests the following approach to evaluating  $p_\sigma(\tilde{\mathbf{x}} | \tilde{\mathbf{x}}_{<i})$ :

- Learn a (discrete) model  $\tilde{p}_\sigma(\mathbf{x}_i | \tilde{\mathbf{x}}_{<i})$ , trained to predict the un-noised value  $\mathbf{x}_i$  given noisy history  $\tilde{\mathbf{x}}_{<i}$ . This model can be learned efficiently by finetuning a pre-trained model  $p(\mathbf{x}_i | \mathbf{x}_{<i})$  on noisy covariates  $\tilde{\mathbf{x}}_{<i}$ . This is visualized in the middle column of Figure 1.
- Evaluate the Gaussian convolution  $\varphi_\sigma * \tilde{p}_\sigma(\cdot | \tilde{\mathbf{x}}_{<i})$  at  $\tilde{\mathbf{x}}_i$  to compute  $p_\sigma(\tilde{\mathbf{x}}_i | \tilde{\mathbf{x}}_{<i})$ . This convolution can be calculated in closed form given  $\tilde{p}_\sigma(\mathbf{x}_i | \tilde{\mathbf{x}}_{<i})$ . This is visualized in the right column of Figure 1.

The convolution  $p_\sigma(\tilde{\mathbf{x}}_i | \tilde{\mathbf{x}}_{<i}) = (\varphi_\sigma * \tilde{p}_\sigma(\cdot | \tilde{\mathbf{x}}_{<i}))(\tilde{\mathbf{x}}_i)$  has a simple closed form given by a Gaussian mixture model

$$p_\sigma(\tilde{\mathbf{x}}_i | \tilde{\mathbf{x}}_{<i}) = \sum_{k=1}^d \tilde{p}_\sigma(e_k | \tilde{\mathbf{x}}_{<i}) \varphi_\sigma(\tilde{\mathbf{x}}_i - e_k). \quad (7)$$

Using the softmax parameterization of  $\tilde{p}(\mathbf{x}_i | \mathbf{x}_{<i})$  given by Equation (4), with fine-tuned logits  $f_{\sigma,i} : \mathcal{X}^{\otimes i} \rightarrow \mathbb{R}^d$ , the log-density of this smoothed conditional density can be written in a numerically stable form:

$$\begin{aligned} \log p_\sigma(\mathbf{x}_i | \mathbf{x}_{<i}) &= -\log \sum_{\ell=1}^d \exp(f_{\sigma,i,\ell}(\mathbf{x}_{<i})) \\ &+ \log \sum_{k=1}^d \exp\left(f_{\sigma,i,k}(\mathbf{x}_{<i}) - \frac{1}{2\sigma^2}(\mathbf{x}_i - e_k)^2\right) + C. \end{aligned} \quad (8)$$

This smoothing procedure requires us to store finetuned models  $\tilde{p}_\sigma(\mathbf{x}_i | \tilde{\mathbf{x}}_{<i})$  for each of the  $L$  noise levels  $\sigma \in \{\sigma_1, \dots, \sigma_L\}$ . This could be avoided by training a single noise-conditioned generative model. The advantage of fine-tuning is that we can directly use standard models, without any adjustment to the network architecture or subsequent hyper-parameter tuning of the modified architecture; this cleanly decouples our approach to conditional sampling from neural architecture design questions. Note that while we store  $L$  copies of the model, there is no additional memory overhead: these models are loaded and unloaded serially during optimization as we anneal the noise levels, so only one model is resident in memory at a time. While memory is a scarce resource, disk space is generally abundant.

**Algorithm 2** Stochastic Parallel and Flexible Sampling

```

Input:  $\mathbf{y}$ ,  $\{\sigma_i\}_{i=1}^L$ ,  $\delta$ ,  $T$ 
Sample  $\mathbf{x} \sim \mathcal{N}(0, \sigma_1^2 I_n)$ 
for  $i \leftarrow 1$  to  $L$  do
     $\eta_i \leftarrow \delta \cdot \sigma_i^2 / \sigma_L^2$ 
    Fork()
    for  $t = 1$  to  $T$  do
        Sample  $j \sim \text{Uniform}\{1, \dots, n\}$ 
        Sample  $\varepsilon_t \sim \mathcal{N}(0, I_c)$ 
        Read  $\mathbf{x}_{j-w:j+c+w}^{(t)} \leftarrow \mathbf{x}_{j-w:j+c+w}$ 
         $\mathbf{g}_j^{(t)} \leftarrow \nabla_{\mathbf{x}_{j:j+c}} \log p_{\sigma_i}(\mathbf{x}_{j-w:j+c+w}^{(t)})$ 
             $+ \nabla_{\mathbf{x}_{j:j+c}} \log p_{\sigma_i}(\mathbf{y}_{j:j+c} | \mathbf{x}_{j-w:j+c+w}^{(t)})$ 
        Write  $\mathbf{x}_{j:j+c} \leftarrow \mathbf{x}_{j:j+c}^{(t)} + \eta_i \mathbf{g}_j^{(t)} + \sqrt{2\eta_i} \varepsilon_t$ 
    end for
    Synchronize()
end for
    
```

### 3.2. Stochastic Gradient Langevin Dynamics

Calculating the Langevin updates described in Equations (2) and (3) requires  $O(n)$  operations to compute  $\log p_\sigma(\mathbf{x})$ , given a sequence  $\mathbf{x} \in \mathbb{R}^n$ . This calculation decomposes into  $n$  calculations of  $\log p_\sigma(\mathbf{x}_i | \mathbf{x}_{<i})$  ( $i = 1, \dots, n$ ) which can each be computed in parallel (this is what allows autoregressive models to be efficiently trained using the maximum likelihood objective). But for very large values of  $n$ , even modern parallel computing devices cannot fully parallelize all  $n$  calculations. In this section, we develop a stochastic variant of PnF sampling (Algorithm 2) that is easily distributed across multiple devices.

