



HAL
open science

Plug-and-Play Split Gibbs Sampler: Embedding Deep Generative Priors in Bayesian Inference

Florentin Coeurdoux, Nicolas Dobigeon, Pierre Chainais

► **To cite this version:**

Florentin Coeurdoux, Nicolas Dobigeon, Pierre Chainais. Plug-and-Play Split Gibbs Sampler: Embedding Deep Generative Priors in Bayesian Inference. *IEEE Transactions on Image Processing*, 2024, 33, pp.3496-3507. 10.1109/TIP.2024.3404338 . hal-04602765

HAL Id: hal-04602765

<https://hal.science/hal-04602765>

Submitted on 5 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Plug-and-Play Split Gibbs Sampler: embedding deep generative priors in Bayesian inference

Florentin Coeurdoux, Nicolas Dobigeon, *Senior Member, IEEE* and Pierre Chainais, *Senior Member, IEEE*

Abstract—This paper introduces a stochastic plug-and-play (PnP) sampling algorithm that leverages variable splitting to efficiently sample from a posterior distribution. The algorithm based on split Gibbs sampling (SGS) draws inspiration from the half quadratic splitting method (HQS) and the alternating direction method of multipliers (ADMM). It divides the challenging task of posterior sampling into two simpler sampling problems. The first problem depends on the likelihood function, while the second is interpreted as a Bayesian denoising problem that can be readily carried out by a deep generative model. Specifically, for an illustrative purpose, the proposed method is implemented in this paper using state-of-the-art diffusion-based generative models. Akin to its deterministic PnP-based counterparts, the proposed method exhibits the great advantage of not requiring an explicit choice of the prior distribution, which is rather encoded into a pre-trained generative model. However, unlike optimization methods (e.g., PnP-ADMM and PnP-HQS) which generally provide only point estimates, the proposed approach allows conventional Bayesian estimators to be accompanied by confidence intervals at a reasonable additional computational cost. Experiments on commonly studied image processing problems illustrate the efficiency of the proposed sampling strategy. Its performance is compared to recent state-of-the-art optimization and sampling methods.

Index Terms—Bayesian inference, plug-and-play prior, deep generative model, diffusion-based model, Markov chain Monte Carlo, inverse problem.

I. INTRODUCTION

MANY scientific problems raise the challenge of inferring an unknown object of interest $\mathbf{x} \in \mathbb{R}^N$ from partial and noisy measurements $\mathbf{y} \in \mathbb{R}^M$. These inverse problems frequently encountered in image processing are typically formulated as the minimisation task

$$\min_{\mathbf{x}} f(\mathbf{x}, \mathbf{y}) + g(\mathbf{x}) \quad (1)$$

where $f(\cdot, \mathbf{y})$ denotes the data-fitting term. Due to the ill-posed or ill-conditioned nature of the inverse problem, it is often not possible to uniquely and stably recover \mathbf{x} from the sole observations \mathbf{y} . Therefore, additional information about the unknown object \mathbf{x} is incorporated in the form of the regularization $g(\cdot)$ to obtain a well-posed estimation problem,

leading to meaningful solutions [1]. Due to the increasing volume, dimensionality, and variety of available data, solving such inference problems can be computationally demanding and may rely on methods such as variational optimization or stochastic sampling.

Until recently, most of the optimization methods have relied on priors designed as explicit model-based regularizations such as the total variation, promoting piecewise constant behaviors, or the ℓ_1 norm, promoting sparsity. In this context, convex optimization algorithms have played an important role and their convergence properties have been well-established [2]–[5]. However, for an always larger family of problems related to image processing, methods based on explicit convex priors are now significantly outperformed by deep learning based approaches. There exist a number of deep neural network architectures that can directly learn a description of the solution space [6]–[10]. Such so-called end-to-end approaches that bypass the problem of explicitly defining the prior knowledge do not even need the knowledge of the forward operator itself. Instead, they are implicitly learnt from a large data set of degraded images (i.e., network input) along with their original versions (i.e., network output) when training the network. However, such end-to-end methods suffer from the lack of interpretability and generality of black-box deep neural networks (DNN). Moreover, they do not take advantage of the generally well-established expertise of the end-users about the acquisition or damaging protocols, which makes the training process particularly energy and data intensive.

To overcome these limitations, more and more deep learning based methods propose to combine DNN with conventional optimization algorithms within the so-called plug-and-play (PnP) framework [11]. The main ingredient of PnP approaches is a variable splitting strategy as implemented by half-quadratic splitting (HQS) [12] or alternating direction method of multipliers (ADMM) [13]. The main idea of this splitting consists in introducing an auxiliary variable \mathbf{z} such that the problem (1) rewrites

$$\min_{\mathbf{x}, \mathbf{z}} f(\mathbf{x}, \mathbf{y}) + g(\mathbf{z}) \quad \text{subject to} \quad \mathbf{x} = \mathbf{z}. \quad (2)$$

The equality constraint ensures that solving (2) is equivalent to solving the initial problem (1). Adopting an alternate minimization strategy, this trick permits to separately deal with the data-fitting term and the regularization [14]. In particular, the subproblem with respect to \mathbf{z} is solved by using the proximal operator of the regularization term, which can be interpreted as a denoising task. Recent PnP methods replace this proximal mapping by a DNN-based denoiser that implicitly encodes the

F. Coeurdoux and N. Dobigeon are with the University of Toulouse, IRIT/INP-ENSEEIH, CNRS 31071, Toulouse Cedex 7, France (e-mail: Florentin.Coeurdoux@irit.fr; Nicolas.Dobigeon@enseeih.fr).

P. Chainais is with the Univ. Lille, CNRS, Centrale Lille, UMR 9189—CRISTAL—Centre de Recherche en Informatique Signal et Automatique de Lille F-59000, Lille, France (e-mail: Pierre.Chainais@centralelille.fr).

This work was supported by the Artificial Natural Intelligence Toulouse Institute (ANITI, ANR-19-PI3A-0004), the AI Sherlock Chair (ANR-20-CHIA-0031-01), the ULNE national future investment programme (ANR-16-IDEX-0004) and the Hauts-de-France Region.

regularization. They now stand as a reference that yield state-of-the-art performance in a variety of applications [15], [16].

However, PnP-based optimization algorithms generally produce point estimates only. More generally, except in special cases [17], optimization methods do not give any information about the posterior distribution

$$\pi(\mathbf{x}) \triangleq p(\mathbf{x}|\mathbf{y}) \propto \exp[-f(\mathbf{x}, \mathbf{y}) - g(\mathbf{x})] \quad (3)$$

associated with (1) and do not quantify uncertainties. Conversely, Bayesian approaches and Markov chain Monte Carlo (MCMC) methods have the great advantage of providing a comprehensive description of (3) in very general settings. In particular, this knowledge permits to derive credibility intervals on the parameter \mathbf{x} of interest. This uncertainty quantification is often of crucial importance, for instance when only very few observations are available [18], when one is interested in extreme events [19] or when no ground truth is available, like in astrophysics. There is still a price to pay: sampling methods and MCMC in particular suffer from their high computational cost which can be prohibitive in high-dimensional problems. Optimization-driven Monte Carlo methods [20]–[22] tentatively overcome this limitation.

More recently, the split Gibbs sampler (SGS) [23] proposes to sample from an augmented distribution defined as an asymptotically exact data augmentation model [24]. By introducing an auxiliary variable as in (2), it yields a divide-to-conquer strategy by splitting the initial sampling problem into individual simpler sampling tasks. Sampling according to the augmented distribution with Gibbs steps permits to deal separately with the distinct components of the problem, i.e., the likelihood on the one hand and the prior on the other hand. Per se, SGS can be seen as a stochastic counterpart of HQS or ADMM algorithms. It both makes the sampling more scalable to high dimensions and significantly improves the mixing properties of the Markov chain.

The main contribution of the work reported in this paper is to provide a straightforward and systematic instantiation of the PnP paradigm within a Monte Carlo sampling framework. This is made possible thanks to the splitting strategy implemented by the SGS scheme. Moreover the timeliness of devising such an approach can be easily justified by the recent advances in the design of powerful deep generative models. The proposed approach coined as PnP-SGS is based on three main rationales. First, as any PnP-based methods, PnP-SGS allows Bayesian inference problems to be solved without explicitly defining a prior distribution, which is rather implicitly encoded into a DNN trained beforehand. Second, we show that diffusion-based or score-based models [25]–[27] initially derived for generative purposes can be diverted to be employed as universal stochastic denoisers. Third, PnP-SGS generate samples that can be used to build confidence intervals, which is not possible with its deterministic counterparts, i.e., PnP-HQS and PnP-ADMM, that only provide point estimates [15], [16]. High dimensional image processing experiments will illustrate the strong potential of the proposed approach when using a diffusion model [28]–[34] as a denoiser. These extensive experiments include various inverse problems such

as inpainting, super-resolution, and deblurring. The experimental results show that the proposed PnP-SGS is a general approach to solve ill-posed inverse problems in high dimension with superior quality and uncertainty quantification.

