



MCMC Correction of Score-Based Diffusion Models for Model Composition

Downloaded from: <https://research.chalmers.se>, 2026-05-01 14:23 UTC

Citation for the original published paper (version of record):

Sjöberg, A., Lindqvist, J., Onnheim, M. et al (2026). MCMC Correction of Score-Based Diffusion Models for Model Composition. *Entropy*, 28(3). <http://dx.doi.org/10.3390/e28030351>

N.B. When citing this work, cite the original published paper.

Article

MCMC Correction of Score-Based Diffusion Models for Model Composition

Anders Sjöberg ^{1,2,*}, Jakob Lindqvist ^{2,†}, Magnus Önnheim ¹, Mats Jirstrand ^{1,2} and Lennart Svensson ²

¹ Fraunhofer-Chalmers Centre, SE-412 88 Gothenburg, Sweden; magnus.onnheim@algeno.se (M.Ö.); mats.jirstrand@fcc.chalmers.se (M.J.)

² Department of Electrical Engineering, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden; jakob@echophi.ai (J.L.); lennart.svensson@chalmers.se (L.S.)

* Correspondence: andsjob@chalmers.se

† These authors contributed equally to this work.

Abstract

Diffusion models can be parameterized in terms of either score or energy function. The energy parameterization is attractive as it enables sampling procedures such as Markov Chain Monte Carlo (MCMC) that incorporates a Metropolis–Hastings (MH) correction step based on energy differences between proposed samples. Such corrections can significantly improve sampling quality, particularly in the context of model composition, where pre-trained models are combined to generate samples from novel distributions. Score-based diffusion models, on the other hand, are more widely adopted and come with a rich ecosystem of pre-trained models. However, they do not, in general, define an underlying energy function, making MH-based sampling inapplicable. In this work, we address this limitation by retaining score parameterization and introducing a novel MH-like acceptance rule based on line integration of the score function. This allows the reuse of existing diffusion models while still combining the reverse process with various MCMC techniques, viewed as an instance of annealed MCMC. Through experiments on synthetic and real-world data, we show that our MH-like samplers yield relative improvements of similar magnitude to those observed with energy-based models, without requiring explicit energy parameterization.

Keywords: diffusion models; energy-based models; annealed MCMC; Metropolis–Hastings correction

1. Introduction

Significant advancements have recently been achieved in generative modeling across various domains [1–3]. These models have become potent priors for a wide range of applications, including code generation [4], text-to-image generation [5], question-answering [2], and many others [6,7]. Among generative models, diffusion models [3,8,9] have arguably emerged as the most powerful class. Diffusion models learn to denoise corrupted inputs in small, gradual steps and are capable of generating samples from complex distributions. They have been successful in many domains, such as generating highly realistic images [10], modeling temporal point processes [11], and even generating neural network parameters [12].

Diffusion models also offer the capability of composed sampling, which combines pre-trained models to generate samples from a new distribution. This approach, known as model composition, has a rich history [13–16]. For diffusion models, the most common form of composition is classifier-guided sampling, where the reverse process is augmented



Academic Editor: Antonio M. Scarfone

Received: 16 January 2026

Revised: 6 March 2026

Accepted: 18 March 2026

Published: 20 March 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

by a separate classifier model [8,10,17], but other compositions have also been explored [18]. The ability to compose new models without having to re-learn the individual components is especially appealing for diffusion models since their ever-increasing size and data hunger make them exceedingly costly to train [19]. Therefore, developing sampling methods that work for pre-trained diffusion models is valuable.

The foundation of composed sampling for diffusion models is score-based, where we interpret diffusion models as predictors of the score function for the marginal distribution at each diffusion step [20]. From this perspective, MCMC methods, such as the Langevin algorithm (LA) [21] or Hamiltonian Monte Carlo (HMC) sampling [22], emerge as viable options to incorporate. Augmenting the standard reverse process with additional MCMC sampling has been shown to improve composed sampling for diffusion models [18,20]. However, we are traditionally restricted to unadjusted variants of these samplers, namely Unadjusted LA (U-LA) and Unadjusted HMC (U-HMC), which only require utilization of the score. For instance, the seminal work by Song and Ermon [9] proposed score-based generative modeling by combining score networks with unadjusted Langevin dynamics. While highly effective, these unadjusted methods lack a rejection mechanism and may introduce discretization bias at larger step sizes. Correcting this bias requires evaluating the unnormalized density for a Metropolis–Hastings (MH) step [23,24]—something that is not directly possible in standard score-parameterized diffusion models since they do not provide access to the underlying energy.

An intriguing alternative to directly modeling the score function is to model the marginal distribution with an energy function, from which the score can be obtained through explicit differentiation [9,25]. This parameterization connects diffusion models and energy-based models (EBMs) [26] and offers several desirable properties. With energy parameterization, we can evaluate the unnormalized density and guarantee a proper score function. This, in turn, enables an MH correction step when employing an MCMC method, where the MH acceptance probability is computed from the energy function. Adding such a correction step has been shown to improve sampling performance in composed models [18]. Nevertheless, the score parameterization remains far more popular, as it avoids the direct computation of the gradient of the log density. In practice, score-based models are significantly easier to train and scale to large datasets and high-dimensional problems than energy-based parameterizations, a practical advantage that has driven their widespread adoption.

In this study, we build on the work in [18] and introduce a novel approach to obtain an MH-like correction step directly from pre-trained diffusion models without relying on an energy-based parameterization. Specifically, we use a connection between the score and the energy to estimate the MH acceptance probability by approximating a line integral along the vector field generated by the score. This enables an improved sampling procedure for various pre-trained score-parameterized diffusion models. We find that our approximate method yields relative improvements of a similar magnitude to those obtained with explicit energy parameterizations, without having to estimate the energy directly.

In summary, our main contributions are:

- We show that MH-like correction sampling can be directly applied to score-based models without requiring additional training.
- We introduce two efficient algorithms to approximate the energy difference used in MH and demonstrate that our pseudo-energy difference more accurately represents analytical energy differences than an explicitly trained energy model in a toy example while performing on par with the energy model on MNIST.
- We establish that the relative improvements in sample quality achieved via MCMC for energy-based models can be closely matched within a purely score-based framework.

In our implementations, the pseudo-MH correction also exhibits favorable practical runtime behavior compared to energy-based MCMC.

2. Background

2.1. Diffusion Models

We consider Gaussian diffusion models initially proposed by [8] and further improved by [3,9]. Starting with a sample from the data distribution $x_0 \sim q(\cdot)$, we construct a Markov chain of latent variables x_1, \dots, x_T by iteratively introducing Gaussian noise to the sample $q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I)$, where $\beta_t \in [0, 1], \forall t = 1, \dots, T$ are known. For large enough T we have $q(x_T) \approx \mathcal{N}(x_T; 0, I)$.

A diffusion model learns to gradually denoise samples by modeling the distribution of the previous sample in the chain $p_\theta(x_{t-1} | x_t), t = 1, \dots, T$. Approximate samples from the data distribution $q(x_0)$ are obtained by starting from $x_T \sim \mathcal{N}(0, I)$ and sequentially sampling less noisy versions of the sample until the noise is removed. This is called the *reverse process*.

The reverse distribution is typically modeled as $p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t))$, since the posterior $q(x_{t-1}|x_t)$ can be well approximated by a Gaussian distribution when the noise magnitude β_t is sufficiently small. The mean is parameterized as $\mu_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sigma_t} \epsilon_\theta(x_t, t) \right)$, where α_t and σ_t are positive and defined by $\{\beta_t\}_{t=1}^T$ [3]. The noise prediction model $\epsilon_\theta(x_t, t)$, typically a neural network, is learned from data. We assume $\Sigma_\theta(x_t, t) = \beta_t I$ throughout unless otherwise stated.

2.2. Energy-Based Models

Energy-based models (EBMs) represent probability distributions with a scalar, non-negative energy function E_θ by assigning low energy to regions of the input space where the probability is high and high energy to regions where the distribution has little or no support:

$$\begin{aligned} p_\theta(x_t, t) &= \frac{1}{Z_\theta(t)} \exp\left(-\frac{1}{\sigma_t} E_\theta(x_t, t)\right), \\ Z_\theta(t) &= \int \exp\left(-\frac{1}{\sigma_t} E_\theta(x_t, t)\right) dx_t. \end{aligned} \quad (1)$$

Here, we define E_θ as a time-dependent function and deliberately choose not to absorb σ_t (introduced in the previous section) into E_θ , to maintain a more explicit connection to diffusion models, as clarified in the next section. This time dependency can be seen as a sequence of energy functions, one for each diffusion step t . The normalization constant Z_θ is typically intractable, prohibiting computing a normalized density. However, Z_θ does not depend on the input x_t , making the so-called *score function* easy to compute,

$$\nabla_x \log p_\theta(x_t, t) = -\frac{1}{\sigma_t} \nabla_x E_\theta(x_t, t), \quad (2)$$

even though the gradient of the energy function can be costly to compute in practice.

2.3. Energy and Score Parameterized Diffusion Models

In the diffusion setting, one approach to training EBMs is denoising score matching (DSM). When the data are perturbed with Gaussian noise, the DSM loss coincides with the diffusion training loss (up to a factor of σ_t^2) [20]. This equivalence arises by identifying the noise prediction model, $\epsilon_\theta(x_t, t)$, with the score of an energy function, i.e., as an EBM:

$$\epsilon_\theta(x_t, t) = \nabla_x E_\theta(x_t, t). \quad (3)$$

Under this identification, $\epsilon_\theta(x_t, t)$ is required to define a proper score. Thus, an EBM and a plain diffusion model differ only in parameterization: the former uses *energy parameterization* via E_θ , while the latter, since ϵ_θ is only a pseudo-score, is referred to as using *score parameterization*.