Instead of making batch updates on a full sequence  $\mathbf{x} \in \mathbb{R}^n$ , consider updating a single coordinate  $j \in \{1, \dots, n\}$ :

$$\mathbf{x}_j^{(t+1)} = \mathbf{x}_j^{(t)} + \eta \nabla_{\mathbf{x}_j} \log p_\sigma(\mathbf{x}^{(t)}) + \sqrt{2\eta} \varepsilon_j^{(t)}. \quad (9)$$

This coordinate-wise derivative is only dependent on the tail of the sequence  $\mathbf{x}_{\geq j}$ :

$$\nabla_{\mathbf{x}_j} \log p_\sigma(\mathbf{x}) = \sum_{i=j}^n \nabla_{\mathbf{x}_j} \log p_\sigma(\mathbf{x}_i | \mathbf{x}_{<i}). \quad (10)$$

This doesn't look promising; calculating an update on a single coordinate  $\mathbf{x}_j$  required  $n - j$  inference calculations  $p_\sigma(\mathbf{x}_i | \mathbf{x}_{<i})$ . But models over very long sequences, including WaveNets, usually make a Markov assumption  $p(\mathbf{x}_i | \mathbf{x}_{<i}) = p(\mathbf{x}_i | \mathbf{x}_{i-w}, \dots, \mathbf{x}_{i-1})$  for some limited contextual window of length  $w$ . In this case, the coordinate-wise derivative requires only  $w$  calls:

$$\nabla_{\mathbf{x}_j} \log p_\sigma(\mathbf{x}) = \sum_{i=j}^{j+w} \nabla_{\mathbf{x}_j} \log p_\sigma(\mathbf{x}_i | \mathbf{x}_{<i}). \quad (11)$$

Calculating a gradient on a contiguous block of  $c$  coordi-

nates leads to a more efficient update

$$\nabla_{\mathbf{x}_{j:j+c}} \log p_\sigma(\mathbf{x}) = \sum_{i=j}^{j+c+w} \nabla_{\mathbf{x}_{j:j+c}} \log p_\sigma(\mathbf{x}_i | \mathbf{x}_{<i}). \quad (12)$$

Calculating Equation (12) requires transmission of a block  $\{\mathbf{x}_{j-w}, \dots, \mathbf{x}_{j+c+w}\}$  of length  $c + 2w$  to the computing device, and  $c + w$  calculations  $p_\sigma(\mathbf{x}_i | \mathbf{x}_{<i})$  in order to compute the gradient of a block of length  $c$ . If we partition a sequence of length  $n$  into  $n/c$  blocks of length  $c$ , then we can distribute computation of  $\nabla_{\mathbf{x}} \log p_\sigma(\mathbf{x})$  with an overhead factor of  $1 + w/c$ . This motivates choosing  $c$  as large as possible, under the constraint that  $c + w$  calculations  $p_\sigma(\mathbf{x}_i | \mathbf{x}_{<i})$  can still be parallelized on a single device.

We can calculate  $\nabla_{\mathbf{x}} \log p_\sigma(\mathbf{x})$  by aggregating  $n/c$  blocks of gradients according to Equation (12), requiring synchronous communication between  $n/c$  machines for every update Equation (2); this is a MapReduce algorithm (Dean & Ghemawat, 2004). We propose a bolder approach in Algorithm 2 based on block-stochastic Langevin dynamics (Welling & Teh, 2011). If  $j \in \{1, \dots, n\}$  is chosen uniformly at random then Equation (12) is an unbiased estimate of  $\nabla_{\mathbf{x}} \log p_\sigma(\mathbf{x})$ . This motivates block-stochastic updates on patches, which multiple devices can perform asynchronously, a Langevin analog to Hogwild! (Niu et al., 2011).

### 3.3. Smoothing a Joint Distribution

We now consider joint distributions  $p(\mathbf{x}, \mathbf{y}) = p(\mathbf{y} | \mathbf{x}) p(\mathbf{x})$  over sources  $\mathbf{x} \in \mathcal{X}^n$  and measurements  $\mathbf{y} \in \mathcal{Y}^n$ . For example,  $\mathbf{y}$  could be a low resolution version of the signal, or an observed mixture of two signals. We are particularly interested in measurement models of the form  $\mathbf{y} = g(\mathbf{x})$ , for some linear function  $g : \mathcal{X}^n \rightarrow \mathcal{Y}^n$ . We can view these measurements  $\mathbf{y}$  as degenerate likelihoods  $p(\mathbf{y} | \mathbf{x})$  of the form

$$p(\mathbf{y} | \mathbf{x}) = \delta(\mathbf{y} - g(\mathbf{x})), \quad (13)$$

where  $\delta$  denotes the Dirac delta function. This family of linear measurement models describes the source separation, in-painting, and super-resolution tasks featured in Section 4, as well as other linear inverse problems including sparse recovery and image colorization.