Section II recalls necessary notions about the split Gibbs sampler (SGS) and Denoising Diffusion Probabilistic Models (DDPMs) that will be used as PnP-denoisers in the sequel. Section III describes how the proposed PnP-SGS adapts to several usual inverse problems frequently encountered in image processing. Section IV describes numerical experiments and reports the performances in comparison with state-of-the-art methods. Section V finally enlightens the contributions.

II. SGS AND GENERATIVE MODELS FOR PNP

A. Split Gibbs sampling (SGS)

Starting from the target posterior distribution (3), the introduction of a splitting variable $\mathbf{z} \in \mathbb{R}^N$ leads to the augmented distribution

$$\begin{aligned} \pi_\rho(\mathbf{x}, \mathbf{z}) &\triangleq p(\mathbf{x}, \mathbf{z}|\mathbf{y}; \rho^2) \\ &\propto \exp\left[-f(\mathbf{x}, \mathbf{y}) - g(\mathbf{z}) - \frac{1}{2\rho^2} \|\mathbf{x} - \mathbf{z}\|_2^2\right] \end{aligned} \quad (4)$$

where ρ is a positive parameter that controls the coupling between \mathbf{x} and \mathbf{z} . As shown in [24], for a large variety of coupling kernels including the quadratic one, the marginal distribution of \mathbf{x} under π_ρ in (4) coincides with the target distribution π in (3) when ρ^2 tends to zero, i.e.,

$$\|\pi - \pi_\rho\|_{\text{TV}} \xrightarrow{\rho^2 \rightarrow 0} 0 \quad (5)$$

where $\|\pi\|_{\text{TV}} \triangleq \int_{\mathbb{R}^M} |\pi(\mathbf{x})| d\mathbf{x}$. This defines an asymptotically exact data augmentation scheme [24]. In other words, the original target distribution $\pi(\mathbf{x})$ in (3) is recovered from the marginal distribution $\pi_\rho(\mathbf{x})$ derived from (4) in the limiting case $\rho \rightarrow 0$. Instead of sampling directly according to $\pi(\mathbf{x})$, SGS proposes to sample according to the augmented distribution $\pi_\rho(\mathbf{x}, \mathbf{z})$ using Gibbs steps. More specifically, the associated conditional distributions to sample from $\pi_\rho(\mathbf{x}, \mathbf{z})$ are given by

$$p(\mathbf{x} | \mathbf{z}, \mathbf{y}; \rho^2) \propto \exp\left[-f(\mathbf{x}, \mathbf{y}) - \frac{1}{2\rho^2} \|\mathbf{x} - \mathbf{z}\|_2^2\right] \quad (6)$$

$$p(\mathbf{z} | \mathbf{x}; \rho^2) \propto \exp\left[-g(\mathbf{z}) - \frac{1}{2\rho^2} \|\mathbf{z} - \mathbf{x}\|_2^2\right] \quad (7)$$

It is now clear that sampling alternatively from (6) and (7) dissociates the potential functions $f(\cdot, \mathbf{y})$ and $g(\cdot)$ associated with the likelihood and the prior distribution, respectively. As a consequence, SGS inherits from well-known advantages already exhibited by its deterministic counterparts (i.e., HQS and ADMM), e.g., easier implementations, faster convergences and possibly distributed computations. In particular, sampling according to (6) can be interpreted as solving the initial problem defined by the same potential function $f(\cdot, \mathbf{y})$ but now granted with a Gaussian prior distribution of mean \mathbf{z} and diagonal covariance matrix $\rho^2 \mathbf{I}$. It is thus expected to be significantly simpler than sampling according to the initial posterior distribution $\pi(\mathbf{x})$ defined by (3).

Moreover, it is worth noting that the conditional distribution (7) can be interpreted as the posterior distribution associated with a Bayesian denoising problem. Its goal is to recover an object \mathbf{z} from a noisy observations \mathbf{x} contaminated by an additive white Gaussian noise with variance ρ^2 . Instead of sampling directly from (7), we propose to resort to deep generative models used as stochastic denoisers. Generative adversarial network (GAN), variational autoencoders (VAE) or more recently denoising diffusion probabilistic models (DDPM) are powerful candidates to tackle this task [25]–[27]. Due to the high interest they have received recently, this work instantiates the PnP-SGS framework and reports experimental results based on DDPM-based denoisers. Note however that any pre-trained probabilistic denoising generative model can be plugged into the proposed approach.

B. Denoising diffusion probabilistic models (DDPM)

Denoising diffusion models [28]–[31] and score based models [32]–[34] are trendy classes of generative models. They have recently drawn significant attention from the community due to their state-of-the-art performances. Although nourished by different inspirations, they share very similar aspects and can be presented as variants of each other [30], [34], [35]. They are often referred to under the generic name *diffusion models*.

1) *DDPM as generative models*: A denoising diffusion probabilistic model [28] makes use of two Markov chains: a forward chain that perturbs data to pure noise, and a backward chain that converts noise back to data. The former is typically model-based designed with the goal to transform any data distribution into a simple prior distribution, i.e., a standard Gaussian. Conversely the latter Markov chain aims at reversing the noising process by learning transition kernels parameterized by a DNN. Once the DNN has been trained, new data points can be generated by first drawing from the prior distribution, and then sampling through the backward Markov chain.

Formally, given a data distribution $\mathbf{u}_0 \sim p(\mathbf{u}_0)$, the forward Markov process generates a sequence of random variables $\mathbf{u}_t \in \mathbb{R}^N$, $t \in \{0, \dots, T\}$ according to the transition kernel $p(\mathbf{u}_t | \mathbf{u}_{t-1})$. Using the probability chain rule and the Markovian property, the joint distribution $p(\mathbf{u}_1, \dots, \mathbf{u}_T | \mathbf{u}_0)$ can be factorized as

$$p(\mathbf{u}_1, \dots, \mathbf{u}_T | \mathbf{u}_0) = \prod_{t=1}^T p(\mathbf{u}_t | \mathbf{u}_{t-1}). \quad (8)$$

In DDPMs, the transition kernel $p(\mathbf{u}_t | \mathbf{u}_{t-1})$ is arbitrarily chosen to incrementally transform the data distribution $p(\mathbf{u}_0)$ into a tractable prior distribution $p(\mathbf{u}_T) \approx \mathcal{N}(\mathbf{u}_T; \mathbf{0}, \mathbf{I})$. One typical design for the transition kernel exploits a Gaussian perturbation and the most common choice for the transition kernel is

$$p(\mathbf{u}_t | \mathbf{u}_{t-1}) = \mathcal{N}\left(\mathbf{u}_t; \sqrt{1 - \beta(t)}\mathbf{u}_{t-1}, \beta(t)\mathbf{I}\right) \quad (9)$$

where $\beta(t) \in (0, 1)$ is a predefined function which plays a key role. It directly adjusts the amount of noise along the process

such that larger values lead to noisier samples. Conventionally, it is chosen as a linearly increasing function [29]. More recent techniques have proposed to use cosine-based functions [36]. Intuitively speaking, this forward process slowly injects noise into data until all structures are lost and only noise prevails.

For generating new data samples, DDPMs start by first drawing a sample \mathbf{u}_T from an instrumental prior distribution $q(\mathbf{u}_T) = \mathcal{N}(\mathbf{u}_T; \mathbf{0}, \mathbf{I})$. Then DDPMs gradually remove noise by running a Markov chain in the reverse time direction. This Markov chain is defined thanks to a kernel modeled by DNNs. The learnable transition kernel $q_\theta(\mathbf{u}_{t-1} | \mathbf{u}_t)$ takes the form of

$$q_\theta(\mathbf{u}_{t-1} | \mathbf{u}_t) = \mathcal{N}(\mathbf{u}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{u}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{u}_t, t)) \quad (10)$$

where the mean $\boldsymbol{\mu}_\theta(\mathbf{u}_t, t)$ and the covariance matrix $\boldsymbol{\Sigma}_\theta(\mathbf{u}_t, t)$ are DNNs parametrized by θ and t with \mathbf{u}_t as an input.

2) *DDPM as stochastic denoisers*: According to the above discussion, it is clear that the forward diffusion process (9) progressively adds noise to a noise-free image \mathbf{u}_0 . Following a discretization scheme generally adopted by these deep generative models, each \mathbf{u}_t corresponds to a scaled version of \mathbf{u}_{t-1} corrupted by a Gaussian noise with covariance matrix $\beta(t)\mathbf{I}$. Thanks to the factorization induced by the direct Markov chain and the Gaussian nature of the transition kernel, the transition from the original image \mathbf{u}_0 to any intermediate noisy image \mathbf{u}_t can be written as

$$p(\mathbf{u}_t | \mathbf{u}_0) = \mathcal{N}\left(\mathbf{u}_t; \sqrt{\bar{\alpha}(t)}\mathbf{u}_0, \alpha(t)\mathbf{I}\right) \quad (11)$$

$$\text{where } \alpha(t) = \prod_{j=1}^t (1 - \beta(j)) \quad (12)$$

and $\bar{\alpha}(t) = 1 - \alpha(t)$. In other words, at any arbitrary time instant $t^* < T$, the image \mathbf{u}_{t^*} resulting from t^* steps of the forward process is a noisy version of the input image \mathbf{u}_0 corrupted by a Gaussian noise of variance $\alpha(t^*)$.