Both parameterizations have their advantages and disadvantages. The energy parameterization can evaluate density $p_\theta(x_t, t)$ up to normalization $Z_\theta(t)$, which enables various MCMC methods. Furthermore, by making the score equal to the gradient of an actual scalar function, we ensure a proper score. On the other hand, to evaluate the score function, E_θ must be explicitly differentiated, which can be costly.

The score parameterization is more flexible as it predicts an arbitrary vector field. While there is some empirical evidence that this improves sampling performance in diffusion processes [18], this difference may primarily stem from model architecture [25]. Nevertheless, the score parameterization’s direct estimation of the score function makes it more efficient for reverse process sampling and remains the more widely adopted approach. In the next section, we describe how these parameterizations affect the design of MCMC samplers for diffusion models.

2.4. MCMC Sampling for Diffusion Models

MCMC sampling is a promising strategy for improving diffusion model sampling since it can be combined with the reverse process. MCMC methods are naturally defined through transition kernels, and just like the reverse process there are MCMC methods that base their transitions on the score function, such as the Unadjusted Langevin Algorithm (ULA) and the Unadjusted Hamiltonian Monte Carlo (UHMC) [27,28]. We let τ denote the index of the MCMC iterations (as opposed to t , which refers to the diffusion timestep).

For ULA we use the kernel

$$k_t(x^{\tau+1} | x^\tau) = \mathcal{N}(x^{\tau+1}; x^\tau + \delta_t \nabla_x \log p_\theta(x^\tau, t), 2\delta_t I),$$

at diffusion step t , where $x^0 = x_t$, δ_t is the step size, and the chain is iterated for L_t steps.

For UHMC we augment the state with momenta $v^\tau \sim \mathcal{N}(0, M_t)$ (diagonal mass M_t), and propose a new state by applying ℓ_t leapfrog steps of size ε_t under the same score field. Writing

$$(x^{\tau+1}, v^{\tau+1}) = \text{LF}_{\ell_t, \varepsilon_t}(x^\tau, v^\tau),$$

for the leapfrog map (which in our case is defined with the score $\nabla_x \log p_\theta(x, t)$ as potential gradient), the kernel for x is given implicitly by this deterministic proposal after marginalizing out v . In practice, we may perform L_t such proposals per diffusion step t .

These methods are called unadjusted since, as L_t grows, the Markov chains converges to the target distribution only in the limit of infinitesimal step sizes. By adding, for instance, a Metropolis–Hastings (MH) correction step, we can sample with larger step sizes and still converge to the target distribution [23,24]. With the correction, we sample a candidate $\hat{x} \sim k_t(\cdot | x^\tau)$ and accept it as the new iterate with probability

$$\alpha = \min\left(1, \frac{p_\theta(\hat{x}, t)}{p_\theta(x^\tau, t)} \frac{k_t(x^\tau | \hat{x})}{k_t(\hat{x} | x^\tau)}\right), \tag{4}$$

so that $x^{\tau+1} = \hat{x}$ with probability α , and $x^{\tau+1} = x^\tau$ otherwise. With the correction, ULA becomes LA (often referred to as MALA) [29], and UHMC becomes the standard HMC algorithm [22,30].

The model p_θ appears only through a ratio, so a normalized density is not required. When p_θ is parameterized as an EBM (see (1)), the ratio simplifies to

$$\frac{p_\theta(\hat{x}, t)}{p_\theta(x^\tau, t)} = \exp\left(\frac{1}{\sigma_t} (E_\theta(x^\tau, t) - E_\theta(\hat{x}, t))\right), \quad (5)$$

which allows us to directly evaluate the MH acceptance probability, making it straightforward to construct an adjusted MCMC sampler. This offers a key advantage over the score parameterization, where only an approximation of the score is accessible which cannot directly be used to compute the probability ratio needed in MH.

2.5. Sampling from Composed Models

Composed sampling is a powerful feature of diffusion models that enables sampling from new target distributions by combining multiple pre-trained models. Rather than retraining a model for every new task or data combination, one can reuse existing components. This flexibility is especially appealing in large-scale settings, where retraining is often prohibitively expensive.

The most common form of composition is *guidance* [10], where the goal is to sample from a distribution conditioned on a class label y ,

$$q(x_0 | y) \propto q(x_0)q(y | x_0). \quad (6)$$

This is implemented by modifying the score function at each diffusion step as

$$\nabla_x \log p_\theta(x_t, t) + \lambda \nabla_x \log p_\phi(y | x_t, t), \quad (7)$$

where p_θ is an unconditional diffusion model and p_ϕ is a classifier predicting class y . Hyperparameter λ controls the strength of the conditioning. We refer to this approach as *classifier-full guidance*. Other variants include reconstruction guidance [31,32] and classifier-free guidance [17].

More generally, ref. [18] explores a range of composition types beyond guidance, including *products*, *negations*, and *mixtures*. A product composition—of which guidance can be seen as a special case—is defined as

$$q^\Pi(x_0) \propto \prod_i q^i(x_0), \quad (8)$$

and leads to the composed model at diffusion step t ,

$$p_\theta^\Pi(x_t, t) \propto \prod_i p_{\theta_i}^i(x_t, t) = \exp\left(-\frac{1}{\sigma_t} \sum_i E_{\theta_i}^i(x_t, t)\right). \quad (9)$$

This distribution is then used as the target in MCMC sampling, resulting in improved sampling performance.

Importantly, the factorization in (8) only strictly holds at $t = 0$; at intermediate diffusion steps, the composed model $p_\theta^\Pi(x_t, t)$ does not generally correspond to the true marginal of any product data distribution; refer to Du et al. [18]. This becomes problematic when relying solely on the reverse process, which assumes access to a valid score function for the true intermediate marginals. However, this construction remains valid and effective from the perspective of *annealed MCMC* [33], where the overall sampling procedure is interpreted as a chain targeting a sequence of gradually evolving distributions. From this viewpoint, the intermediate distributions $p_\theta^\Pi(x_t, t)$ are treated as design choices that guide the chain toward the final target $q^\Pi(x_0)$, and asymptotic correctness is still preserved. In

practice, since diffusion models are trained using denoising score matching, the sampling process converges to a denoised version of $q^\Pi(x_0)$, which can be made arbitrarily close to the true distribution by construction.

Note that for models using a score-based parameterization, a pseudo-score for this type of composition is equal to $-\frac{1}{\sigma_t} \sum_i \epsilon_{\theta_t}^i(x_t, t)$.

3. MCMC Correction Step for Score Parameterization

We propose combining the energy parameterization properties with the performance and practical accessibility of score parameterization. Instead of using energy parameterization and computing the score by differentiation, we take the complementary approach: using score parameterization and computing the change in (pseudo-)energy by integrating the score.

3.1. Pseudo-Energy Difference and MH-like Correction

This section describes how MCMC acceptance probabilities can be approximated given only access to a score function. The Metropolis–Hastings acceptance probability in (4) depends on the *relative* probability of a proposed state \hat{x} compared to the current state x^τ . Since the proposal kernel $k_t(\cdot | \cdot)$ is straightforward to evaluate, the acceptance ratio reduces to comparing unnormalized target densities.

For an energy-based model, this ratio depends only on the difference and—because E_θ is a scalar potential—this difference can be written as a path-independent line integral,

$$E_\theta(x^\tau, t) - E_\theta(\hat{x}, t) = - \int_{\mathcal{C}} \nabla_{\zeta} E_\theta(\zeta, t) \cdot d\zeta, \tag{10}$$

for any differentiable curve \mathcal{C} connecting x^τ and \hat{x} . This provides a natural interpretation of MH in terms of integrating the model’s score field along a path.

For score-parameterized diffusion models, the situation is reversed: the model provides a score field $\epsilon_\theta(x, t)$ but no explicit energy function. Motivated by the fact that denoising score matching trains ϵ_θ to approximate the gradient of the log-density under Gaussian perturbations, we define an analogous *pseudo-energy difference*

$$\Delta \tilde{E}_{\mathcal{C}}(x \rightarrow x', t) = - \int_{\mathcal{C}} \epsilon_\theta(\zeta, t) \cdot d\zeta, \tag{11}$$

where \mathcal{C} is any smooth curve connecting x and x' . This construction can be viewed as integrating the vector field ϵ_θ along a path, thereby approximating the change in an underlying scalar potential—if such a potential existed.

Based on this quantity, we define an MH-like acceptance probability

$$\alpha_{\text{MH-like}} = \min\left(1, \exp\left[\frac{1}{\sigma_t} \Delta \tilde{E}_{\mathcal{C}}(x \rightarrow x', t)\right] \frac{k_t(x | x')}{k_t(x' | x)}\right). \tag{12}$$

If $\epsilon_\theta(x, t) = \nabla_x F_\theta(x, t)$ for some potential F_θ , then $\Delta \tilde{E}_{\mathcal{C}}$ equals $F_\theta(x', t) - F_\theta(x, t)$ for all curves \mathcal{C} , and (12) exactly recovers the true MH acceptance probability by the fundamental theorem of line integrals.

In practice, ϵ_θ is not perfectly conservative and the pseudo-energy difference depends on the chosen path \mathcal{C} . We therefore consider two practical choices: (i) a straight line between x and x' and (ii) a curved trajectory following the leapfrog path of an HMC proposal. The latter allows reusing score evaluations already computed during proposal generation, yielding higher numerical accuracy without additional model evaluations.

In both cases, the integral in (11) is approximated using the trapezoidal rule, with the number of line segments treated as a hyperparameter n . This requires score evalua-

tions at internal points of \mathcal{C} , but avoids differentiating the model; by contrast, the energy parameterization evaluates E_θ at the endpoints but must differentiate it to obtain the score.