Extending our analysis in Section 3.1, we can smooth the joint density  $p(\mathbf{x}, \mathbf{y})$  by convolving  $\mathbf{x}$  with a spherical Gaussian  $\mathcal{N}(0, \sigma^2 I_n)$ . Let  $\tilde{\mathbf{x}} = \mathbf{x} + \varepsilon_\sigma$  where  $\varepsilon_\sigma \sim \mathcal{N}(0, \sigma^2 I_n)$ . Note that  $\tilde{\mathbf{x}}$  is conditionally independent of  $\mathbf{y}$  given  $\mathbf{x}$  and therefore the joint distribution over  $\mathbf{x}$ ,  $\tilde{\mathbf{x}}$ , and  $\mathbf{y}$  can be factored as

$$p_\sigma(\mathbf{x}, \mathbf{y}, \tilde{\mathbf{x}}) = p_\sigma(\tilde{\mathbf{x}} | \mathbf{x}) p(\mathbf{y} | \mathbf{x}) p(\mathbf{x}). \quad (14)$$

We will work with the smoothed marginal  $p_\sigma(\tilde{\mathbf{x}}, \mathbf{y})$  of the joint distribution  $p_\sigma(\mathbf{x}, \mathbf{y}, \tilde{\mathbf{x}})$ . If  $\phi_\sigma(\mathbf{x})$  denotes the density of

a spherical Gaussian  $\mathcal{N}(0, \sigma^2 I_n)$  on  $\mathbb{R}^n$  then the marginal  $p_\sigma(\tilde{\mathbf{x}}, \mathbf{y})$  of Equation (14) can be expressed by

$$p_\sigma(\tilde{\mathbf{x}}, \mathbf{y}) = \int \phi_\sigma(\tilde{\mathbf{x}} - \mathbf{x}) p(\mathbf{x}, \mathbf{y}) d\mathbf{x} \quad (15)$$

This density approximates the original distribution  $p(\mathbf{x}, \mathbf{y})$  in the sense that  $p_\sigma(\tilde{\mathbf{x}}, \mathbf{y}) \rightarrow p(\mathbf{x}, \mathbf{y})$  in total variation as  $\sigma^2 \rightarrow 0$ .

The smoothed density  $p_\sigma(\tilde{\mathbf{x}}, \mathbf{y})$  can be factored as  $p_\sigma(\mathbf{y}|\tilde{\mathbf{x}})p_\sigma(\tilde{\mathbf{x}})$ . The density  $p_\sigma(\tilde{\mathbf{x}})$  is simply the smoothed Gaussian convolution of  $p(\mathbf{x})$  given by Equation (5). For general likelihoods  $p(\mathbf{y}|\mathbf{x})$ , because  $\mathbf{y}$  is conditionally independent of  $\tilde{\mathbf{x}}$  given  $\mathbf{x}$ , we can write the density  $p_\sigma(\mathbf{y}|\tilde{\mathbf{x}})$  by marginalizing over  $\mathbf{x}$  as

$$p_\sigma(\mathbf{y}|\tilde{\mathbf{x}}) = \int p_\sigma(\mathbf{y}|\tilde{\mathbf{x}}, \mathbf{x}) p_\sigma(\mathbf{x}|\tilde{\mathbf{x}}) d\mathbf{x} \quad (16)$$

$$= \int p(\mathbf{y}|\mathbf{x}) \phi_\sigma(\mathbf{x} - \tilde{\mathbf{x}}) d\mathbf{x}. \quad (17)$$

This integral is difficult to calculate directly. One way to evaluate  $p_\sigma(\mathbf{y}|\tilde{\mathbf{x}})$  is to take the same approach described in Section 3.1 to evaluate  $p_\sigma(\tilde{\mathbf{x}})$ : fine-tune the model  $p(\mathbf{y}|\mathbf{x})$  to a model  $\tilde{p}(\mathbf{y}|\tilde{\mathbf{x}})$  that predicts  $\mathbf{y}$  given noisy covariates  $\tilde{\mathbf{x}} = \mathbf{x} + \varepsilon_\sigma$ . A similar procedure (using a noise-conditioned architecture rather than finetuning) is described in Section 5 of Song et al. (2021).

For linear measurement models with the form given by Equation (13), the smoothed density  $p_\sigma(\mathbf{y}|\tilde{\mathbf{x}})$  can be calculated in closed form. Writing the linear function  $g : \mathcal{X}^n \rightarrow \mathcal{Y}^n$  as a matrix  $g(\mathbf{x}) = A\mathbf{x}$ , we have

$$\mathbf{y} = g(\mathbf{x}) = g(\tilde{\mathbf{x}}) + g(-\varepsilon_\sigma) \sim \mathcal{N}(g(\tilde{\mathbf{x}}), \sigma^2 AA^T). \quad (18)$$

This smoothing given by Equation (18) generalizes the smoothing proposed in Jayaram & Thickstun (2020) for source separation. That work proposed separately smoothing the prior and likelihood, resulting in a smoothed likelihood  $p(\tilde{\mathbf{y}}|\mathbf{x})$  over  $\tilde{\mathbf{y}} = \mathbf{y} + \varepsilon_y$ , where  $\varepsilon_y \sim \mathcal{N}(0, \sigma^2 I_n)$ . This is equivalent to Equation (18) in the case of source separation, for which  $g(\mathbf{x}) = \frac{1}{2}\mathbf{x}_1 + \frac{1}{2}\mathbf{x}_2$  and therefore  $g(\mathbf{x}) \sim \mathcal{N}(g(\tilde{\mathbf{x}}), \sigma^2 I)$ .

For general likelihoods  $p(\mathbf{y}|\mathbf{x})$  (e.g. a classifier) the conditioning values  $\mathbf{y}$  may depend on the whole sequence  $\mathbf{x}$ . In this case, stochastic PnF must read the entire sequence  $\mathbf{x}$  in order to calculate the posterior

$$\nabla_{\mathbf{x}_{j:j+c}} \log p(\mathbf{x}|\mathbf{y}) = \nabla_{\mathbf{x}_{j:j+c}} (\log p(\mathbf{x}) + \log p(\mathbf{y}|\mathbf{x})). \quad (19)$$

But for long sequences  $\mathbf{x}$  such as audio, the conditioning information  $\mathbf{y}$  is often a local function of the sequence  $\mathbf{x}$ . In this case,  $\mathbf{x} \in \mathcal{X}^n$ ,  $\mathbf{y} \in \mathcal{Y}^n$ ,  $\mathbf{y}_i = g(\mathbf{x}_{N(i)})$  where  $N(i)$  is

a local neighborhood of indices near  $i$ , and the likelihood decomposes via conditional independence into

$$\log p(\mathbf{y}|\mathbf{x}) = \sum_{i=1}^n \log p(\mathbf{y}_i|\mathbf{x}_{N(i)}). \quad (20)$$

All experiments presented in Section 4 feature this conditioning pattern. For spectrogram conditioning,  $N(i)$  is the set of indices (centered at  $i$ ) required to compute a short-time Fourier transform. For source separation, super-resolution, and inpainting,  $N(i) = i$ . This allows us to compute block gradients of the conditional likelihood (Algorithm 2).