Therefore, it appears that a trained DDPM can be used as a stochastic Gaussian denoiser. Contrary to the normal use of a DDPM as a generator (see above), the key idea is rather to start the backward diffusion process from a noisy image \mathbf{u}_{t^*} for some t^* and not as usual from a realization of noise \mathbf{u}_T . The noise-free image \mathbf{u}_0 can be recovered by applying the backward process defined by (10) from time instant t^* .

C. Proposed DDPM-based PnP-SGS algorithm

In a nutshell, the proposed PnP-SGS alternatively samples according to the conditional posterior distributions (6) and (7). Along this iterative process, SGS generates a set of N_{MC} samples $\{\mathbf{x}^{(n)}, \mathbf{z}^{(n)}\}_{n=1}^{N_{MC}}$ asymptotically distributed according to the augmented posterior $\pi_\rho(\mathbf{x}, \mathbf{z})$. From this set of samples, various Bayesian quantities can be approximated, such as Bayesian estimators and credibility intervals. In particular, the samples $\{\mathbf{x}^{(n)}\}_{n=1}^{N_{MC}}$ are marginally distributed according to $\pi_\rho(\mathbf{x})$. Thus the minimum mean square estimator (MMSE or posterior mean) $\hat{\mathbf{x}}_{MMSE} = \mathbb{E}[\mathbf{x}|\mathbf{y}]$ associated with π_ρ can be easily approximated by the empirical average

$$\hat{\mathbf{x}}_{MMSE} \approx \frac{1}{N_{MC} - N_{bi}} \sum_{n=N_{bi}+1}^{N_{MC}} \mathbf{x}^{(n)} \quad (13)$$

where N_{bi} is the number of burn-in iterations.

Regarding the first step of SGS, sampling according to (6) is problem dependent and should be suitably adapted to the targeted task. For illustration purpose, it will be explicitly specified for various imaging problems in Section III. As expected and already pointed out in Section II-A, it will be shown that sampling according to (6) is significantly simpler than directly sampling according to the target posterior distribution $\pi(\mathbf{x})$ defined by (3).

Regarding the second step of SGS, at the n th iteration of the algorithm, sampling according to (7) is interpreted as a stochastic denoising of the current value $\mathbf{x}^{(n)}$. This sampling according to (7) is performed in a PnP manner thanks to a previously trained DDPM, following the strategy detailed in Section II-B2. With the notations adopted in the previous paragraph, it assigns the current sample $\mathbf{x}^{(n)}$ to the variable \mathbf{u}_{t^*} for some t^* at iteration n and then iterates the backward diffusion (10). After t^* steps, the produced denoised image \mathbf{u}_0 is allocated to the new sample $\mathbf{z}^{(n)}$ according to (7) of the current SGS iteration. Note that DDPMs used as generators are known to be generally computationally demanding due to the number T of overall steps involved in the backward process. The proposed approach obviates this impediment by initiating the process from a generally weakly noisy image, which significantly reduces the necessary number $t^* \ll T$ of denoising steps to be applied [37]. Next section provides some insights into this number t^* and proposes a systematic and reliable strategy to adjust it.

Remark. *One alternative of the proposed DDPM embedding would have been to directly train a novel surrogate probabilistic model $q_\theta(\mathbf{z}|\mathbf{x})$ for instance parametrized by its two first moments (mean and covariance). However, it is worth raising the following shortcomings: i) restricting the description of a distribution to its two first moments boils down to assuming a Gaussian distribution, which may lack of expressiveness in high dimensional problems such as those addressed in this manuscript, ii) such denoisers are known to be significantly outperformed by diffusion-based models, which replace the probabilistic model $q_\theta(\mathbf{z}|\mathbf{x})$ by successive incremental denoising steps and iii) such a strategy would require to design and train a new network while the proposed approach can reuse any pre-trained probabilistic denoising generative model.*

D. Some insights into the number t^* of backward steps

This section discusses the role and the tuning of the time instant t^* which defines the number of denoising steps to be applied at a given iteration of the SGS sampler. As already stated, Eq. (11) shows that the variance of the noise corrupting \mathbf{u}_0 after t^* transitions of the forward Markov chain is $\alpha(t^*)$. This variance is defined by the product (12) of continuous strictly monotone functions $\beta(\cdot)$, thus it is also continuous and strictly monotone. This has two consequences: *i)* a level of noise $\alpha(t^*)$ is associated to a unique instant t^* of the forward diffusion process (i.e., the noise scheduling function $t \mapsto \alpha(t)$ is invertible) and *ii)* the larger t^* , the noisier the image \mathbf{u}_{t^*} . Reciprocally, when applying the backward diffusion to a noisy

image, the larger t^* , the higher the impact of the denoising, that is of the regularization. Note that the DDPM, that is used for regularization here, has no explicit hyperparameter. An important consequence is that, within the framework of PnP-SGS, the number t^* of denoising steps can be interpreted as the hyperparameter that adjusts the amount of imposed regularization, the coupling parameter ρ being kept fixed.

The proposed approach capitalizes on the explicit and unequivocal mapping between the hyperparameter t^* and the variance $\alpha(t^*)$ of the noise contained in \mathbf{u}_{t^*} . This relationship permits a simple and efficient strategy to set the number t^* of required denoising steps (10) when sampling according to (7). Given a current sample $\mathbf{x}^{(n)}$ generated by SGS, the identification of the appropriate instant t^* to generate $\mathbf{z}^{(n)}$ according to (7) boils down to estimating the level $\alpha(t^*)$ of the noise corrupting the sample $\mathbf{x}^{(n)}$. This is possible using any good conventional estimator $\hat{\sigma} = \Phi(\mathbf{x}^{(n)})$ of the noise level in $\mathbf{x}^{(n)}$ [38]–[40], see Appendix A for implementation details. Since the noise scheduling function $t \mapsto \alpha(t)$ is invertible, one can finally set $\hat{t}^* = \alpha^{-1}(\hat{\sigma}^2)$ to start the backward diffusion (10). Appendix B discusses technical details of the inversion of the noise scheduling function $\alpha(\cdot)$.

In practice, during the experiments reported in Section IV, the number \hat{t}^* of achieved steps has been shown to stabilize at a fixed value after the burn-in period of PnP-SGS. Therefore the transition kernel associated with the denoising procedure becomes invariant, which ensures that SGS converges towards a stationary distribution π_ρ ; recall that ρ is fixed, typically of order 1, see Appendix C. The resulting distribution π_ρ is eventually similar to (4) where the role of the explicit regularizing potential $g(\cdot)$ has been implicitly replaced by the DDPM.

Algorithm 1 describes the final sampling PnP-SGS algorithm using a DDPM for the denoising step, with the proposed strategy to set the hyperparameter t^* .

Algorithm 1: PnP-SGS using DDPM

Input : Parameter ρ^2 , total number of iterations N_{MC} , number of burn-in iterations N_{bi} , pre-trained DDPM $s_\theta(\cdot, \cdot)$, noise scheduling function $\alpha(\cdot)$, initialization $\mathbf{z}^{(0)}$

1 for $n \leftarrow 1$ **to** N_{MC} **do**

2 # Sampling the variable of interest $\mathbf{x}^{(n)}$

3 Draw $\mathbf{x}^{(n)} \sim p(\mathbf{x} | \mathbf{z}, \mathbf{y}; \rho^2)$ according to (6)

4 # Estimating noise level in $\mathbf{x}^{(n)}$

5 Set $\hat{\sigma} = \Phi(\mathbf{x}^{(n)})$ using [38]

6 # Setting the number of diffusion steps to denoise $\mathbf{x}^{(n)}$

7 Set $\hat{t}^* = \alpha^{-1}(\hat{\sigma}^2)$

8 # Sampling the splitting variable $\mathbf{z}^{(n)}$ according to (7)

9 Set $\mathbf{u}_{\hat{t}^*} = \mathbf{x}^{(n)}$

10 **for** $j \leftarrow \hat{t}^*$ **downto** 1 **do**

11 | Draw $\mathbf{u}_{j-1} \sim q_\theta(\mathbf{u}_{j-1} | \mathbf{u}_j)$ according to (10)

12 **end for**

13 Set $\mathbf{z}^{(n)} = \mathbf{u}_0$

14 end for

Output: Collection of samples $\{\mathbf{x}^{(n)}, \mathbf{z}^{(n)}\}_{n=N_{\text{bi}}+1}^{N_{\text{MC}}}$ asymptotically distributed according to (4).