An overview of the full sampling procedure is given in Algorithm 1. At each diffusion step, an optional reverse update is followed by an MCMC refinement targeting the intermediate distribution. This structure aligns naturally with the annealed MCMC framework, where both the reverse update and the MCMC kernel act as design choices guiding the chain toward the final distribution. Including the reverse step typically improves sample quality [18].

Algorithm 1 Annealed MCMC with MH-like correction

Require: Score function $\epsilon_\theta(\cdot, t)$; reverse-diffusion schedule $(\alpha_t, \beta_t, \sigma_t)_{t=1}^T$; kernel family $\{k_t(\cdot | \cdot)\}_{t=1}^T$ (LA or HMC) with kernel hyperparameters; MCMC steps per diffusion step $\{L_t\}$; integration mode $\text{mode} \in \{\text{line}, \text{curve}\}$; line segments n (if line); per-leapfrog subsegments m (if curve).

- 1: Sample initial $x_T \sim \mathcal{N}(0, I)$.
 - 2: **for** $t = T$ **to** 1 **do**
 - 3: *(Optional) Reverse step:* update x_{t-1} from x_t using the standard reverse-diffusion update at time t
 - 4: **if** $t > 1$ **then**
 - 5: $t' \leftarrow t - 1$ ▷ all MCMC quantities live at time t'
 - 6: Initialize MCMC: $x^0 \leftarrow x_{t'}$.
 - 7: **for** $\tau = 1$ **to** L_t **do**
 - 8: Propose candidate $\hat{x} \sim k_{t'}(\cdot | x^{\tau-1})$ ▷ LA or HMC kernel
 - 9: $\Delta \tilde{E}_C \leftarrow \text{EstimatePseudoEnergy}(x^{\tau-1}, \hat{x}, t'; \text{mode}, n, m)$ ▷ Algorithm 2 or Algorithm 3
 - 10: Compute acceptance
 - $$\alpha_{\text{MH-like}} = \min\left(1, \exp\left(\frac{1}{\sigma_{t'}} \Delta \tilde{E}_C\right) \cdot \frac{k_{t'}(x^{\tau-1} | \hat{x})}{k_{t'}(\hat{x} | x^{\tau-1})}\right)$$
 - 11: With probability $\alpha_{\text{MH-like}}$: $x^\tau \leftarrow \hat{x}$; otherwise $x^\tau \leftarrow x^{\tau-1}$.
 - 12: **end for**
 - 13: Set $x_{t'} \leftarrow x^{L_t}$.
 - 14: **end if**
 - 15: **end for**
 - 16: **return** x_0
-

Algorithm 2 Estimate Pseudo-Energy (straight-line path)

Require: Current x , candidate \hat{x} , time t , number of segments $n \geq 2$; access to $\epsilon_\theta(\cdot, t)$. Note that $\epsilon_\theta(x, t)$ and $\epsilon_\theta(\hat{x}, t)$ are available from cache (Algorithm 1). The routine may reuse them instead of re-evaluating

- 1: $\Delta r \leftarrow \frac{1}{n-1} (\hat{x} - x)$
- 2: $r_{\text{prev}} \leftarrow x$; $g_{\text{prev}} \leftarrow \epsilon_\theta(x, t)$; $\Delta \tilde{E}_C \leftarrow 0$
- 3: **for** $j = 1$ **to** $n - 1$ **do**
- 4: $r_j \leftarrow r_{\text{prev}} + \Delta r$
- 5: $g_j \leftarrow \epsilon_\theta(r_j, t)$ ▷ reuse cached value at endpoints
- 6: $\Delta \tilde{E}_C \leftarrow \Delta \tilde{E}_C - \frac{1}{2} (g_{\text{prev}} + g_j) \cdot \Delta r$
- 7: $r_{\text{prev}} \leftarrow r_j$; $g_{\text{prev}} \leftarrow g_j$
- 8: **end for**
- 9: **return** $\Delta \tilde{E}_C$

Complexity (extra score evaluations). If endpoint scores are reused, this routine performs $n - 2$ new $\epsilon_\theta(\cdot, t)$ calls; otherwise n calls.

Algorithm 3 Estimate Pseudo-Energy (HMC-curved path)

Require: Current x , candidate \hat{x} , time t ; leapfrog trajectory $\{x^{(i)}\}_{i=0}^{\ell_t}$ with $x^{(0)}=x$, $x^{(\ell_t)}=\hat{x}$; subsegments per leapfrog step $m \geq 2$; access to $\epsilon_\theta(\cdot, t)$. Note that $\{\epsilon_\theta(x^{(i)}, t)\}_{i=0}^{\ell_t}$ are available from cache (Algorithm 1) and may be reused instead of re-evaluating.

```

1:  $\Delta\tilde{E}_C \leftarrow 0$ 
2: for  $i = 1$  to  $\ell_t$  do ▷ integrate along each LF segment  $x^{(i-1)} \rightarrow x^{(i)}$ 
3:    $r_{\text{prev}} \leftarrow x^{(i-1)}$ ;  $g_{\text{prev}} \leftarrow \epsilon_\theta(x^{(i-1)}, t)$ 
4:    $\Delta r_i \leftarrow \frac{1}{m-1}(x^{(i)} - x^{(i-1)})$ 
5:   for  $j = 1$  to  $m - 1$  do
6:      $r_{i,j} \leftarrow r_{\text{prev}} + \Delta r_i$ 
7:      $g_{i,j} \leftarrow \epsilon_\theta(r_{i,j}, t)$  ▷ reuse cached leapfrog scores if available
8:      $\Delta\tilde{E}_C \leftarrow \Delta\tilde{E}_C - \frac{1}{2}(g_{\text{prev}} + g_{i,j}) \cdot (r_{i,j} - r_{\text{prev}})$ 
9:      $r_{\text{prev}} \leftarrow r_{i,j}$ ;  $g_{\text{prev}} \leftarrow g_{i,j}$ 
10:  end for
11: end for
12: return  $\Delta\tilde{E}_C$ 

```

Complexity. With cached leapfrog scores, this routine performs $\ell_t(m-2)$ new $\epsilon_\theta(\cdot, t)$ calls (one per interior subsegment); without caching it performs $\ell_t(m-1)$ calls. Setting $m = 2$ reduces to a trapezoid rule that uses only leapfrog endpoints.

3.2. MH Correction for Composition Models

The pseudo-energy difference for compositions can be derived based on their specific definitions. Our proposed method applies directly to product compositions. We calculate a pseudo-energy difference, corresponding to $E_\theta^\Pi(x^\tau, t) - E_\theta^\Pi(\hat{x}, t)$ for an EBM (defined in (9)), as

$$- \int_{\mathcal{C}} \sum_i \epsilon_{\theta_i}^i(\zeta, t) \cdot d\zeta. \quad (13)$$

Guidance is a specific case of product composition, where the pseudo-score is composed of two terms according to (7): the unconditional diffusion model $\epsilon_\theta(x_t, t)$ and the score of a classifier $p_\varphi(y | x_t, t)$. Since $p_\varphi(y | x_t, t)$ can be evaluated directly, only the pseudo-energy difference for $\epsilon_\theta(x_t, t)$ requires computation using the line integral in (13).

The pseudo-energy difference for a negation composition (as defined in [18]) can be computed analogously to products, as negations follow a similar additive structure in their pseudo-scores.

Mixture compositions (as defined in [18]), on the other hand, cannot be expressed as a pseudo-energy difference, since mixtures do not naturally conform to an additive structure analogous to products or negations. However, mixtures can be addressed by first sampling a component distribution according to the mixture definition and then generating a sample from that distribution. The MH correction can subsequently be applied to this sampled distribution, providing a seamless way to handle mixture compositions within our framework.

This generalization allows our method to support advanced use cases such as classifier guidance, multi-modal fusion, and spatially structured prompts, without requiring retraining or access to energy-based models.

4. Results

In this section, we present an empirical evaluation of our MH-like correction method (all code used to produce the results in this section is publicly available at https://github.com/FraunhoferChalmersCentre/mcmc_corr_score_diffusion, accessed on 4 January 2025), examining both the accuracy of the pseudo-energy differences and the quality of the generated samples. The experiments are designed to span a spectrum of difficulty: from

controlled, low-dimensional setups where models can be trained from scratch and analytical solutions are available to more realistic high-dimensional scenarios involving pre-trained models. Our two primary objectives are (1) to compare our proposed approach against a true energy parameterization when available and (2) to assess the sampling improvements achieved over the standard reverse process when augmented with MCMC steps.

The experiments in Sections 4.1 and 4.2 and the first part of Section 4.3 involve training diffusion models using both energy and score parameterizations. The score parameterization follows a noise prediction model, $\epsilon_\theta(x_t, t)$, while the energy parameterization defines an energy function as $E_\theta(x_t, t) = \|x_t - s_\theta(x_t, t)\|_2^2$, as in [18]. We use identical network architectures for ϵ_θ and s_θ . Both models are trained with the standard diffusion loss [3], with the energy model's score function obtained through explicit differentiation.

The later experiments utilize only pre-trained score-based diffusion models, as pre-trained energy-based models are unavailable for direct comparison. We evaluate both unadjusted and MH-corrected versions of Langevin and Hamiltonian Monte Carlo, comparing them against the standard reverse process, which serves as the baseline.

For the MH-like correction, we examine two types of integration paths: a straight line between x^τ and \hat{x} (Algorithm 2) and the trajectory defined by the HMC leapfrog steps (Algorithm 3). Both approaches rely on a trapezoidal rule where the number of intermediate points is treated as a hyperparameter.

Complete training details, hyperparameter settings, and implementation specifics are deferred to Appendix A.