## 4. Experiments

We present qualitative and quantitative results of PnF sampling for WaveNet models of audio (van den Oord et al., 2016a) and a PixelCNN++ model of images (Salimans et al., 2017). In Section 4.2 we show that PnF sampling can produce samples of comparable quality to ancestral sampling. In Section 4.3 we show that stochastic PnF sampling is faster than ancestral sampling, when parallelized across a modest number of devices. We go on to demonstrate how PnF sampling can be applied to a variety of image and audio restoration tasks: source separation (Section 4.4), super-resolution (Section 4.5), and inpainting (Section 4.6). We encourage the reader to browse the supplementary material for qualitative examples of PnF audio sampling.

### 4.1. Datasets

For audio experiments we use the VCTK dataset (Veaux et al., 2016) consisting of 44 hours of speech, as well as the Supra Piano dataset (Shi et al., 2019) consisting of 52 hours of piano recordings. We use a random 80-20 train-test split of VCTK speakers and piano recordings for evaluation. Audio sequences are sampled at a 22kHz, with 8-bit  $\mu$ -law encoding (CCITT, 1988), except for source separation where 8-bit linear encoding is used. Sequences used for quantitative evaluation are 50k sample excerpts, approximately 2.3 seconds of audio, chosen randomly from the longer test set recordings. For image experiments we use the CIFAR-10 dataset (Krizhevsky, 2009) with the standard train-test split. Additional training and hyperparameter details can be found in the appendix.

### 4.2. Quality of Generated Samples

To evaluate the quality of samples generated by PnF sampling, we follow a similar procedure to Holtzman et al. (2020). We compare log-likelihoods, calculated using the noiseless model  $p(\mathbf{x})$ , of sequences generated by PnF sampling to sequences generated by ancestral sampling from the lowest-noise model  $p_{\sigma_L}(\mathbf{x})$ . Because PnF-sampled sequences are continuous, we quantize these samples to 8-bit

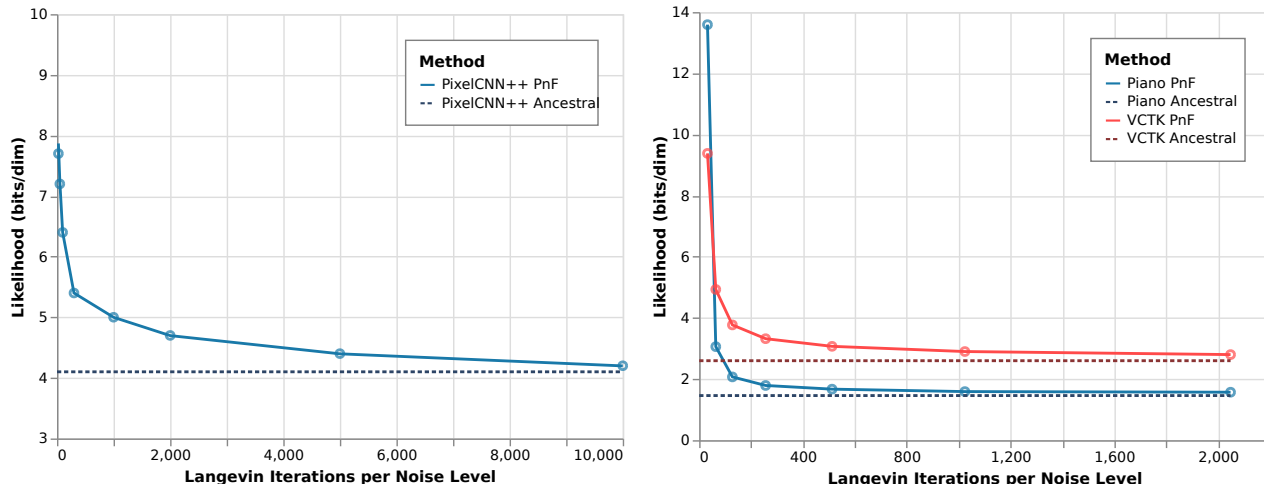


Figure 2. As the number of Langevin iterations  $T$  increases, the log-likelihood of sequences generated by PnF sampling approaches the log-likelihood of test set sequences. Left: sampling from a PixelCNN++ model trained on CIFAR-10. Right: sampling from WaveNet models trained on the Supra Piano and VCTK speech datasets.

values when evaluating their likelihood under the noiseless model. We consider PnF sampling to be successful if it generates sequences with comparable log-likelihoods to ancestral generations.

In Figure 2 we present quantitative results for PnF sampling using an unconditional PixelCNN++ model of CIFAR-10, and a spectrogram-conditioned WaveNet model of both voice and piano datasets. We evaluate 1,000 generations (length  $n = 50,000$  sequences for the WaveNet models) using various numbers of Langevin iterations, and report the median log likelihood of quantizations of these sequences under the noiseless model. Asymptotically, as the iterations  $T$  of Langevin dynamics increase, the likelihood of PnF samples approaches the likelihood of ancestral samples. Audio PnF samples for various  $T$  are presented on the project website.