III. APPLICATION TO BAYESIAN INVERSE PROBLEMS

The proposed PnP-SGS method is now instantiated for three different imaging problems, namely deblurring, inpainting and superresolution, following the protocols already considered in [23]. The considered linear Gaussian inverse problems define an archetypal class of problems that can efficiently be tackled by the proposed method. More specifically, a degraded image \mathbf{y} is observed and one wants to infer a restored image \mathbf{x} under the linear model

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n} \quad (14)$$

where \mathbf{H} is a forward operator and \mathbf{n} accounts for noise or error modeling. Assuming that \mathbf{n} is a Gaussian random vector with covariance matrix $\mathbf{\Omega}^{-1}$, the likelihood function associated with the observation \mathbf{y} writes

$$p(\mathbf{y} | \mathbf{x}) \propto \exp \left[-\frac{1}{2} (\mathbf{H}\mathbf{x} - \mathbf{y})^T \mathbf{\Omega} (\mathbf{H}\mathbf{x} - \mathbf{y}) \right].$$

In most applicative contexts, inferring the unknown parameter vector \mathbf{x} from the observation vector \mathbf{y} under the linear model (14) is known to be an ill-posed or ill-conditioned inverse problem. A common approach to tackle such problems consists in using some regularization defined through the choice of a prior distribution $p(\mathbf{x}) \propto \exp[-g(\mathbf{x})]$, leading to the posterior distribution (3). Instead of explicitly specifying the potential function $g(\cdot)$ in (3), the proposed PnP-SGS algorithm targets an augmented posterior similar to (4) to capitalize on a pre-trained denoising diffusion model presented in Section II-B.

The three considered tasks mainly differ by the nature of the linear operator \mathbf{H} . Following the SGS algorithmic scheme, a special care should be taken to ensure an efficient sampling according to the conditional posterior (6) which involves \mathbf{H} , see Algo. 1, line 3. Since the sampling according to (7) does not depend on the forward operator, it is achieved in a unique manner from a DDPM. Thus the sequel of this section is only devoted to the technical derivations associated with (6). Experimental results obtained by the proposed PnP-SGS will be reported in Section IV.

A. Image deblurring

In this setup, the operator \mathbf{H} is assumed to be an $N \times N$ circulant convolution matrix hereafter denoted \mathbf{B} associated to a spatially invariant blurring kernel. The noise covariance matrix is assumed to be diagonal, i.e., $\mathbf{\Omega}^{-1} = \text{diag}[\sigma_1^2, \dots, \sigma_N^2]$ where distinct diagonal elements mimic a spatially-variant noise level. The proposed PnP-SGS algorithm yields the conditional distribution (6) defined here as

$$p(\mathbf{x} | \mathbf{z}, \mathbf{y}; \rho^2) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_{\mathbf{x}}, \mathbf{Q}_{\mathbf{x}}^{-1}) \quad (15)$$

with

$$\begin{cases} \mathbf{Q}_{\mathbf{x}} = \mathbf{B}^T \mathbf{\Omega} \mathbf{B} + \frac{1}{\rho^2} \mathbf{I}_N \\ \boldsymbol{\mu}_{\mathbf{x}} = \mathbf{Q}_{\mathbf{x}}^{-1} \left(\mathbf{B}^T \mathbf{\Omega} \mathbf{y} + \frac{1}{\rho^2} \mathbf{z} \right). \end{cases} \quad (16)$$

Thanks to the splitting trick inherent to the proposed PnP-SGS algorithm, this step does not depend on $g(\cdot)$ and boils down to a high-dimensional Gaussian sampling task [41]. Direct sampling according to this distribution may remain a challenging task, mainly due to the presence of the precision

matrix $\mathbf{\Omega}$ which prevents a direct computation in the Fourier domain. It can be efficiently achieved by introducing an auxiliary variable and sampling according to the augmented distribution, as detailed in [42]. Finally, sampling from (7) is straightforward using the pre-trained network as discussed in Section II-B.

B. Image inpainting

Image inpainting problems aim at recovering an original image $\mathbf{x} \in \mathbb{R}^N$ from the noisy and partial measurements $\mathbf{y} \in \mathbb{R}^M$ under the linear model (14). The operator $\mathbf{H} \in \{0, 1\}^{M \times N}$ now stands for a binary matrix associated with an irregular subsampling with $M \ll N$. The noise is assumed to be white and Gaussian such that $\mathbf{\Omega}^{-1} = \sigma^2 \mathbf{I}_M$. As for the deblurring task, the conditional distribution (6) is (15) with

$$\begin{cases} \mathbf{Q}_{\mathbf{x}} = \frac{1}{\sigma^2} \mathbf{H}^T \mathbf{H} + \frac{1}{\rho^2} \mathbf{I}_N \\ \boldsymbol{\mu}_{\mathbf{x}} = \mathbf{Q}_{\mathbf{x}}^{-1} \left(\frac{1}{\sigma^2} \mathbf{H}^T \mathbf{y} + \frac{1}{\rho^2} \mathbf{z} \right). \end{cases} \quad (17)$$

It is worth noting that the direct operator \mathbf{H} is an $M \times N$ binary matrix that can be obtained by extracting a subset of M rows of the $N \times N$ identity matrix. Thus, the matrix product $\mathbf{H}^T \mathbf{H}$ is a diagonal matrix with M non-zero values equal to 1. As a consequence the precision matrix $\mathbf{Q}_{\mathbf{x}}$ is diagonal and the conditional distribution is easy to sample from.

C. Image super-resolution

Image super-resolution is characterized by a forward model composed of a blurring kernel followed by a subsampling step. The forward operator writes

$$\mathbf{H} = \mathbf{S} \mathbf{B} \quad (18)$$

where \mathbf{B} is a $N \times N$ circulant convolution matrix, as in Section III-A, and $\mathbf{S} \in \{0, 1\}^{M \times N}$ is associated with a binary mask, as in Section III-B. The noise \mathbf{n} is assumed to be white and Gaussian. To fully benefit from the advantages of the SGS, two auxiliary variables \mathbf{z}_1 and \mathbf{z}_2 are introduced to define the augmented posterior distribution

$$p(\mathbf{x}, \mathbf{z}_1, \mathbf{z}_2 | \mathbf{y}; \rho_1^2, \rho_2^2) \propto \exp \left[-\frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{S} \mathbf{z}_1\|^2 - \frac{1}{2\rho_1^2} \|\mathbf{z}_1 - \mathbf{B} \mathbf{x}\|^2 - g(\mathbf{z}_2) - \frac{1}{2\rho_2^2} \|\mathbf{z}_2 - \mathbf{x}\|^2 \right] \quad (19)$$

This double splitting leads to a SGS algorithm which samples alternatively according to the conditional distributions

$$p(\mathbf{z}_1 | \mathbf{x}, \mathbf{y}) \propto \exp \left[-\frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{S} \mathbf{z}_1\|^2 - \frac{1}{2\rho_1^2} \|\mathbf{z}_1 - \mathbf{B} \mathbf{x}\|^2 \right] \quad (20)$$

$$p(\mathbf{x} | \mathbf{z}_1, \mathbf{z}_2) \propto \exp \left[-\frac{1}{2\rho_1^2} \|\mathbf{z}_1 - \mathbf{B} \mathbf{x}\|^2 - \frac{1}{2\rho_2^2} \|\mathbf{z}_2 - \mathbf{x}\|^2 \right] \quad (21)$$

$$p(\mathbf{z}_2 | \mathbf{x}) \propto \exp \left[-g(\mathbf{z}_2) - \frac{1}{2\rho_2^2} \|\mathbf{z}_2 - \mathbf{x}\|^2 \right]. \quad (22)$$

The two distributions (20) and (21) correspond to the previously discussed tasks of inpainting and deblurring, respectively. Sampling according to the last one (22) is achieved thanks to a DDPM used as a stochastic PnP denoiser.

TABLE I

FFHQ 256×256 DATA SET: IMAGE RECONSTRUCTION (PSNR, SSIM) OBTAINED BY THE COMPARED METHODS. **BOLD**: BEST, UNDERLINE: SECOND.