4.1. Evaluating Pseudo-Energy Differences

To evaluate the accuracy of pseudo-energy differences, we conducted experiments on a synthetic 2D dataset, generated from a bivariate Gaussian distribution to allow access to analytical solutions, and a higher-dimensional dataset, MNIST [34]. For each experiment, we trained 10 independent score models and 10 independent energy models from scratch. For evaluation, we sampled 2000 pairs of points (x_t^1, x_t^2) independently via the forward process at various diffusion steps t . These pairs were used to compute the score-based pseudo-energy difference ($\Delta \tilde{E}_{\text{score}}$) and the explicit EBM energy difference (ΔE_{EBM}), as well as the analytical difference (ΔE_{true}) when available. The pseudo-energy difference was computed along a straight-line path connecting the two points, using five discretization points for the numerical integration.

2D Gaussian: For the 2D Gaussian dataset, the relative error metric is defined as $|\Delta E_{\text{pred}} - \Delta E_{\text{true}}| / |\Delta E_{\text{true}}|$, where ΔE_{pred} corresponds to either the explicitly predicted difference from the energy model (ΔE_{EBM}) or the pseudo-energy difference from the score model ($\Delta \tilde{E}_{\text{score}}$), and ΔE_{true} is the analytical energy difference. The median relative error was calculated across all sampled pairs for each trained model, and the mean and standard deviation of this metric were computed across the 10 models. Interestingly, the score model achieved a lower relative error (0.071 ± 0.005) compared to the energy model (0.084 ± 0.004), demonstrating that the line-integral approximation aligns slightly better with the true energy differences in this setting.

MNIST: For the MNIST dataset, where analytical energy differences are unavailable, we used a symmetric relative discrepancy metric defined as $2|\Delta \tilde{E}_{\text{score}} - \Delta E_{\text{EBM}}| / (|\Delta \tilde{E}_{\text{score}}| + |\Delta E_{\text{EBM}}|)$. The median relative discrepancy was calculated across all sampled pairs for each trained model, and the mean and standard deviation were computed across the 10 models. This yielded a mean symmetric relative discrepancy of 0.030 ± 0.002 , indicating that the pseudo-energy differences predicted by the score models align closely with the explicit EBM predictions, even in this higher-dimensional setting.

4.2. Two-Dimensional Composition

To investigate the effectiveness of our MH-like correction in a controlled yet expressive setting, we replicate the 2D composition experiment introduced by Du et al. [18] using their publicly available codebase (https://github.com/yilundu/reduce_reuse_recycle, accessed on 4 January 2025) as a foundation. Only minor modifications are necessary, ensuring a faithful reproduction of their setup. Apart from differences in evaluation metrics, our setup is identical to theirs.

A 2D density pair is composed via multiplication into a complex distribution, as in (9): a Gaussian mixture with eight modes in a circle and a uniform distribution covering two of the modes. For a visual representation of the two individual distributions and their resulting product distribution together with samples from the reverse diffusion and HMC corrected samples, see Figure 1. The baseline reverse diffusion process uses $T = 100$ steps. In the MCMC variants, following [18], we omit the optional reverse step for a fair comparison. MCMC sampling runs for $L_t = 10$ at each t , with (U-)HMC using three leapfrog steps per MCMC step.

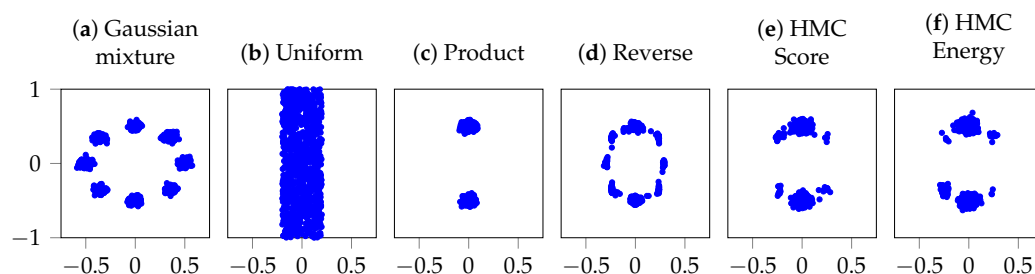


Figure 1. Samples from: (a,b) the component distributions: a Gaussian mixture and a uniform distribution, (c) the true product distribution, (d) a standard score-parameterized reverse process, (e,f) HMC sampling using score and energy parameterization, respectively.

We evaluate performance using three metrics: (1) negative log-likelihood (NLL), (2) a Gaussian mixture model (GMM) comparison, and (3) the Wasserstein-2 distance (W_2). The suffixes in Table 1 refer to the choice of integration path and the number of evaluation points: “L” denotes the straight-line path from Algorithm 2, while “C” denotes the curve path from Algorithm 3. The number indicates trapezoidal evaluation points (m or n).

Table 1 reports performance metrics averaged over 10 independent trials. In each trial, we train the diffusion models from scratch and sample 2000 points using different MCMC methods. The corrected sampling methods consistently outperform the unadjusted ones. HMC variants yield the best results across all metrics. Score- and energy-parameterized samplers show similar NLL and GMM performance, while HMC with score parameterization achieves a substantially lower W_2 . Performance also saturates with as few as three integration points.

In addition to performance, we measure runtime and memory consumption. These results are reported separately in Table 2. The experiment is implemented in JAX (v0.4.30) [35] and run on a desktop computer equipped with an NVIDIA GeForce RTX 3060 GPU. Score-based parameterization is more than twice as memory-efficient as energy-based parameterization, and—except for LA with $m = 10$ —also faster for corresponding MCMC methods. Notably, the HMC curve variant is significantly faster. Although score-based corrections require more model evaluations, they do not necessarily incur higher runtime or memory costs.

As discussed by Du et al. [18], directly adding score functions does not yield a valid product composition, which explains why the reverse sampler performs poorly in this setting. This apparent failure is expected and highlights the motivation for annealed MCMC: by treating intermediate distributions as design choices that guide the chain toward the

target, annealing achieves improved results while preserving asymptotic correctness. This accounts for the large gap between the reverse method and the annealed MCMC variants reported in Table 1.

Table 1. Performance metrics (NLL, GMM, and W_2) for the 2D composition experiment. The downward arrow (\downarrow) indicates that lower values are better. Values represent the mean with the standard deviation in parentheses over 10 independent trials. Bold numbers indicate the best performance in each metric for both Energy and Score samplers.

	Sampler	NLL \downarrow	GMM \downarrow	W_2 \downarrow
Energy	Reverse	8.22 (0.21)	27.01 (1.34)	5.81 (0.19)
	U-LA	7.52 (0.22)	14.61 (1.35)	4.19 (0.45)
	LA	6.50 (0.30)	14.66 (1.46)	4.24 (0.55)
	U-HMC	5.72 (0.18)	6.53 (0.91)	4.19 (1.25)
	HMC	4.09 (0.14)	3.33 (0.65)	4.12 (1.44)
Score	Reverse	8.15 (0.24)	26.88 (1.20)	5.80 (0.20)
	U-LA	7.57 (0.12)	14.99 (0.62)	4.44 (0.63)
	LA-3L	6.45 (0.20)	14.28 (1.07)	4.03 (0.52)
	LA-5L	6.61 (0.17)	15.19 (0.92)	4.22 (0.46)
	LA-10L	6.53 (0.17)	14.75 (0.91)	4.20 (0.51)
	U-HMC	5.77 (0.12)	6.90 (0.71)	3.39 (0.77)
	HMC-3L	4.29 (0.13)	3.72 (0.61)	2.92 (1.02)
	HMC-5L	4.07 (0.13)	3.08 (0.69)	2.68 (1.20)
	HMC-10L	4.07 (0.14)	3.17 (0.56)	2.87 (0.89)
HMC-2C	4.07 (0.12)	3.06 (0.54)	2.94 (0.90)	

Table 2. Runtime (seconds) and peak memory usage (MiB) for the 2D composition experiment.

	Sampler	Time	Memory
Energy	Reverse	0.22(0.00)	5252
	U-LA	1.54(0.01)	5252
	LA	9.13(0.08)	5252
	U-HMC	2.36(0.05)	5254
	HMC	21.02(0.06)	5256
Score	Reverse	0.11 (0.00)	2178
	U-LA	0.93 (0.01)	2180
	LA-3L	4.77 (0.13)	2180
	LA-5L	7.60 (0.02)	2180
	LA-10L	10.53 (0.15)	2180
	U-HMC	1.19 (0.01)	2180
	HMC-3L	7.08 (0.02)	2180
	HMC-5L	9.56 (0.01)	2180
	HMC-10L	10.64 (0.01)	2180
HMC-2C	1.48 (0.04)	2180	

4.3. Guided Diffusion

We evaluate our proposed sampling methods for guided diffusion on the CIFAR-100 [36] and ImageNet [37] datasets. The sampling process is based on a score function defined in (7). For both datasets, the marginal score, $\nabla_x \log q(x_t)$, is estimated using an unconditional diffusion model parameterized by a UNet architecture. For the guidance model, we use classifier-full guidance, training a time-dependent classifier to predict class labels across all diffusion steps, $p_\phi(y | x_t, t)$. This classifier shares its architecture with the encoder part of the UNet used for the diffusion model and is extended with a dense output layer. The guidance scale is set to $\lambda = 20.0$ across all experiments, which is a commonly used default in classifier-

guided diffusion models. This choice is also consistent with the settings provided in the public implementation of [18] (https://github.com/yilundu/reduce_reuse_recycle, accessed on 4 January 2025). Sampling is based on the standard reverse process with $T = 1000$, and additional MCMC steps are incorporated to refine the generated samples. The *Reverse* baseline reported in the tables corresponds to the standard classifier-full guided reverse diffusion process without any additional MCMC refinement steps.

To quantify generation quality, we use three evaluation metrics: the Fréchet Inception Distance (FID) [38], which compares the distribution of generated and real images; classification accuracy, based on a separate pre-trained classifier applied to generated samples; and, for ImageNet, an additional top 5 accuracy metric.