### 4.3. Speed and Parallelism

Ancestral sampling has  $O(n)$  serial runtime in the length  $n$  of the generated sequence. Using the stochastic PnF sampler described in Section 3.2, the serial runtime is  $O(T)$ , where  $T$  is the number of Langevin iterations at each level of smoothing. We find empirically that we can set  $T$  independent of  $n$ , so in principle the serial runtime of stochastic PnF is constant as a function of sequence length. In practice, we do not have an infinite supply of parallel devices, so the serial runtime of stochastic PnF grows inversely proportional to the number of devices. This behavior is demonstrated in Figure 3 for spectrogram-conditioned WaveNet stochastic PnF sampling using a cluster of 8 Nvidia Titan Xp GPU’s and  $T = 256$ . Each GPU can calculate Equation (12) for a block of  $c = 50,000$  samples (2.3 seconds of audio). For PixelCNN++, we find that  $T > n$  and therefore the PnF

sampler does not improve sampling speed for this model.

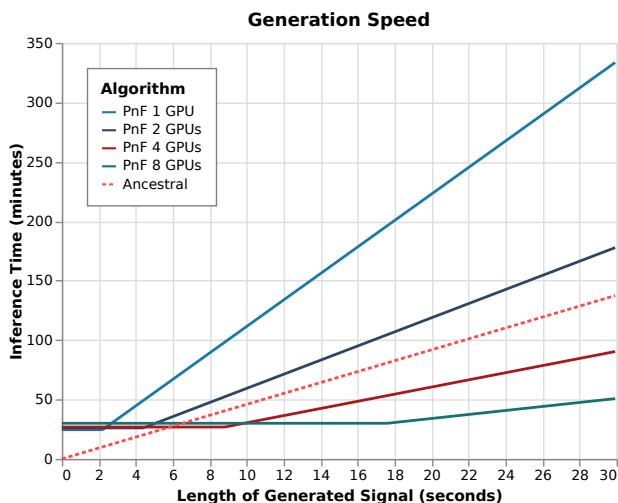


Figure 3. PnF sampling can be parallelized across multiple devices, resulting in faster inference time than ancestral sampling. Beyond a threshold level of computation, PnF sampling time is inversely proportional to the number of devices.

Stochastic PnF sampling depends upon asynchronous writes being sparse so that memory overwrites, when two workers update overlapping blocks, are rare. This situation is analogous to the sparse update condition required for Hogwild! If blocks are length  $c$  and the number of devices is substantially less than  $n/c$ , then updates are sufficiently sparse. But if the number of devices is larger than  $n/c$ , memory overwrites become common, and stochastic PnF sampling fails to converge. This imposes a floor on generation time determined by  $c$ , exhibited in Figure 3. We cannot substantially reduce this floor by decreasing  $c$  because of the tradeoff

between  $c$  and the model’s Markov window  $w$  described in Section 3.2.

In general, PnF sampling becomes faster than ancestral sampling for long sequences  $n$ , in which case the stochastic variant of PnF sampling becomes necessary in order to distribute the calculation of  $n$  conditional likelihoods. For shorter sequences, accurate unconditional samples can be produced more quickly with the ancestral sampler. Unconditional  $32 \times 32$  CIFAR-10 generation using the PixelCNN++ mode requires  $T = 10,000$  Langevin iterations per noise level (Figure 2) for accurate samples; annealing through  $L = 20$  levels requires a total of  $L \times T = 200,000$  serial queries to the PixelCNN++ model, far more than  $n = 3 \times 32 \times 32$  serial queries to PixelCNN++ for ancestral sampling. Note also that PixelCNN++ conditions on a full image ( $w = n$ ) so the stochastic variant of PnF sampling is not applicable to this model.

#### 4.4. Source Separation

The single-channel source separation problem (Davies & James, 2007) asks us to recover unobserved sources  $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}^n = \mathbb{R}^{2 \times n}$  given an observed mixture  $\mathbf{y} = g(\mathbf{x}) = \mathbf{x}_1 + \mathbf{x}_2$ . Like Jayaram & Thickstun (2020), we view source separation as a linear Bayesian inverse problem: recover  $(\mathbf{x}_1, \mathbf{x}_2)$  given  $\mathbf{y}$  and a prior  $p(\mathbf{x}_1, \mathbf{x}_2) = p(\mathbf{x}_1)p(\mathbf{x}_2)$ . We consider three variants of this task: (1) audio separation of mixtures of voice (VCTK test set) and piano (Supra Piano test set), (2) visual separation of mixtures of CIFAR-10 test set “animal” images and “machine” images, and (3) class-agnostic visual separation of mixtures of CIFAR-10 test set images.

We compare PnF audio separation to results using the Demucs (Défossez et al., 2019) and Conv-Tasnet (Luo & Mesgarani, 2019) source separation models. Both Demucs and Conv-Tasnet are supervised models, trained specifically for the source separation task, that learn to output source components given an input mixture. An advantage of PnF sampling is that it does not rely on pairs of source signals and mixes like these supervised methods. We train the supervised models on 10K mixtures of VCTK and Supra Piano samples and measure results on 1K test set mixtures using the standard Scale Invariant Signal-to-Distortion Ratio (SI-SDR) metric for audio source separation (Le Roux et al., 2019). Results in Table 1 show that PnF sampling is competitive with these specialized source separation models. Qualitative comparisons are provided in the supplement. We do not compare results on the popular MusDB dataset (Rafii et al., 2017) because this dataset has insufficient single-channel audio to train WaveNet generative models.