		PnP-SGS	SPA	TV-ADMM	PnP-ADMM	DPIR	Score-SDE	DDRM	MCG	DPS
Inpainting	PSNR \uparrow	32.59	24.09	22.03	8.41	24.41	13.52	9.19	21.57	<u>25.23</u>
	SSIM \uparrow	0.913	0.524	0.784	0.325	0.809	0.437	0.319	0.751	<u>0.851</u>
	FID \downarrow	<u>37.36</u>	71.12	181.56	123.61	52.73	76.54	69.71	29.26	38.82
	LPIPS \downarrow	0.144	0.785	0.463	0.692	0.398	0.612	0.587	<u>0.286</u>	0.262
Deblurring (Gaussian)	PSNR \uparrow	27.96	23.17	22.37	24.93	<u>26.09</u>	7.12	23.36	6.72	24.25
	SSIM \uparrow	0.837	0.499	0.801	<u>0.812</u>	<u>0.820</u>	0.109	0.767	0.051	0.811
	FID \downarrow	59.667	78.67	186.74	90.42	80.18	109.07	<u>74.92</u>	101.2	62.72
	LPIPS \downarrow	0.331	0.452	0.507	0.441	0.392	0.403	<u>0.332</u>	0.340	0.444
Deblurring (motion)	PSNR \uparrow	28.46	17.73	21.36	24.65	27.33	6.58	N/A	6.72	20.92
	SSIM \uparrow	0.828	0.211	0.751	0.825	<u>0.814</u>	0.102	N/A	0.055	0.808
	FID \downarrow	60.01	103.87	152.39	89.08	<u>78.95</u>	292.28	N/A	310.5	56.08
	LPIPS \downarrow	0.294	0.446	0.508	0.405	<u>0.386</u>	0.657	N/A	0.702	0.389
Superres. ($\times 4$)	PSNR \uparrow	<u>25.99</u>	N/A	23.86	25.55	26.14	17.62	25.36	19.97	25.67
	SSIM \uparrow	0.862	N/A	0.803	<u>0.865</u>	0.889	0.617	0.835	0.703	0.852
	FID \downarrow	<u>58.82</u>	N/A	110.64	66.52	63.98	96.72	62.15	87.64	52.82
	LPIPS \downarrow	0.279	N/A	0.428	0.353	0.319	0.563	<u>0.294</u>	0.520	0.337

IV. EXPERIMENTS

A. Experimental setup

Experiments have been conducted on two data sets each composed of 1000 RGB images of size 256×256 with various characteristics, namely FFHQ 256×256 [43], and Imagenet 256×256 [44]. Pre-trained diffusion models have been directly taken from [28], [45] and used without any additional fine-tuning. The test images have never been seen by the model while training to avoid any bias due to potentially over-fitted pre-trained models. All images are normalized to the range $(0, 1)$. For the inversion tasks described in Section III, the forward measurement operators have been designed as follows:

- *deblurring*: two blurring kernels are considered, namely a Gaussian blur with a kernel size of 61×61 with standard deviation of 3.0, and a randomly generated motion blur¹ with size 61×61 and intensity value 0.5,
- *inpainting*: 80% of the total pixels have been randomly masked across all RGB channels,
- *superresolution*: the operator \mathbf{S} corresponds to a down-sampling factor $d = 4$ in both directions and the operator \mathbf{B} stands for a Gaussian blur with a kernel size of 9×9 and a standard deviation of 1.5.

For the inpainting and superresolution tasks, the images have been corrupted with a Gaussian noise whose variance is adjusted to reach a signal-to-noise ratio of $\text{SNR} = 40\text{dB}$. For the deblurring task, to mimic a spatially-variant noise, the diagonal elements σ_i^2 of noise precision matrix Ω^{-1} have been independently randomly drawn such that $\sigma_i = 40$ with probability 0.35 and $\sigma_i = 13$ with probability 0.65.

B. Compared methods & figures-of-merit

The proposed method has been compared to state-of-the-art methods related to the rationales motivating PnP-SGS:

- SPA [23]: split-and-augmented Gibbs sampler is an extension of SGS; in our experiments, it is used with a usual

Tikhonov regularizer for deblurring and superresolution and with total-variation (TV) for inpainting;

- TV-ADMM: ADMM with a TV regularization;
- PnP-ADMM [46]: ADMM with a PnP regularization chosen as DnCNN [7];
- DPIR [15]: HQS with PnP regularization chosen as DRUNet; this can be interpreted as the deterministic counterpart of PnP-SGS;
- Score-SDE [34]: implemented using the same DDPM as the one used by PnP-SGS;
- DDRM [47]: the denoising diffusion restoration model is implemented using the same DDPM as PnP-SGS;
- MCG [48]: manifold constrained gradients;
- DPS [49]: the diffusion posterior sampling is implemented using the same DDPM as the one used by PnP-SGS.

PnP-ADMM, TV-ADMM, DPIR, DDRM and Score-SDE yield point estimates only. In contrast, PnP-SGS provides a comprehensive description of the targeted posterior distribution so that it permits to quantify uncertainties. It yields variances and credibility intervals and multiple statistics of the posterior in addition to the point estimate. Unless otherwise specified, the results associated with the proposed PnP-SGS method correspond to the MMSE estimate $\hat{\mathbf{x}}_{\text{MMSE}}$ approximated by the empirical average (13). Implementation details are reported in Appendix C.

Remark. *The proposed PnP-SGS and its deterministic counterparts PnP-ADMM and DPIR do not leverage pre-trained models with the same architectures, namely the UNet architecture of [29] for PnP-SGS, DnCNN for PnP-ADMM [7] and DRUNet for DPIR [15]. This can be justified by the very different nature of the two approaches. PnP-SGS requires to sample from the conditional posterior (7) and embeds a stochastic denoiser. Using a deterministic denoiser would only provide a point estimate associated to this conditional posterior, i.e., its mean in case of a MMSE denoiser or its mode in case of a MAP denoiser. As a consequence, the use of such conventional deterministic denoisers (such as DnCNN and*

¹Following the code available online at <https://github.com/LeviBorodenko/motionblur>.

TABLE II
 IMAGENET 256 × 256 DATA SET: IMAGE RECONSTRUCTION (PSNR, SSIM) OBTAINED BY THE COMPARED METHODS. **BOLD**: BEST, UNDERLINE: SECOND.

		PnP-SGS	SPA	TV-ADMM	PnP-ADMM	DPIR	Score-SDE	DDRM	MCG	DPS
Inpainting	PSNR ↑	25.22	23.14	20.96	8.39	22.08	18.62	14.29	19.03	18.90
	SSIM ↑	0.870	<u>0.802</u>	0.676	0.300	0.762	0.517	0.403	0.546	0.794
	FID ↓	34.28	41.33	189.3	114.7	37.47	127.1	114.9	39.19	<u>35.87</u>
	LPIPS ↓	0.297	0.323	0.510	0.677	0.448	0.659	0.665	0.414	0.303
Deblurring (Gaussian)	PSNR ↑	21.76	21.08	19.99	<u>21.81</u>	<u>21.81</u>	15.97	22.73	16.32	24.25
	SSIM ↑	0.701	0.577	0.634	0.669	0.612	0.436	<u>0.705</u>	0.441	0.811
	FID ↓	<u>64.12</u>	98.78	155.7	100.6	98.1	120.3	63.02	95.04	64.72
	LPIPS ↓	0.399	0.537	0.588	0.519	0.499	0.667	<u>0.427</u>	0.550	0.444
Deblurring (motion)	PSNR ↑	21.47	20.49	20.79	21.98	<u>22.49</u>	7.21	N/A	5.89	24.92
	SSIM ↑	0.695	0.681	0.677	0.702	<u>0.731</u>	0.120	N/A	0.037	0.859
	FID ↓	47.57	91.51	138.8	89.76	76.11	98.25	N/A	186.9	<u>56.08</u>
	LPIPS ↓	0.372	0.538	0.525	0.483	0.448	0.591	N/A	0.758	<u>0.389</u>
Superres. (×4)	PSNR ↑	24.33	N/A	22.17	23.75	24.30	12.25	<u>24.96</u>	13.39	25.67
	SSIM ↑	<u>0.772</u>	N/A	0.679	0.761	0.769	0.256	0.790	0.227	0.752
	FID ↓	<u>59.09</u>	N/A	130.9	97.27	88.85	170.7	59.57	144.5	50.66
	LPIPS ↓	<u>0.418</u>	N/A	0.523	0.433	0.424	0.701	0.339	0.637	0.337

DRUNet is incompatible with the SGS scheme. Conversely, PnP-ADMM and DPIR require to compute the MAP estimate associated with the conditional distribution (7). Thus, at each iteration of PnP-ADMM and DPIR, using one sample or averaging several samples drawn from the DDPM to perform the denoising step makes the convergence of the resulting PnP-ADMM and DPIR algorithms to be hardly ensured. Moreover, this strategy has been shown to provide poor results (not reproduced in this manuscript for brevity).

The results are first qualitatively evaluated through visual inspection. Quantitative comparisons are also conducted based on four widely-used metrics. The first two criteria are standard image reconstruction metrics, namely peak signal-to-noise-ratio (PSNR) and structural similarity index (SSIM). The two other criteria are perceptual metrics: Fréchet Inception Distance (FID), and Learned Perceptual Image Patch Similarity (LPIPS) distance. Results are averaged over 1000 test images.

C. Experimental results

Tables I and II report the quantitative results in terms of image reconstruction and perceptual metrics for the two data sets FFHQ and Imagenet, respectively. Results appearing as N/A correspond to tasks which have not been implemented by the original authors. The proposed method compares favourably with all the other methods in terms of SNR as well as visual perception metrics. When considering the deblurring or superresolution tasks, DPS and DPIR provide comparable results. However these methods fail to recover good restored images when considering the inpainting task. Some methods perform poorly for some tasks on some data sets, e.g., Score-SDE and MCG for deblurring or PnP-ADMM and DDRM for inpainting on FFHQ images. Such unreliable behaviors have been consistently encountered during the experiments conducted for this work but this seems to be in agreement with the results previously reported in the literature [49]. Note again that DDRM, Score-SDE and DPS rely on a DDPM where the

pre-trained generative model is exactly the same as the one implemented in PnP-SGS.