CIFAR-100: For CIFAR-100, we trained the diffusion models from scratch using the same UNet architecture and training settings as in [3], which were originally designed for CIFAR-10 [36]. The MCMC samplers add $L_t = 2$ or 6 extra MCMC steps at each diffusion step t for (U-)HMC and (U-)LA, respectively, with (U-)HMC using three leapfrog steps per MCMC step, following the configuration used by Du et al. [18] in their guided diffusion experiments.

For this experiment, we used denser meshes in the trapezoidal rule compared to the 2D setting. For HMC, we followed the curve defined by the leapfrog steps (Algorithm 3) and set $n = 3$, i.e., the leapfrog points plus an additional midpoint evaluation. For LA, we followed the straight line path (Algorithm 2) with $m = 10$ evenly spaced points, corresponding to eight additional evaluations per step.

Recognizing the impact of the step length on MCMC methods, we followed Du et al. [18] by parameterizing it as a function of the beta-schedule. For Langevin dynamics, the stepsize is given by $\delta_t = a\beta_t^b$, and for HMC, the leapfrog stepsize is denoted $\varepsilon_t = a\beta_t^b$. To ensure a fair comparison across all baselines and proposed methods, we conducted a seeded parameter search over a predefined range of (a, b) values shared across all variants. Parameters were selected by minimizing FID on a validation subset, with classification accuracy monitored as a secondary metric. The search revealed that unadjusted samplers (U-LA, U-HMC) achieved their best performance with comparatively smaller step sizes. In contrast, the adjusted variants (both energy and score) attained their lowest FID at larger step sizes. For these settings, the average empirical acceptance rates were approximately 5% for HMC and 2% for LA, consistent across both energy and score parameterizations.

The results are shown in Table 3. Average accuracy is obtained using a separate classifier trained exclusively on noise-free pairs (x_0, y) , following the VGG-13-BN architecture [39]. The table shows a general trend of improvement over the baseline reverse process when additional MCMC steps are added. In particular, the MH-corrected samplers LA and HMC show significant improvements in FID scores, which are arguably the more important metric for image generation.

Comparing the score and energy parameterizations, their performances share similar characteristics. Interestingly, the reverse process favors the score parameterization, supporting the claim that this less restricted approach better models the score function. However, the energy parameterization sees larger improvements from the added MCMC steps. This indicates, perhaps, that direct energy estimation provides a better correction step compared to our method of approximating the pseudo-energy difference from ε_θ . Although the energy-based method performs slightly better in this setting, our MH-corrected sampling methods achieve comparable improvements without requiring an energy model.

ImageNet: For ImageNet, training diffusion models from scratch requires substantial computational resources, so we relied on widely used pre-trained models. Score-based models are publicly available through the OpenAI GitHub repository (<https://github.com/openai/guided-diffusion>, accessed on 4 January 2025), as provided by Dhariwal

and Nichol [10]. However, to our knowledge, there are currently no equivalent publicly available, pre-trained energy-parameterized models at this scale. This limitation highlights the practical value of our approach: by operating directly on pre-trained score models, our method enables MCMC-based refinement even in large-scale settings where explicit energy parameterizations are not readily available. Given the high computational demands of large-scale diffusion models, we focused solely on evaluating HMC and compared it to the reverse process. The HMC sampler adds $L_t = 2$ MCMC steps per diffusion step t , with each step consisting of three leapfrog steps, following the setup of the guided diffusion experiment in Du et al. [18]. For the trapezoidal rule, we used the curve-based integration from Algorithm 3 with $m = 4$. The step length parameterization and tuning follow the same procedure as in CIFAR-100.

Table 3. Accuracy and FID score for classifier-full guidance on CIFAR-100. The metrics are based on 50k generated samples for each sampling method with both energy and score models. The upward (\uparrow) and downward (\downarrow) arrows indicate that higher and lower values are better, respectively. Bold numbers indicate the best performance in each metric for both Energy and Score samplers.

	Sampler	Accuracy [%] \uparrow	FID \downarrow
Energy	Reverse	72.6	33.4
	U-LA	87.3	24.6
	LA	80.0	12.7
	U-HMC	87.2	25.4
	HMC	84.9	12.4
Score	Reverse	74.2	31.8
	U-LA	82.9	25.9
	LA-10L	75.2	15.5
	U-HMC	79.0	28.6
	HMC-3C	75.8	13.3

The results can be seen in Table 4. Accuracy metrics are computed using a pre-trained RegNetX-8.0GF [40] classifier. The reverse process and HMC perform very similarly in average accuracy, but our method shows a slight improvement in top 5 average accuracy. HMC obtains a significantly better FID score.

Table 4. Average accuracy, top 5 accuracy, and FID score for classifier-full guidance on ImageNet. The metrics are based on 50k generated samples for both sampling methods with score parameterizations. The upward (\uparrow) and downward (\downarrow) arrows indicate that higher and lower values are better, respectively. Bold numbers indicate the best performance in each metric.

	Sampler	Acc [%] \uparrow	Acc-5 [%] \uparrow	FID \downarrow
Score	Reverse	50.0	83.9	14.5
	HMC-4C	49.9	85.1	11.6

4.4. Image Tapestry

We conduct a final, image tapestry experiment, similar to [18] and based on their code (https://github.com/yilundu/reduce_reuse_recycle, accessed on 4 January 2025), with only minor modifications to incorporate our MH-like correction. The goal is to generate a coherent image composed of spatially localized content, each region conditioned on different prompts. This task involves both classifier-free guidance and model composition—specifically the combination of multiple overlapping text-to-image diffusion models, each responsible for a portion of the scene.

We use a pre-trained DeepFloyd-IF model (<https://huggingface.co/DeepFloyd/IF-I-XL-v1.0>, accessed on 4 January 2025) as the base diffusion model. To refine the generated

samples, we apply Langevin dynamics with our MH-like correction. For each diffusion step ($T = 100$), we include 15 additional Langevin steps. The pseudo-energy difference is approximated via line integration using three additional evaluation points per step. We set the classifier-free guidance scale to $\lambda = 20.0$.

The resulting image is presented in Figure 2a, which showcases the generated tapestry with different regions displaying distinct visual content. Figure 2b provides a schematic overview of the used prompts and their spatial layout. In total, nine content regions are specified: four located in the corners of the image, each with unique prompts, and five overlapping in the center, all guided by the same prompt to create a unified visual theme.

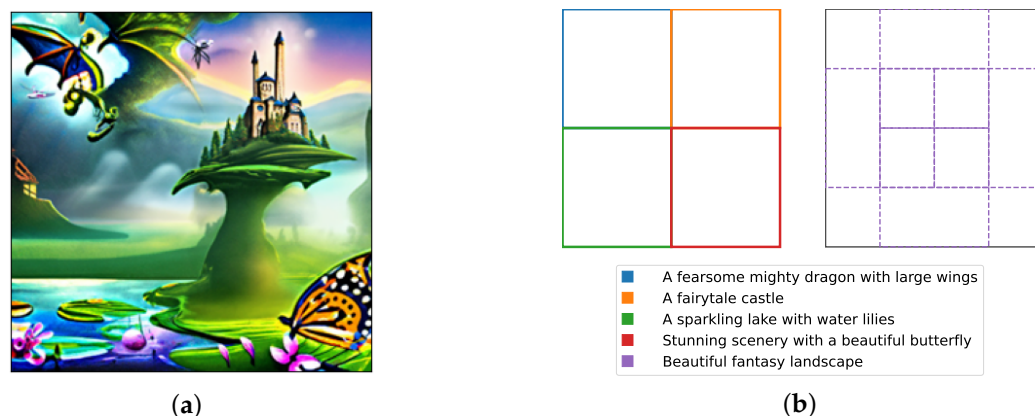


Figure 2. In (a), the generated tapestry image is shown with different content at various locations. In (b), the specified content and their positions are illustrated.

5. Discussion

The choice between score and energy parameterizations remains an intriguing and nuanced topic within diffusion-based generative modeling. In this work, we provided additional empirical evidence suggesting that the score parameterization performs better in the standard reverse process.

At the same time, we showed that performance gains often attributed to the energy parameterization can, in fact, be recovered within a score-based framework. This is achieved by approximating pseudo-energy differences using a line integral of the model's noise predictions. Notably, this allows us to incorporate MH-like correction steps into a variety of MCMC samplers—without the need to explicitly train an energy-based model—achieving relative improvements in sample quality of a similar magnitude. This represents a crucial practical benefit, as score-based models are fundamentally easier to train and scale to high-dimensional datasets compared to energy-based parameterizations.

A particularly interesting observation is that using a curve composed only of model evaluations from the HMC sampler appears to perform on par with using a straight-line path. This suggests that the proposed correction comes at virtually no additional computational cost in this case. However, it is worth noting that in higher-dimensional settings, additional intermediate points along the integration path may be required to maintain accuracy, which could increase the computational burden. This challenge might be addressed through more efficient numerical integration techniques, or by working in a lower-dimensional latent space, as in latent diffusion models. One persistent drawback of the energy parameterization is that it always requires an explicit gradient computation to recover the score function.

Another important consideration is computational cost. As we observed in the experiments, adding MCMC updates on top of the reverse diffusion process introduces a clear overhead in terms of additional score evaluations. Our approach is therefore not

intended as a replacement for accelerated solvers, e.g., PF-ODE, which aim to minimize the number of function evaluations. Instead, we view MH-like corrections as an orthogonal contribution: they can be applied on top of any diffusion model and any sampler. In this sense, our method is complementary to existing work, targeting improved sample quality while remaining agnostic to the choice of backbone or solver.