For CIFAR-10, we follow the experimental methodology described in Jayaram & Thickstun (2020). Table 2 shows that PixelCNN++ performs comparably to Glow as a prior,

but underperforms NCSN. This makes sense, given the relative strength of NCSN as a prior over CIFAR-10 images in comparison to PixelCNN++ and Glow. Given the strong correlation between the quality of a generative model and the quality of separations using that model as a prior, we anticipate that more recent innovations in autoregressive image models based on transformers (Parmar et al., 2018; Child et al., 2019) will lead to stronger separation results once implementations of these models that match the results reported in these papers become public. Select qualitative image separation results are presented in Figure 4.

Table 1. Quantitative results for audio source separation of mixtures of Supra piano and VCTK voice samples. Results are measured using SI-SDR (higher is better).

Algorithm	Test SI-SDR (dB)		
	All	Piano	Voice
PnF (WaveNet)	17.07	13.92	20.25
Conv-Tasnet	17.48	20.02	15.50
Demucs	14.18	16.67	12.75

Table 2. Quantitative results for visual sources separation on CIFAR-10. Results are measured using Inception Score / FID Score of 25,000 separations (50,000 separated images) of two overlapping CIFAR-10 images. In Class Split one image comes from the category of animals and other from the category of machines. NES results (Halperin et al., 2019) and BASIS results are as reported in Jayaram & Thickstun (2020).

Algorithm	Inception Score	FID
Class Split		
NES	$5.29 \pm 0.08$	51.39
BASIS (Glow)	$5.74 \pm 0.05$	40.21
<b>PnF (PixelCNN++)</b>	$5.86 \pm 0.07$	40.66
Average	$6.14 \pm 0.11$	39.49
BASIS (NCSN)	$7.83 \pm 0.15$	29.92
Class Agnostic		
BASIS (Glow)	$6.10 \pm 0.07$	37.09
<b>PnF (PixelCNN++)</b>	$6.14 \pm 0.15$	37.89
Average	$7.18 \pm 0.08$	28.02
BASIS (NCSN)	$8.29 \pm 0.16$	22.12

#### 4.5. Super-Resolution

The super-resolution problem asks us to recover unobserved data  $\mathbf{x}$  given a down-sampled observation  $\mathbf{y} = g(\mathbf{x})$ . For 1-dimensional (audio) super-resolution,  $\mathbf{x} \in \mathbb{R}^n$ ,  $\mathbf{y} \in \mathbb{R}^{n/r}$ , and  $\mathbf{y}_i = \mathbf{x}_{r_i}$  (Kuleshov et al., 2017; Eskimez et al., 2019). For 2-dimensional (image) super-resolution  $\mathbf{x} \in \mathbb{R}^n = \mathbb{R}^{k \times k \times 3}$ ,  $\mathbf{y} \in \mathbb{R}^{k/r \times k/r \times 3}$ , and  $\mathbf{y}_{i,j} = \mathbf{x}_{r_i, r_j}$  (Dahl et al., 2017; Zhang et al., 2018). Like source separation, super-resolution can be viewed as a Bayesian linear inverse prob-



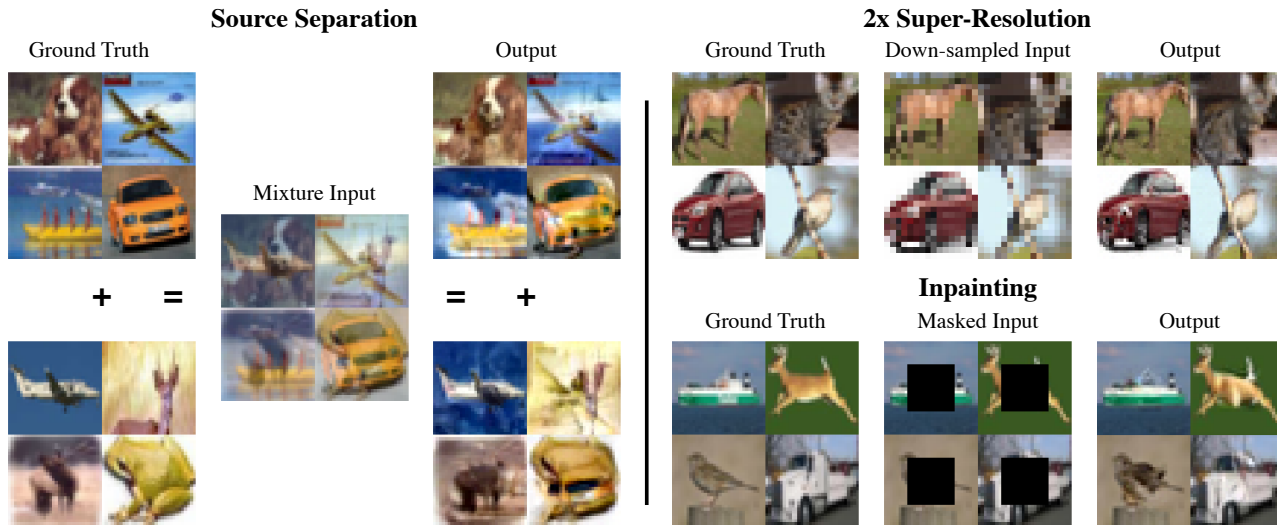


Figure 4. PnF sampling applied to visual source separation (Section 4.4) super-resolution (Section 4.5) and inpainting (Section 4.6) using a PixelCNN++ prior over images trained on CIFAR-10. Ground-truth images in this figure are taken from the CIFAR-10 test set.