Fig. 1 permits to assess the performances by visual inspection when inpainting 4 test images taken from the FFHQ and Imagenet data sets. In particular, PnP-SGS is compared to state-of-the-art methods which are known to be robust to measurement noise. PnP-SGS is able to provide high-quality reconstructions that are crisp and realistic. In particular it is able to recover more granular details.

As already stated, the proposed PnP-SGS generate samples asymptotically distributed according to the posterior distribution. These samples can be used to approximate various Bayesian estimators but also to derive credibility intervals. Fig. 2 illustrates this advantage by depicting various restored images (in term of MMSE estimates) as well as 90% credibility intervals for different tasks. This added value cannot be provided by optimization-based methods, e.g., TV-ADMM, PnP-ADMM and DPIR, which provide point estimates only. Besides, stochastic samplers such as DDRM, MCG, Score-SDE and DPS are not able to provide this information either. Indeed, they do not generate multiple samples drawn from a stationary posterior. Several runs of these methods produce outputs that may be individually relevant but that are not consistent between them in their details, in particular because they originate from different noise realizations. This is also why averaging multiple outputs of these methods does not yield reliable MMSE estimators but rather tends to recover blurred images, as illustrated in Fig. 3 (6th right panel for MCG).

When targeting (4), PnP-SGS generates two sets of samples $\{\mathbf{x}^{(n)}\}_n$ and $\{\mathbf{z}^{(n)}\}_n$ that are marginally distributed according to the marginals of $\pi_\rho(\mathbf{x}, \mathbf{z})$. It follows a splitting strategy where the variables \mathbf{x} and \mathbf{z} are coupled thanks to a quadratic kernel that is controlled by the parameter ρ . Thus the posterior means $\hat{\mathbf{x}}_{\text{MMSE}} = \mathbb{E}[\mathbf{x}|\mathbf{y}]$ and $\hat{\mathbf{z}}_{\text{MMSE}} = \mathbb{E}[\mathbf{z}|\mathbf{y}]$ should be very similar up to some variations adjusted by the coupling parameter ρ . Fig. 3 depicts the two estimates as well as pixel-wise 90% credibility intervals. As expected, slight differences




Fig. 1. Inpainting task on the FFHQ (two top rows) and Imagenet data sets (two bottom rows), from left to right: measurement, true image, PnP-SGS, SPA [23], DDRM [47], MCG [48].



Fig. 2. From left to right: measurement, true image, MMSE estimate, pixel-wise 90% credibility intervals. From top to bottom: Gaussian blur, motion blur, superresolution. The first three and the last images are from the FFHQ and Imagenet data sets, respectively.

TABLE III
INPAINTING: COMPUTATIONAL TIMES (S.) OF THE COMPARED METHODS.

PnP-SGS	SPA	PnP-ADMM	DPIR	Score-SDE	DDRM	MCG	DPS
13.81	218.90	3.63	4.18	36.71	29.03	80.10	43.89

are observed. In particular, the point estimate $\hat{\mathbf{x}}_{\text{MMSE}}$ seems to be characterized by sharper details (better viewed by zooming on screen). Recall that this estimate is closer to the observation, while $\hat{\mathbf{z}}_{\text{MMSE}}$ is closer to the prior, therefore smoother. Fig. 4 (top) depicts the estimate $\hat{\mathbf{x}}_{\text{MMSE}}$ as a function of the number n of Gibbs iterations. This evolution can be compared in Fig. 4 (bottom) to that of the estimate of the denoised image $\hat{\mathbf{u}}_0 \triangleq \mathbb{E}[\mathbf{u}_0 | \mathbf{x}^{(n)}]$, which is approximated by averaging DDPM outputs. The two trajectories are shown to evolve synchronously along time.

Table III reports the execution times for the task of inpainting of the various methods implemented on a single GTX 2080Ti GPU. Noticeably, the computational time of PnP-SGS is similar to its competitors. In particular, this stochastic MCMC method (13.81s) is more than twice faster than Score-SDE (36.71s). It remains within a factor less than 4 with respect to PnP-ADMM (3.63s) and DPIR (4.18s), its deterministic counterparts. These results may appear surprising since the pre-trained DnCNN and DRUNet models are used only once per iteration of PnP-ADMM and DPIR (for about ten iterations) while the denoising model is used several



Fig. 3. Image inpainting, from left to right: measurement, true image, MMSE estimate $\hat{\mathbf{x}}_{\text{MMSE}}$, MMSE estimate $\hat{\mathbf{z}}_{\text{MMSE}}$, pixel-wise 90% credibility intervals, averaged samples generated by MCG [48]



Fig. 4. Top: estimate $\hat{\mathbf{x}}_{\text{MMSE}}$ as a function of the number n of Gibbs iteration. Bottom: estimate of the denoised image $\hat{\mathbf{u}}_0 \triangleq \mathbb{E}[\mathbf{u}_0 | \mathbf{x}^{(n)}]$ produced by the DDPM (approximated by Monte Carlo integration).

dozen times per iteration of PnP-SGS (for about a hundred iterations). This finding should be mitigated by the fact that the efforts dedicated by different research groups to optimize their respective models are not comparable. Yet, the price to pay to get quantified uncertainties sounds very reasonable. It is worth noting that using a DDPM-based PnP with SGS significantly reduces the number of iterations required by the sampler to reach the steady regime, which explains the reduced computational cost with respect to SPA.

V. CONCLUSION

This work proposes the plug-and-play split Gibbs sampler (PnP-SGS) as a stochastic counterpart of the well-known PnP-ADMM. Thanks to the SGS divide-to-conquer strategy, the PnP-SGS algorithm permits to target a posterior distribution that involves an implicit PnP prior where the regularization is ensured by some efficient stochastic denoiser. The proposed methodology can make use of any well-suited PnP prior, depending on the final application. For instance, it can be based on a denoising diffusion probabilistic model (DDPM), as proposed here, since it appears that a DDPM can be turned into a Bayesian sampler of a denoising problem. With the same versatility as PnP-ADMM, sampling from the posterior distribution noticeably permits to build credibility intervals on top of point estimates. Extensive numerical experiments show that the proposed approach competes favourably with existing state-of-the-art models on typical imaging problems, namely

deblurring, inpainting and superresolution. The quantitative performances are at least comparable when not better, while the computational times remain very moderate as well. PnP-SGS appears as a scalable MCMC sampling method that can benefit from the most recent progress in machine (deep) learning at the price of a reasonable computational cost. One open question deals with the theoretical finding which may justify the replacement of the sampling according to the conditional posterior (7) by the backward diffusion (10). This requires to identify the sufficient conditions to guarantee that the transition kernel $q_\theta(\mathbf{u}_0 | \mathbf{u}_t)$ properly defines a posterior distribution associated to a denoising task. In the framework of PnP-ADMM, this question is equivalent to discussing the conditions required by the denoiser to be a proximal operator.

APPENDIX A ESTIMATING THE NOISE LEVEL

The stochastic denoising task corresponding to the conditional distribution (7) requires to adjust the number t^* of backward steps to be performed (see Algo. 1, line 10). As discussed in Section II-D, this number is related to the level $\hat{\sigma}$ of the Gaussian noise assumed to affect the current state $\mathbf{x}^{(n)}$ through the noise scheduling function $\alpha(\cdot)$, i.e., $\hat{\sigma}^2 = \alpha(t^*)$. As already stated, the conditional posterior (7) can be interpreted as a Bayesian formulation of a denoising task, i.e., recovering a noise-free image \mathbf{z} from the a version corrupted by an additive white Gaussian noise with variance

ρ^2 . Thus, one natural choice would consist in prescribing ρ^2 as the amount of noise, i.e., $\hat{\sigma}^2 = \rho^2$. However, in practice, this strategy is empirically shown to be less efficient than an adaptive tuning of this noise level. More precisely, the noise level is estimated directly from the current state $\mathbf{x}^{(n)}$ at each iteration of the Gibbs sampler, i.e., $\hat{\sigma} = \Phi(\mathbf{x}^{(n)})$ where $\Phi(\cdot)$ is a noise level estimator. The problem of estimating the level of the noise corrupting natural images has motivated plenty of research works, see [38]–[40]. In our implementations, this estimation has been carried out following the strategy proposed in [38]. This robust wavelet-based estimator is already implemented in the library `scikit-image` (aka `skimage`) as the function `estimate_sigma()`. When handling RGB natural images, this function has been used with the parameter `average_sigmas=True` to average the noise level estimates over the three channels. It is worth noting that the estimated level of noise tends to stabilize beyond the burn-in period. This leads to a simultaneous stabilization of the estimated number t^* of backward steps, as discussed in Appendix C-A (see also Fig. 5).