Our empirical step length parameter search on CIFAR-100 highlights a qualitative difference between adjusted and unadjusted samplers. The unadjusted variants (U-LA, U-HMC), which lack a rejection mechanism, achieved their best performance only within a relatively narrow range of smaller step sizes. Outside this regime, image quality deteriorated rapidly during sampling. In contrast, the adjusted variants consistently produced semantically coherent samples across a broader portion of the explored parameter range. Empirically, the rejection step appears to mitigate the destabilizing effect of overly aggressive proposals, effectively filtering out unfavorable transitions. For the best performing configurations, the adjusted samplers exhibited relatively low average acceptance rates (around 2 or 5%). Notably, both energy-based and score-based adjusted samplers achieved peak performance at similar acceptance levels, indicating comparable dynamical behavior between explicit and pseudo-energy formulations.

In line with prior work such as Du et al. [18], our experimental design deliberately focused on widely used baselines rather than the strongest available backbones. This choice allowed us to isolate the effect of the MH-like correction without confounding factors from architectural or solver improvements. While this comes at the cost of absolute FID values that are below the state of the art, the relative improvements we observe consistently demonstrate the benefit of our approach. We expect that applying our method to stronger architectures will yield proportionally similar gains, but leave this as an exciting direction for future work.

One limitation of the score parameterization is that the learned vector field is not guaranteed to be conservative, and therefore our line-integral construction does not, in general, yield a true Metropolis–Hastings correction. Consequently, the resulting acceptance rule should not be interpreted as guaranteeing detailed balance with respect to the target distribution. In practice, the samples are generated by the sequential application of the reverse diffusion updates together with MCMC steps and the practically motivated MH-like acceptance rule. The resulting distribution should therefore be understood as the implicit distribution induced by this procedure. Importantly, our construction recovers a true Metropolis–Hastings correction in the special case where the score is conservative, thereby aligning with the energy-based formulation. More broadly, recent work by Horvat and Pfister [41] highlights that score representations may admit non-conservative components without necessarily undermining their effectiveness for generative modeling. This perspective aligns with the empirical success of score-based generative models that operate without explicitly modeling an energy function. Consistent with this perspective, our MH-like correction mechanism—though built on a generally non-conservative field—yields consistent improvements across the evaluated datasets and samplers when applied to the reverse process.

To further assess the practical impact of the non-conservative nature of the learned score field, we conducted a series of empirical sanity checks (detailed in Appendices B.1 and B.2). First, to examine potential path dependence, we compared pseudo-energy differences obtained from straight-line and curved integration paths. On both trained MNIST models and the higher-dimensional CIFAR-100 model, the observed discrepancies were small relative to the magnitude of the pseudo-energy differences. This indicates limited path sensitivity at the proposal scale relevant for our MCMC refinement steps. Furthermore, we evaluated the numerical stability of the line-integral approximation via a mesh convergence analysis on

CIFAR-100 (detailed in Appendix B.3). We observed rapid convergence of the discretization scheme: the median symmetric relative discrepancy falls below 1% when using $n \geq 10$ integration points compared to a higher-resolution reference. Together, these empirical results suggest that the learned score behaves approximately conservatively on the local scales relevant for MCMC proposals, as small variations in the integration path yield similar pseudo-energy differences.

Still, the lack of theoretical guarantees may explain the slightly superior performance of the energy parameterization observed in the CIFAR-100 experiment. Developing better techniques for estimating pseudo-energy differences from score-based models—without requiring an explicitly trained energy function—thus remains a highly relevant and promising direction for future research.

6. Conclusions

We introduced a method for extending the reverse diffusion process with MCMC sampling based on an MH-like correction step computed from the score function. This approach enables improved sampling for composed diffusion models without requiring an energy-based parameterization.

While previous work, Du et al. [18], demonstrated the benefits of MH correction under an energy parameterization, our method instead defines a pseudo-energy difference derived from the score, estimated via numerical integration. This allows us to apply MH-like corrections in the score-based setting—by far the most common in practice—and thereby make use of existing pre-trained diffusion models for composition tasks.

Our method can reuse intermediate evaluations from samplers such as HMC to compute the correction with little to no additional cost. In general, the accuracy of the MH-like correction depends on the numerical integration of the score, which may require more intermediate points as the dimensionality increases. While this can introduce some overhead, energy-based methods incur their own costs, such as differentiating the energy function. In practice, our corrected score-based samplers consistently yield relative improvements of a similar magnitude to energy-based methods across a range of tasks, making them a practical alternative in settings where score-based models are already available.

Overall, our work extends the applicability of corrected MCMC sampling to the broad class of score-based diffusion models and opens the door to more flexible and modular composition of generative models.

Author Contributions: Conceptualization, A.S. and M.Ö.; Methodology, A.S. and J.L.; Software, A.S. and J.L.; Validation, A.S., J.L., M.Ö. and L.S.; Formal analysis, A.S., J.L. and M.Ö.; Investigation, A.S. and J.L.; Resources, M.J. and L.S.; Writing—original draft preparation, A.S. and J.L.; Writing—review and editing, A.S., J.L., M.Ö., L.S. and M.J.; Visualization, A.S. and J.L.; Supervision, M.Ö., L.S. and M.J.; Project administration, L.S. and M.J.; Funding acquisition, L.S. and M.J. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council (grant no. 2022-06725).

Data Availability Statement: All datasets used in this study are either publicly available from the sources cited in the manuscript, or can be fully reproduced from the code provided.

Acknowledgments: During the preparation of this manuscript, the authors used ChatGPT (GPT-5 and later versions) for assistance in improving the clarity and readability of the manuscript. The authors have reviewed and edited all AI-assisted content and take full responsibility for the final version of the manuscript.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

MCMC	Markov Chain Monte Carlo
MH	Metropolis–Hastings
LA	Langevin Algorithm
HMC	Hamiltonian Monte Carlo
U-LA	Unadjusted Langevin Algorithm
U-HMC	Unadjusted Hamiltonian Monte Carlo
EBM	Energy-Based Models
DSM	Denosing Score Matching

Appendix A. Experimental Details

Here we provide more details about our different conducted experiments: Evaluating pseudo-energy differences, two-dimensional composition, guided diffusion, and image tapestry.

The earlier experiments were conducted on a machine equipped with an NVIDIA GeForce RTX 3060, while the later experiments were run on a computing cluster with NVIDIA A100 Tensor Core GPUs.

Appendix A.1. Evaluating Pseudo-Energy Difference

All models in this section were trained with the Adam optimizer using a learning rate of 10^{-3} , together with a StepLR scheduler with step size 1 and decay factor 0.99.

2D Gaussian: We generated samples from a bivariate Gaussian distribution with mean $\mu = (2, 0)^\top$ and covariance $\Sigma = 0.1I$, where I is the identity matrix.

The diffusion models use $T = 100$ timesteps, with the noise schedule β_t following the cosine schedule proposed in [42].

We used the same neural network architectures as the base for both the score and energy models. It is a residual network consisting of a linear layer (dim $2 \rightarrow 128$) followed by four blocks, and concluding with a linear layer (dim $128 \rightarrow 2$). Within each block, the input x passes through a normalization layer, a SiLU activation, and a linear layer (dim $128 \rightarrow 256$). Subsequently, it is supplemented with an embedded t (dim 32) that has undergone a linear layer transformation (dim $32 \rightarrow 256$). The resulting sum passes through a SiLU activation and is further processed by a linear layer (dim $256 \rightarrow 256$). After that, another SiLU activation is applied, followed by a final linear layer (dim $256 \rightarrow 128$). The output of this linear layer is then added to the original input x within the block. The embedding of t is also learnable.

MNIST: The diffusion models use $T = 1000$ timesteps, with the noise schedule β_t following the cosine schedule. For score parameterization, we trained a UNet-based architecture adapted to 28×28 grayscale images. The network uses a time embedding of dimension 112, implemented as sinusoidal position embeddings followed by two fully connected layers with GELU activations. The model begins with a 1×1 convolution mapping the input image to dimension 28 and proceeds through three down-sampling stages, a middle block, and three up-sampling stages. Each down-/upsampling stage consists of two residual blocks with time conditioning, a linear or full attention layer, and either a strided convolution (downsampling) or nearest-neighbor upsampling with convolution (upsampling). The middle block contains two ResNet blocks and one full attention layer. Skip connections are applied between corresponding encoder and decoder layers, following the standard UNet design. The output stage concatenates the upsampled features with the

initial projection and applies a residual block followed by a 1×1 convolution to map back to the image space.

For energy parameterization, the architecture is identical, but the output head is replaced with an energy function whose gradient defines the score.

Appendix A.2. Two-Dimensional Composition

The composed distribution is defined by a product of two components, a Gaussian mixture and a uniform distribution with non-zero values on

$$\square = \{x \in \mathbb{R}^2 : -s_i \leq x_i \leq s_i, i = 1, 2\}, \tag{A1}$$

where s_1 and s_2 are equal to 0.2 and 1.0, respectively. The eight modes of the Gaussian mixture are evenly distributed on a circle with a radius of 0.5 at the angles $\frac{\pi}{4}i$ for $i = 0, \dots, 7$, respectively. The covariance matrix at each mode is $0.03^2 \cdot I$, where I is the identity matrix.

We use the same network architecture setup for score and energy as in the 2D Gaussian case (see Appendix A.1).

The metric log-likelihood is ill-defined as we may generate samples where the true distribution has no support (due to the uniform distribution). We address this problem by expanding the definition set of the uniform distribution and redistributing one percent of the probability mass into this extended region. The whole set is defined as (A1) except $s_1 = s_2 = 1.1$. Note that the 99 percent probability mass remains inside the original definition set \square .

The parameter β_t follows the cosine schedule. For (U-)HMC, the damping coefficient is set to 0.5, the mass diagonal matrix has all diagonal elements equal to 1, and the stepsize for each t is 0.03. For (U-)LA, the stepsize for each t is set to 0.001.