Table 3. Quantitative results for audio super-resolution at three different scales on the Supra piano and VCTK voice datasets. Results are measured using PSNR (higher is better). KEE refers to the method described in Kuleshov et al. (2017)

Ratio	Piano			Voice		
	Spline	KEE	PnF	Spline	KEE	PnF
4x	23.07	22.25	29.78	15.8	16.04	15.47
8x	13.58	15.79	23.49	10.7	11.15	10.03
16x	7.09	6.76	14.23	6.4	7.11	5.32

lem, and we can recover solutions to this problem via PnF sampling. In the audio domain, the down-sampling operation  $g(\mathbf{x})$  can be interpreted as a low-pass filter.

We measure audio super-resolution performance using peak signal-to-noise ratio (PSNR) and compare against a deep learning baseline (Kuleshov et al., 2017) as well as a simple cubic B-spline. Quantitative audio results are presented in Table 3, which show that we outperform these baselines on piano data and produce similar quality reconstructions on voice data. Qualitative audio samples are available in the supplement, where we also show examples of 32x super resolution—beyond the reported ability of existing methods. Select qualitative visual results are presented in Figure 4.

#### 4.6. Inpainting

Inpainting problems involve the recovery of unobserved data  $\mathbf{x}$  given a masked observation  $g(\mathbf{x}) = \mathbf{m} \odot \mathbf{x}$ , where  $\mathbf{m} \in \{0, 1\}^n$  (Adler et al., 2011; Pathak et al., 2016). This family of problems includes completion tasks (finishing a sequence given a prime) pre-completion tasks (generating a prefix to

a sequence) and outpainting tasks. Ancestral sampling can only be applied to completion tasks, whereas PnF sampling can be used to fill in any pattern of masked occlusions. Qualitative results for audio inpainting are available in the supplement. Select qualitative results for image inpainting are presented in Figure 4.

## 5. Conclusion

In this paper we introduced PnF sampling, a parallelizable approach to sampling from autoregressive models that can be flexibly adapted to conditional sampling tasks. The flexibility of PnF sampling decouples the (unconditional) generative modeling problem from the details of specific conditional sampling tasks. Using WaveNet models, we demonstrated a reduction in wall-clock sampling time using PnF sampling in comparison to ancestral sampling, as well as PnF’s ability to solve a variety of practical audio processing problems: source separation, super-resolution, and inpainting. We anticipate that PnF conditional sampling results will improve as developments in generative modeling empower us to incorporate stronger models as priors. More broadly, we are inspired by ongoing research in generative modeling that, coupled with PnF sampling, will continue to drive performance improvements for practical conditional image and audio restoration tasks.

## Acknowledgements

We thank Zaid Harchaoui, Sham M. Kakade, Steven Seitz, and Ira Kemelmacher-Shlizerman for valuable discussion and computing resources. This work was supported by a Qualcomm Innovation Fellowship.

## References

- Adler, A., Emiya, V., Jafari, M. G., Elad, M., Gribonval, R., and Plumbley, M. D. Audio inpainting. *IEEE Transactions on Audio, Speech, and Language Processing*, 2011.
- Blake, A. and Zisserman, A. *Visual Reconstruction*. MIT Press, 1987.
- Bora, A., Jalal, A., Price, E., and Dimakis, A. G. Compressed sensing using generative models. In *International Conference on Machine Learning*, 2017.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, 2020.
- CCITT. Pulse code modulation (pcm) of voice frequencies. International Telecommunication Union, 1988.
- Chen, N., Zhang, Y., Zen, H., Weiss, R. J., Norouzi, M., and Chan, W. Wavegrad: Estimating gradients for waveform generation. In *International Conference on Learning Representations*, 2021.
- Child, R., Gray, S., Radford, A., and Sutskever, I. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- Dahl, R., Norouzi, M., and Shlens, J. Pixel recursive super resolution. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017.
- Davies, M. E. and James, C. J. Source separation using single channel ica. *Signal Processing*, 2007.
- Dean, J. and Ghemawat, S. Mapreduce: Simplified data processing on large clusters. *Communications of the ACM*, 2004.
- Défossez, A., Usunier, N., Bottou, L., and Bach, F. Music source separation in the waveform domain. *arXiv preprint arXiv:1911.13254*, 2019.
- Dhariwal, P., Jun, H., Payne, C., Kim, J. W., Radford, A., and Sutskever, I. Jukebox: A generative model for music. *arXiv preprint arXiv:2005.00341*, 2020.
- Donahue, C., McAuley, J., and Puckette, M. Adversarial audio synthesis. In *International Conference on Learning Representations*, 2019.
- Du, Y. and Mordatch, I. Implicit generation and modeling with energy based models. In *Advances in Neural Information Processing Systems*, 2019.
- Eskimez, S. E., Koishida, K., and Duan, Z. Adversarial training for speech super-resolution. *IEEE Journal of Selected Topics in Signal Processing*, 2019.
- Frank, M. and Ilse, M. Problems using deep generative models for probabilistic audio source separation. In “*I Can’t Believe It’s Not Better!*” *NeurIPS 2020 workshop*, 2020.
- Hadjeres, G., Pachet, F., and Nielsen, F. Deepbach: a steerable model for bach chorales generation. In *International Conference on Machine Learning*, 2017.
- Halperin, T., Ephrat, A., and Hoshen, Y. Neural separation of observed and unobserved distributions. In *International Conference on Machine Learning*, 2019.
- Holtzman, A., Buys, J., Du, L., Forbes, M., and Choi, Y. The curious case of neural text degeneration. In *International Conference on Learning Representations*, 2020.
- Jayaram, V. and Thickstun, J. Source separation with deep generative priors. In *International Conference on Machine Learning*, 2020.
- Kim, S., Lee, S.-G., Song, J., Kim, J., and Yoon, S. Flowavenet: A generative flow for raw audio. In *International Conference on Machine Learning*, 2018.
- Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P. Optimization by simulated annealing. *Science*, 1983.
- Knapik, B. T., Van Der Vaart, A. W., van Zanten, J. H., et al. Bayesian inverse problems with gaussian priors. *The Annals of Statistics*, 2011.
- Kong, Z., Ping, W., Huang, J., Zhao, K., and Catanzaro, B. Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations*, 2021.
- Krizhevsky, A. Learning multiple layers of features from tiny images. 2009.
- Kuleshov, V., Enam, S. Z., and Ermon, S. Audio super-resolution using neural nets. In *International Conference on Learning Representations (Workshop Track)*, 2017.
- Kumar, K., Kumar, R., de Boissiere, T., Gestin, L., Teoh, W. Z., Sotelo, J., de Brébisson, A., Bengio, Y., and Courville, A. Melgan: Generative adversarial networks for conditional waveform synthesis. In *Advances in Neural Information Processing Systems*, 2019.
- Larochelle, H. and Murray, I. The neural autoregressive distribution estimator. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 2011.