APPENDIX B

INVERTING THE NOISE SCHEDULING FUNCTION $\alpha(\cdot)$

Given the current estimate of the noise level $\hat{\sigma}$, sampling according to (7) is achieved by performing the backward diffusion with kernel (10) from a time instant \hat{t}^* such that $\hat{\sigma}^2 = \alpha(\hat{t}^*)$ where $\alpha(t)$ is defined by (12). This noise scheduling function is controlled by the function $b(\cdot)$ that adjusts the variance of the forward transition kernel (9) from $t - 1$ to t . Various choices of $b(\cdot)$ exist in the literature.

The pre-trained DDPM used during the experiments conducted on FFHQ is governed by the linearly increasing function

$$b(t) = b(0) + rt \quad (23)$$

where $b(0) = 10^{-4}$ and the slope r has been adjusted such that $b(T) = 2.0 \times 10^{-2}$ [29].

Conversely, when handling the ImageNet data set, the DDPM used in the experiments was trained using a cosine-based variance schedule [36]

$$\alpha(t) = 1 - \frac{\gamma(t)}{\gamma(0)} \quad (24)$$

with $\gamma(t) = \cos\left(\frac{\pi}{2} \frac{t/T+s}{1+s}\right)^2$ where s is a small offset preventing $\alpha(t)$ to be too small when $t \rightarrow 0$. In both cases, an explicit inverse scheduling function $\alpha^{-1}(\cdot)$ can be derived, which yields $\hat{t}^* = \alpha^{-1}(\hat{\sigma}^2)$. For more complex noise scheduling functions, an alternative is to use a tabbing strategy, which saves computation cost as well. Given a pre-computed list of $T + 1$ values $\alpha = \{\alpha(0), \dots, \alpha(T)\}$, the diffusion start time is set to

$$\hat{t}^* = \underset{t \in \{0, \dots, T\}}{\operatorname{argmin}} |\alpha(t) - \hat{\sigma}^2|.$$

For our experiments, the scheduling functions were sampled on $T = 1000$ regularly spaced time instants.

APPENDIX C EXPERIMENTAL DETAILS

A. Proposed PnP-SGS

1) *Pre-trained models*: For the experiments on the FFHQ data set, the pre-trained DDPM has been taken from [45] also available online². For the Imagenet data set, we have used the pre-trained DDPM of [28] and available online³. As detailed in Appendix B, these two pre-trained models differ by their respective noise scheduling functions $\alpha(\cdot)$.

2) *Coupling parameter*: The coupling parameter ρ defining the augmented distribution (4) controls the trade-off between the speed of the sampling algorithm and the quality of the approximation used. It is a sensitive hyperparameter whose impact on the asymptotic and pre-asymptotic convergences of the SGS has been deeply studied in [50]. Some recommendations are also provided in [50] to fix this parameter for a posterior distribution defined by a potential $U(\mathbf{x}) = -f(\mathbf{x}, \mathbf{y}) - g(\mathbf{x})$. In this work, since the potential $g(\mathbf{x})$ associated to the prior is not explicitly known, such recommendations can be hardly followed. Moreover, using a naive sequence of decreasing values of the parameter, i.e., $\rho_k \xrightarrow{k \rightarrow \infty} 0$ in a fashion similar to the strategy adopted by HQS [15], does not ensure the invariance of the Gibbs kernels, thus breaking the convergence of the Monte Carlo algorithm. Also, this strategy is empirically shown to perform poorly in practice. Thus, this hyperparameter has been adjusted following a grid-search procedure to reach the best results, namely $\rho = 0.7$ for FFHQ and $\rho = 1.625$ for ImageNet.

3) *Number of iterations*: For all experiments, the number of iterations of the PnP-SGS has been fixed as $N_{\text{MC}} = 100$ including $N_{\text{bi}} = 20$ burn-in iterations.

4) *Number of denoising steps*: In Section IV, the estimated number \hat{t}^* of denoising steps is automatically adjusted by the procedure described in Section II-D. At the first iteration of the PnP-SGS, \hat{t}^* is usually a fraction of T . Along the iterations of the PnP-SGS, this number reduces and then stabilizes around a small fraction of T , as illustrated in Fig. 5 where the initial value of \hat{t}^* is around $0.07T$ for the inpainting task. Moreover, during the burn-in period, instead of applying the kernel (10) for $t = \hat{t}^*, \dots, 1$, we suspend the process in the middle of the diffusion, i.e., $t = \hat{t}^*/2, \dots, \frac{\hat{t}^*}{2}$. Stopping the reverse diffusion at time $t = \hat{t}^*/2$ (instead of $t = 1$), and thus defining $\mathbf{z}^{(n)}$ as the noisy image $\mathbf{u}_{\hat{t}^*/2}$ (instead of the fully denoised image \mathbf{u}_0), leads to best results. Intuitively, at the beginning of the sampling procedure, this strategy avoids the potential hallucination behavior of the DDPM. It prevents the DDPM to sample a new image from an uninformative $\mathbf{x}^{(n)}$ (as long as the samples drawn from the first kernel (6) are not sufficiently informative). This early-stopping trick not only provides empirically better results but also allows the computational burden to be lightened by reducing the number of DNN evaluations.

²https://github.com/jychoi118/ilvr_admin

³<https://github.com/openai/guided-diffusion>

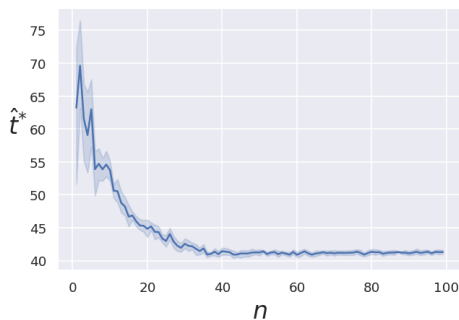


Fig. 5. Inpainting: evolution of \hat{t}^* along the PnP-SGS iterations ($T = 1000$). Results have been averaged over 100 runs conducted on the same image. Shaded areas stand for the corresponding standard deviation.

B. Compared methods

The methods DDRM, MCG and Score-SDE are implemented using the same pre-trained model as PnP-SGS (see above). Additional details are listed below:

- TV-ADMM uses the isotropic regularization. The regularization parameter λ and some penalty parameter ρ linked to the splitting have been adjusted by grid search to reach the best performance. Final values are $(\lambda, \rho) = (2.7 \times 10^{-2}, 1.4 \times 10^1)$ for deblurring, $(\lambda, \rho) = (2.7 \times 10^{-2}, 1.0 \times 10^{-2})$ for inpainting and $(\lambda, \rho_1, \rho_2) = (2.7 \times 10^{-2}, 1.0 \times 10^{-2})$ for superresolution which requires a double splitting.
- PnP-ADMM: the implementation is from the `SCICO`⁴ library. The parameters are set to $\rho = 0.2$ (ADMM penalty parameter) and $\text{maxiter} = 12$. The proximal mappings use the pre-trained DnCNN denoiser [7].
- DPIR: the implementation is from the official repository⁵. The number of iterations is set to $\text{maxiter} = 15$ and the coupling parameter is fixed to $\lambda = 0.23$ for each experiment. The proximal mappings employ the pre-trained DRUNet denoiser, which is accessible through the `DeepInverse` library⁶
- DDRM: all experiments have been performed with the default setting $\eta_B = 1.0$ and $\eta = 0.85$. For the Gaussian deblurring task, the forward model was implemented by separable 1D convolutions for efficient SVD.
- MCG: the noise scheduling function $\alpha(\cdot)$ has been chosen as the one used by PnP-SGS. At each step, complementary data consistency steps are applied as Euclidean projections onto the measurement set $\mathcal{C} = \{\mathbf{x}_i \mid \mathbf{H}\mathbf{x}_i = \mathbf{y}_i, \mathbf{y}_i \sim p(\mathbf{y}_i \mid \mathbf{y}_0)\}$
- Score-SDE solves the inverse problems by iteratively applying a denoising step followed by data consistency projections onto the measurement set \mathcal{C} , as in MCG.
- DPS blends a diffusion step with manifold constrained gradient without a strict measurement consistency projection step.