Appendix A.3. Guided Diffusion for CIFAR-100

The parameter β_t has a linear schedule as originally proposed in [3]. For (U-)HMC, the damping coefficient is equal to 0.9 and the diagonal elements in the mass matrix are equal to β_t for each t . The values of the stepsize parameters a and b were determined through a simple parameter search for the different MCMC methods and they can be found in Table A1. This was performed for both the score and energy parameterizations, where the stepsize is defined as $\delta_t = a\beta_t^b$.

Table A1. The values of the stepsize parameters a and b obtained from a random parameter search for the different MCMC methods for both score and energy parameterization in the CIFAR-100 experiment, where the stepsize is defined as $\delta_t = a\beta_t^b$.

		MCMC	Stepsize Parameters	
			a	b
Energy		U-LA	9.22	1.40
		LA	9.84	0.83
		U-HMC	0.26	1.53
		HMC	9.33	1.48
Score		U-LA	1.96	1.04
		LA	9.84	0.83
		U-HMC	0.26	1.53
		HMC	4.03	1.34

To complement the quantitative results in the main text, we report in Table A2 the theoretical number of function evaluations (NFE) per generated sample for CIFAR-100. We separate forward passes (FPs) and backward passes (BPs), since energy parameterizations require BPs for score evaluations whereas score parameterizations only require FPs. Counts are computed per algorithm step and correspond to the maximum evaluations assuming all proposals are accepted in MH.

Table A2. Theoretical number of function evaluations (NFE) per generated sample on CIFAR-100. We report forward passes (FPs) and backward passes (BPs) separately. Counts are computed per algorithm step, independent of accept/reject outcomes.

	Sampler	NFE (FPs)	NFE (BPs)
Energy	Reverse	0	1000
	U-LA	0	6994
	LA	6993	7993
	U-HMC	0	6994
	HMC	2997	7993
	Score	Reverse	1000
	U-LA	6994	0
	LA-10L	55,945	0
	U-HMC	6994	0
	HMC-3C	13,982	0

To provide a qualitative comparison, we show representative generated samples in Figure A1. Each pair of images (a) and (b) is conditioned on the same class label and initialized from the same noise realization x_T , but generated with different samplers.

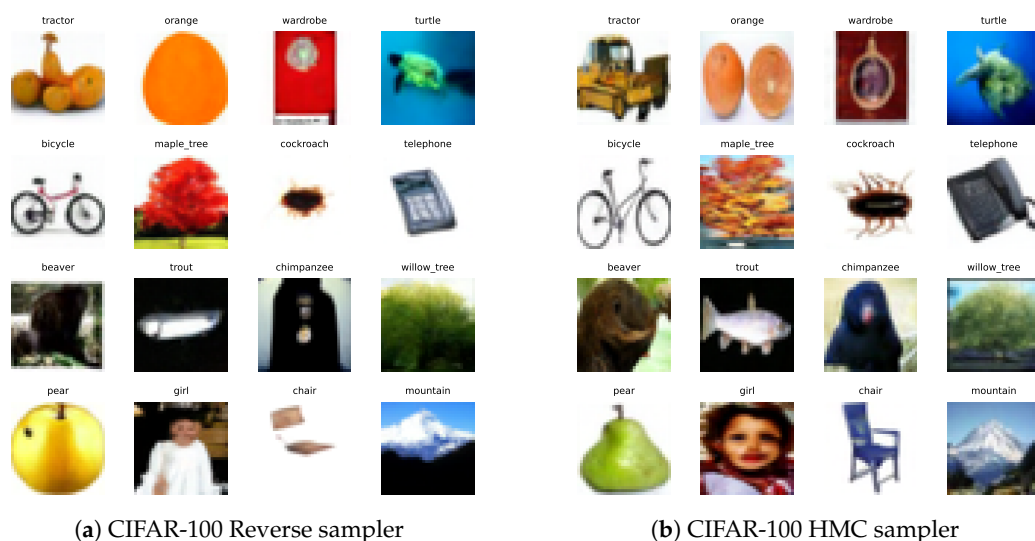


Figure A1. Generated CIFAR-100 samples using (a) reverse and (b) HMC.

Appendix A.4. Guided Diffusion for ImageNet

Again, the parameter β_t follows a linear schedule. The hyperparameters for the HMC include a damping coefficient set to 0.9, with the diagonal elements of the mass matrix being equal to β_t for each t . The stepsize parameters for HMC, obtained from a simple parameter search, are $a = 1.87$ and $b = 1.51$.

The ImageNet dataset (<https://image-net.org/>, accessed on 4 January 2025) used to compute the FID score is available for free to researchers for non-commercial use.

For the ImageNet experiment with classifier guidance, only score parameterizations were considered. Table A3 reports the corresponding NFEs per generated sample, again

separating forward (FPs) and backward (BPs) passes. Here we compare the baseline reverse sampler with our MH-like corrected HMC variant (HMC-6C).

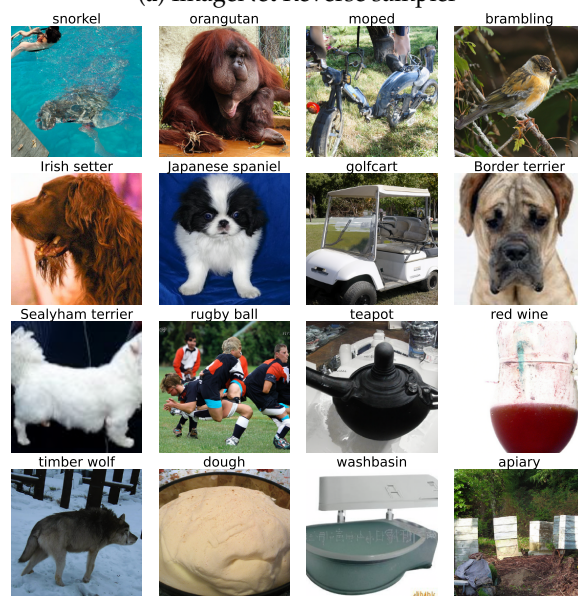
For ImageNet, representative samples are shown in Figure A2. As in the CIFAR-100 case, the pairs are conditioned on the same class label and share the same starting noise realization x_T , enabling a direct visual comparison between the samplers.

Table A3. Theoretical number of function evaluations (NFE) per generated sample on ImageNet. We report forward passes (FPs) and backward passes (BPs) separately. Counts are computed per algorithm step, independent of accept/reject outcomes.

	Sampler	NFE (FP)	NFE (BP)
Score	Reverse	1000	0
	HMC-4C	19,981	0



(a) ImageNet Reverse sampler



(b) ImageNet HMC sampler

Figure A2. Generated ImageNet samples using (a) reverse and (b) HMC. All labels correspond to standard ImageNet class names.

Appendix A.5. Image Tapestry

A cosine schedule is used for the parameter β_t . The stepsize parameters in this case is simply $a = 1$ and $b = 1$, i.e., $\delta_t = \beta_t$.

Appendix B. Additional Experiments

Appendix B.1. Sanity Check: Line vs. Curve Path on MNIST

We conducted an additional experiment on the MNIST dataset to further assess the effect of path choice when estimating pseudo-energy differences. The setup followed the same spirit as in Appendix A.1. Specifically, we trained 10 independent score models and sampled 2000 states from the forward process at various timesteps t . From each state, we performed a single HMC proposal consisting of 3 leapfrog steps with stepsize $\varepsilon_t = a\beta_t^b$, where $a = 4.03$ and $b = 1.34$ (as in the CIFAR-100 experiments). This produced a proposed state \hat{x}_t for each sampled pair (x_t, \hat{x}_t) .

We then computed the pseudo-energy difference between x_t and \hat{x}_t using two different paths: (i) a straight line between the points, with $n = 10$ (Algorithm 2), and (ii) the curved path defined by the leapfrog steps themselves, with $m = 3$ (Algorithm 3).

The discrepancy between the two estimates was measured using the symmetric relative error

$$\frac{2|\Delta E_{\text{line}} - \Delta E_{\text{curve}}|}{|\Delta E_{\text{line}}| + |\Delta E_{\text{curve}}|}. \quad (\text{A2})$$

For each model, we computed the median relative error across the 2000 sampled pairs. We then reported the mean and standard deviation across the 10 trained models, yielding a value of 0.022 ± 0.002 . This small error indicates that line and curve integration paths produce very similar pseudo-energy differences, supporting the interpretation that the score behaves approximately conservatively on the local scales relevant for MCMC proposals.

Appendix B.2. Sanity Check: Line vs. Curve Path on CIFAR-100

To assess path sensitivity in a higher-dimensional setting, we repeat the same analysis using the trained CIFAR-100 score model employed in our main experiments.

We sample 2000 states from the forward diffusion process at various timesteps t . From each state, we perform a single HMC proposal consisting of 3 leapfrog steps with step size $\varepsilon_t = a\beta_t^b$, using the parameters selected in Table A1. This produces pairs (x_t, \hat{x}_t) analogous to the MNIST experiment.

We compute the pseudo-energy difference between x_t and \hat{x}_t using both a straight-line path and the curved leapfrog path, and measure the symmetric relative discrepancy defined in Equation (A2).

Across the 2000 sampled pairs, we obtain a median discrepancy of 0.038. While slightly larger than the MNIST result (0.022 ± 0.002), the discrepancy remains small relative to the magnitude of the pseudo-energy differences. This suggests that the score continues to behave approximately conservatively on the local scales relevant for our MCMC proposals, resulting in limited path sensitivity even in the higher-dimensional CIFAR-100 setting.

Appendix B.3. Integration Mesh Convergence on CIFAR-100

To assess numerical stability of the line-integral approximation, we compare pseudo-energy differences computed with varying numbers of discretization points n against a high-resolution reference with $n_{\text{ref}} = 40$.