- Le Roux, J., Wisdom, S., Erdogan, H., and Hershey, J. R. Sdr-half-baked or well done? In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- Luo, Y. and Mesgarani, N. Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2019.
- Mehri, S., Kumar, K., Gulrajani, I., Kumar, R., Jain, S., Sotelo, J., Courville, A., and Bengio, Y. Samplernn: An unconditional end-to-end neural audio generation model. In *International Conference on Learning Representations*, 2017.
- Neal, R. M. et al. Mcmc using hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2011.
- Niu, F., Recht, B., Ré, C., and Wright, S. J. Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems*, 2011.
- Parmar, N., Vaswani, A., Uszkoreit, J., Kaiser, Ł., Shazeer, N., Ku, A., and Tran, D. Image transformer. In *International Conference on Machine Learning*, 2018.
- 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*, 2016.
- Ping, W., Peng, K., and Chen, J. Clarinet: Parallel wave generation in end-to-end text-to-speech. In *International Conference on Learning Representations*, 2019.
- Ping, W., Peng, K., Zhao, K., and Song, Z. Waveflow: A compact flow-based model for raw audio. In *International Conference on Machine Learning*, 2020.
- Prenger, R., Valle, R., and Catanzaro, B. Waveglow: A flow-based generative network for speech synthesis. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners. *OpenAI blog*, 2019.
- Rafii, Z., Liutkus, A., Stöter, F.-R., Mimilakis, S. I., and Bittner, R. The MUSDB18 corpus for music separation, 2017. URL <https://doi.org/10.5281/zenodo.1117372>.
- Raj, A., Li, Y., and Bresler, Y. Gan-based projector for faster recovery with convergence guarantees in linear inverse problems. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019.
- Razavi, A., van den Oord, A., and Vinyals, O. Generating diverse high-fidelity images with vq-vae-2. In *Advances in Neural Information Processing Systems*, 2019.
- Rick Chang, J., Li, C.-L., Poczos, B., Vijaya Kumar, B., and Sankaranarayanan, A. C. One network to solve them all—solving linear inverse problems using deep projection models. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017.
- Salimans, T., Karpathy, A., Chen, X., and Kingma, D. P. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. In *International Conference on Learning Representations*, 2017.
- Shi, Z., Sapp, C., Arul, K., McBride, J., and Smith III, J. O. Supra: Digitizing the stanford university piano roll archive. In *International Symposium on Music Information Retrieval*, 2019.
- Song, Y. and Ermon, S. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems*, 2019.
- Song, Y. and Ermon, S. Improved techniques for training score-based generative models. *Advances in Neural Information Processing Systems*, 2020.
- Song, Y., Meng, C., Liao, R., and Ermon, S. Nonlinear equation solving: A faster alternative to feedforward computation. *arXiv preprint arXiv:2002.03629*, 2020.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- Theis, L. and Bethge, M. Generative image modeling using spatial lstms. In *Advances in Neural Information Processing Systems*, 2015.
- Tropp, J. A. and Wright, S. J. Computational methods for sparse solution of linear inverse problems. *Proceedings of the IEEE*, 2010.
- Uria, B., Murray, I., and Larochelle, H. Rnade: The real-valued neural autoregressive density-estimator. In *Advances in Neural Information Processing Systems*, 2013.
- van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., and Kavukcuoglu, K. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016a.
- van den Oord, A., Kalchbrenner, N., and Kavukcuoglu, K. Pixel recurrent neural networks. In *International Conference on Machine Learning*, 2016b.

- van den Oord, A., Vinyals, O., et al. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, 2017.
- van den Oord, A., Li, Y., Babuschkin, I., Simonyan, K., Vinyals, O., Kavukcuoglu, K., Driessche, G., Lockhart, E., Cobo, L., Stimberg, F., et al. Parallel wavenet: Fast high-fidelity speech synthesis. In *International Conference on Machine Learning*, 2018.
- Veaux, C., Yamagishi, J., MacDonald, K., et al. Superseded-cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit. 2016.
- Wang, Z., Bardsley, J. M., Solonen, A., Cui, T., and Marzouk, Y. M. Bayesian inverse problems with  $l_1$  priors: a randomize-then-optimize approach. *SIAM Journal on Scientific Computing*, 2017.
- Welling, M. and Teh, Y. W. Bayesian learning via stochastic gradient langevin dynamics. In *International Conference on Machine Learning*, 2011.
- Wiggers, A. J. and Hoogeboom, E. Predictive sampling with forecasting autoregressive models. In *International Conference on Machine Learning*, 2020.
- Xu, Y., Song, Y., Garg, S., Gong, L., Shu, R., Grover, A., and Ermon, S. Anytime sampling for autoregressive models via ordered autoencoding. In *International Conference on Learning Representations*, 2021.
- Zhang, Y., Tian, Y., Kong, Y., Zhong, B., and Fu, Y. Residual dense network for image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.