REFERENCES

- [1] J. Kaipio and E. Somersalo, *Statistical and computational inverse problems*. Springer Science & Business Media, 2006, vol. 160.
- [2] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, “The convex geometry of linear inverse problems,” *Foundations of Computational mathematics*, vol. 12, pp. 805–849, 2012.
- [3] A. Repetti, M. Pereyra, and Y. Wiaux, “Scalable Bayesian uncertainty quantification in imaging inverse problems via convex optimization,” *SIAM J. Imag. Sci.*, vol. 12, no. 1, pp. 87–118, 2019.
- [4] A. Chambolle, “An algorithm for total variation minimization and applications,” *J. Math. Imag. Vision*, vol. 20, pp. 89–97, 2004.
- [5] C. Louchet and L. Moisan, “Posterior expectation of the total variation model: properties and experiments,” *SIAM J. Imag. Sci.*, vol. 6, no. 4, pp. 2640–2684, 2013.
- [6] C. Dong, C. C. Loy, K. He, and X. Tang, “Learning a deep convolutional network for image super-resolution,” in *Proc. IEEE Eur. Conf. Computer Vision (ECCV)*. Springer, Sept. 2014, pp. 184–199.
- [7] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, “Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising,” *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, 2017.
- [8] K. Zhang, W. Zuo, and L. Zhang, “FFDNet: Toward a fast and flexible solution for CNN-based image denoising,” *IEEE Trans. Image Process.*, vol. 27, no. 9, pp. 4608–4622, 2018.
- [9] E. Schwartz, R. Giryes, and A. M. Bronstein, “Deepisp: Toward learning an end-to-end image processing pipeline,” *IEEE Trans. Image Process.*, vol. 28, no. 2, pp. 912–923, 2018.
- [10] K. Gregor and Y. LeCun, “Learning fast approximations of sparse coding,” in *Proc. Int. Conf. Machine Learning (ICML)*, 2010, pp. 399–406.
- [11] S. V. Venkatakrisnan, C. A. Bouman, and B. Wohlberg, “Plug-and-play priors for model based reconstruction,” in *Proc. IEEE Global Conf. Signal and Information Processing*. IEEE, 2013, pp. 945–948.
- [12] D. Geman and C. Yang, “Nonlinear image recovery with half-quadratic regularization,” *IEEE Trans. Image Process.*, vol. 4, no. 7, pp. 932–946, 1995.
- [13] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein *et al.*, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends in Machine learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [14] N. Parikh and S. Boyd, “Proximal algorithms,” *Foundations and Trends in Optimization*, vol. 1, no. 3, p. 127–239, 2014.
- [15] K. Zhang, Y. Li, W. Zuo, L. Zhang, L. Van Gool, and R. Timofte, “Plug-and-play image restoration with deep denoiser prior,” *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 44, no. 10, pp. 6360–6376, 2022.
- [16] R. Ahmad, C. A. Bouman, G. T. Buzzard, S. Chan, S. Liu, E. T. Reehorst, and P. Schniter, “Plug-and-play methods for magnetic resonance imaging: Using denoisers for image recovery,” *IEEE Trans. Signal Process.*, vol. 37, no. 1, pp. 105–116, 2020.
- [17] M. G. Kendall *et al.*, *The advanced theory of statistics*. Charles Griffin and Co., Ltd., London, 1946, no. 2nd Ed.
- [18] X. Cai, M. Pereyra, and J. D. McEwen, “Uncertainty quantification for radio interferometric imaging– I. Proximal MCMC methods,” *Month. Not. Roy. Astronomical Soc.*, vol. 480, no. 3, pp. 4154–4169, 2018.
- [19] E. Gaume, L. Gaál, A. Viglione, J. Szolgay, S. Kohnová, and G. Blöschl, “Bayesian MCMC approach to regional flood frequency analyses involving extraordinary flood events at ungauged sites,” *Journal of hydrology*, vol. 394, no. 1-2, pp. 101–117, 2010.
- [20] M. Pereyra, P. Schniter, É. Chouzenoux, J. Pesquet, J. Tourneret, A. O. Hero, and S. McLaughlin, “A survey of stochastic simulation and optimization methods in signal processing,” *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 2, pp. 224–241, 2016.
- [21] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, “Hybrid Monte Carlo,” *Phys. Lett. B*, vol. 195, no. 2, pp. 216–222, 1987.
- [22] M. Pereyra, “Maximum-a-posteriori estimation with Bayesian confidence regions,” *SIAM J. Imag. Sci.*, vol. 10, no. 1, pp. 285–302, 2017.
- [23] M. Vono, N. Dobigeon, and P. Chainais, “Split-and-augmented Gibbs sampler – Application to large-scale inference problems,” *IEEE Trans. Signal Process.*, vol. 67, no. 6, pp. 1648–1661, 2019.
- [24] —, “Asymptotically exact data augmentation: Models, properties, and algorithms,” *J. Comput. Graph. Stat.*, vol. 30, no. 2, pp. 335–348, 2020.
- [25] D. Im Im, S. Ahn, R. Memisevic, and Y. Bengio, “Denoising criterion for variational auto-encoding framework,” in *Proc. Ass. Adv. Artificial Intelligence (AAAI)*, vol. 31, no. 1, 2017.
- [26] L. D. Tran, S. M. Nguyen, and M. Arai, “Gan-based noise model for denoising real images,” in *Proc. Asian Conf. Computer Vision (ACCV)*, 2020.

⁴<https://scico.readthedocs.io>

⁵<https://github.com/cszn/DPIR/tree/master>

⁶<https://deepinv.github.io/deepinv/deepinv.denoisers.html>

- [27] Y. Song, L. Shen, L. Xing, and S. Ermon, "Solving inverse problems in medical imaging with score-based generative models," in *Proc. IEEE Int. Conf. Learn. Represent. (ICLR)*, 2022.
- [28] P. Dhariwal and A. Nichol, "Diffusion models beat GANs on image synthesis," in *Adv. in Neural Information Process. Systems (NeurIPS)*, vol. 34, 2021, pp. 8780–8794.
- [29] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Adv. in Neural Information Process. Systems (NeurIPS)*, vol. 33, 2020, pp. 6840–6851.
- [30] D. Kingma, T. Salimans, B. Poole, and J. Ho, "Variational diffusion models," in *Adv. in Neural Information Process. Systems (NeurIPS)*, vol. 34, 2021, pp. 21 696–21 707.
- [31] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *Proc. Int. Conf. Machine Learning (ICML)*. PMLR, 2015, pp. 2256–2265.
- [32] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," in *Adv. in Neural Information Process. Systems (NeurIPS)*, vol. 32, 2019.
- [33] ———, "Improved techniques for training score-based generative models," in *Adv. in Neural Information Process. Systems (NeurIPS)*, vol. 33, 2020, pp. 12 438–12 448.
- [34] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," in *Proc. IEEE Int. Conf. Learn. Represent. (ICLR)*, 2021.
- [35] C.-W. Huang, J. H. Lim, and A. C. Courville, "A variational perspective on diffusion-based generative models and score matching," in *Adv. in Neural Information Process. Systems (NeurIPS)*, vol. 34, 2021, pp. 22 863–22 876.
- [36] A. Q. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models," in *Proc. Int. Conf. Machine Learning (ICML)*. PMLR, 2021, pp. 8162–8171.
- [37] H. Chung, B. Sim, and J. C. Ye, "Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction," in *Proc. Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 12 413–12 422.
- [38] D. L. Donoho and J. M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika*, vol. 81, no. 3, pp. 425–455, 1994.
- [39] X. Guo, F. Liu, and X. Tian, "Gaussian noise level estimation for color image denoising," *J. Opt. Soc. Am. A*, vol. 38, no. 8, pp. 1150–1159, 2021.
- [40] Y. Li, C. Liu, X. You, and J. Liu, "A single-image noise estimation algorithm based on pixel-level low-rank low-texture patch and principal component analysis," *Sensors*, vol. 22, no. 22, p. 8899, 2022.
- [41] M. Vono, N. Dobigeon, and P. Chainais, "High-dimensional Gaussian sampling: a review and a unifying approach based on a stochastic proximal point algorithm," *SIAM Review*, vol. 64, no. 1, pp. 3–56, 2022.
- [42] Y. Marnissi, E. Chouzenoux, A. Benazza-Benyahia, and J.-C. Pesquet, "An auxiliary variable method for Markov chain Monte Carlo algorithms in high dimension," *Entropy*, vol. 20, no. 2, 2018.
- [43] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proc. Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 4401–4410.
- [44] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Proc. Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 248–255.
- [45] J. Choi, S. Kim, Y. Jeong, Y. Gwon, and S. Yoon, "ILVR: Conditioning method for denoising diffusion probabilistic models," *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, 2021.
- [46] S. H. Chan, X. Wang, and O. A. Elgendy, "Plug-and-play ADMM for image restoration: Fixed-point convergence and applications," *IEEE Trans. Comput. Imag.*, vol. 3, no. 1, pp. 84–98, 2016.
- [47] B. Kawar, M. Elad, S. Ermon, and J. Song, "Denoising diffusion restoration models," in *Adv. in Neural Information Process. Systems (NeurIPS)*, 2022.
- [48] H. Chung, B. Sim, D. Ryu, and J. C. Ye, "Improving diffusion models for inverse problems using manifold constraints," in *Adv. in Neural Information Process. Systems (NeurIPS)*, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Eds., 2022.
- [49] H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, and J. C. Ye, "Diffusion posterior sampling for general noisy inverse problems," in *Proc. IEEE Int. Conf. Learn. Represent. (ICLR)*, 2023.
- [50] M. Vono, D. Paulin, and A. Doucet, "Efficient MCMC sampling with dimension-free convergence rate using ADMM-type splitting," *Journal of Machine Learning Research*, vol. 23, no. 25, pp. 1–69, 2022.