We measure the symmetric relative discrepancy

$$r_{\text{mesh}}(n) = \frac{2 \left| \Delta E_{\text{line}}^{(n)} - \Delta E_{\text{line}}^{(40)} \right|}{\left| \Delta E_{\text{line}}^{(n)} \right| + \left| \Delta E_{\text{line}}^{(40)} \right|}. \quad (\text{A3})$$

Across 2000 sampled pairs, we obtain median discrepancies $r_{\text{mesh}}(5) = 0.032$, $r_{\text{mesh}}(10) = 0.0064$, and $r_{\text{mesh}}(20) = 8.2 \times 10^{-4}$. These results indicate rapid convergence of the numerical integration scheme. In particular, for $n \geq 10$, the discretization error is already below 1% relative to the high-resolution reference and becomes negligible for $n \geq 20$.

References

1. Brock, A.; Donahue, J.; Simonyan, K. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In Proceedings of the International Conference on Learning Representations (ICLR), New Orleans, LA, USA, 6–9 May 2019.
2. Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language models are few-shot learners. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 1877–1901.
3. Ho, J.; Jain, A.; Abbeel, P. Denoising diffusion probabilistic models. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 6840–6851.
4. Li, Y.; Choi, D.; Chung, J.; Kushman, N.; Schrittwieser, J.; Leblond, R.; Eccles, T.; Keeling, J.; Gimeno, F.; Dal Lago, A.; et al. Competition-level code generation with alphacode. *Science* **2022**, *378*, 1092–1097. [[CrossRef](#)] [[PubMed](#)]
5. Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E.L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. Photorealistic text-to-image diffusion models with deep language understanding. *Adv. Neural Inf. Process. Syst.* **2022**, *35*, 36479–36494.
6. Güngör, A.; Dar, S.U.; Öztürk, Ş.; Korkmaz, Y.; Bedel, H.A.; Elmas, G.; Ozbey, M.; Çukur, T. Adaptive diffusion priors for accelerated MRI reconstruction. *Med. Image Anal.* **2023**, *88*, 102872. [[CrossRef](#)]
7. Wynn, J.; Turmukhambetov, D. DiffusioNeRF: Regularizing neural radiance fields with denoising diffusion models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Vancouver, BC, Canada, 18–22 June 2023; pp. 4180–4189.
8. Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In Proceedings of the International Conference on Machine Learning (ICML), Lille, France, 6–11 July 2015; pp. 2256–2265.
9. Song, Y.; Ermon, S. Generative modeling by estimating gradients of the data distribution. In Proceedings of the Advances in Neural Information Processing Systems 32, Vancouver, BC, Canada, 8–14 December 2019.
10. Dhariwal, P.; Nichol, A. Diffusion models beat GANs on image synthesis. *Adv. Neural Inf. Process. Syst.* **2021**, *34*, 8780–8794.
11. Lüdke, D.; Biloš, M.; Shchur, O.; Lienen, M.; Günnemann, S. Add and Thin: Diffusion for Temporal Point Processes. In Proceedings of the Thirty-Seventh Conference on Neural Information Processing Systems (NeurIPS), New Orleans, LA, USA, 10–16 December 2023.
12. Wang, K.; Xu, Z.; Zhou, Y.; Zang, Z.; Darrell, T.; Liu, Z.; You, Y. Neural Network Diffusion. *arXiv* **2024**, arXiv:2402.13144. [[CrossRef](#)]
13. Jacobs, R.A.; Jordan, M.I.; Nowlan, S.J.; Hinton, G.E. Adaptive mixtures of local experts. *Neural Comput.* **1991**, *3*, 79–87. [[CrossRef](#)] [[PubMed](#)]
14. Hinton, G.E. Training products of experts by minimizing contrastive divergence. *Neural Comput.* **2002**, *14*, 1771–1800. [[CrossRef](#)]
15. Mayraz, G.; Hinton, G.E. Recognizing hand-written digits using hierarchical products of experts. In Proceedings of the 13th International Conference on Neural Information Processing Systems (NeurIPS), Denver, CO, USA, 27 November–2 December 2000.
16. Liu, N.; Li, S.; Du, Y.; Torralba, A.; Tenenbaum, J.B. Compositional visual generation with composable diffusion models. In *Proceedings of the Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, 23–27 October 2022*; Proceedings, Part XXVII; Springer: Cham, Switzerland, 2022; pp. 423–439.
17. Ho, J.; Salimans, T. Classifier-Free Diffusion Guidance. In Proceedings of the NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications, Virtual Event, 14 December 2021.
18. Du, Y.; Durkan, C.; Strudel, R.; Tenenbaum, J.B.; Dieleman, S.; Fergus, R.; Sohl-Dickstein, J.; Doucet, A.; Grathwohl, W.S. Reduce, Reuse, Recycle: Compositional Generation with Energy-Based Diffusion Models and MCMC. In *Proceedings of the 40th International Conference on Machine Learning (ICML), Honolulu, HI, USA, 23–29 July 2023*; Proceedings of Machine Learning Research: Cambridge, MA, USA, 2023; Volume 202, pp. 8489–8510.
19. Aghajanyan, A.; Yu, L.; Conneau, A.; Hsu, W.N.; Hambardzumyan, K.; Zhang, S.; Roller, S.; Goyal, N.; Levy, O.; Zettlemoyer, L. Scaling laws for generative mixed-modal language models. In *Proceedings of the 40th International Conference on Machine Learning (ICML), Honolulu, HI, USA, 23–29 July 2023*; Proceedings of Machine Learning Research: Cambridge, MA, USA, 2023; Volume 202.

20. Song, Y.; Sohl-Dickstein, J.; Kingma, D.P.; Kumar, A.; Ermon, S.; Poole, B. Score-Based Generative Modeling through Stochastic Differential Equations. In Proceedings of the International Conference on Learning Representations (ICLR), Virtual Event, 3–7 May 2021.
21. Roberts, G.O.; Stramer, O. Langevin Diffusions and Metropolis–Hastings Algorithms. *Methodol. Comput. Appl. Probab.* **2002**, *4*, 337–357. [[CrossRef](#)]
22. Duane, S.; Kennedy, A.; Pendleton, B.J.; Roweth, D. Hybrid Monte Carlo. *Phys. Lett. B* **1987**, *195*, 216–222. [[CrossRef](#)]
23. Metropolis, N.; Rosenbluth, A.W.; Rosenbluth, M.N.; Teller, A.H. Equation of State Calculations by Fast Computing Machines. *J. Chem. Phys.* **1953**, *21*, 1087–1092. [[CrossRef](#)]
24. Hastings, W.K. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika* **1970**, *57*, 97–109. [[CrossRef](#)]
25. Salimans, T.; Ho, J. Should EBMs model the energy or the score? In Proceedings of the Energy-Based Models Workshop at ICLR 2021, Virtual Event, 7 May 2021.
26. LeCun, Y.; Chopra, S.; Hadsell, R.; Ranzato, M.; Huang, F. A tutorial on energy-based learning. In *Predicting Structured Data*; MIT Press: Cambridge, MA, USA, 2006; Volume 1.
27. Roberts, G.O.; Tweedie, R.L. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli* **1996**, *2*, 341–363. [[CrossRef](#)]
28. Geffner, T.; Domke, J. MCMC variational inference via uncorrected Hamiltonian annealing. *Adv. Neural Inf. Process. Syst.* **2021**, *34*, 639–651.
29. Besag, J. Comments on “Representations of knowledge in complex systems” by U. Grenander and M. I. Miller. *J. R. Stat. Soc. Ser. B (Methodol.)* **1994**, *56*, 591–592.
30. Neal, R.M.; Diggle, P.; Fienberg, S. *Bayesian Learning for Neural Networks*; Lecture Notes in Statistics Series; Springer: New York, NY, USA, 1996; Volume 118.
31. Chung, H.; Kim, J.; Mccann, M.T.; Klasky, M.L.; Ye, J.C. Diffusion Posterior Sampling for General Noisy Inverse Problems. In Proceedings of the The Eleventh International Conference on Learning Representations (ICLR), Kigali, Rwanda, 1–5 May 2023.
32. Ho, J.; Salimans, T.; Gritsenko, A.; Chan, W.; Norouzi, M.; Fleet, D.J. Video Diffusion Models. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*; Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., Oh, A., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2022; Volume 35, pp. 8633–8646.
33. Neal, R.M. Annealed importance sampling. *Stat. Comput.* **2001**, *11*, 125–139. [[CrossRef](#)]
34. Deng, L. The MNIST Database of Handwritten Digit Images for Machine Learning Research [Best of the Web]. *IEEE Signal Process. Mag.* **2012**, *29*, 141–142. [[CrossRef](#)]
35. Bradbury, J.; Frostig, R.; Hawkins, P.; Johnson, M.J.; Leary, C.; Maclaurin, D.; Necula, G.; Paszke, A.; VanderPlas, J.; Wanderman-Milne, S.; et al. JAX: Composable Transformations of Python+NumPy Programs. Version 0.4.30. 2018. Available online: <https://github.com/google/jax> (accessed on 4 January 2025).
36. Krizhevsky, A.; Hinton, G. *Learning Multiple Layers of Features from Tiny Images*; Technical Report ; University of Toronto: Toronto, ON, Canada, 2009.
37. Deng, J.; Dong, W.; Socher, R.; Li, L.J.; Li, K.; Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Miami, FL, USA, 20–25 June 2009; pp. 248–255.
38. Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, 6626–6637.
39. Simonyan, K.; Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv* **2014**, arXiv:1409.1556.
40. Radosavovic, I.; Kosaraju, R.P.; Girshick, R.; He, K.; Dollár, P. Designing network design spaces. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 13–19 June 2020; pp. 10428–10436.
41. Horvat, C.; Pfister, J.P. On gauge freedom, conservativity and intrinsic dimensionality estimation in diffusion models. *arXiv* **2024**, arXiv:2402.03845. [[CrossRef](#)]
42. Nichol, A.Q.; Dhariwal, P. Improved denoising diffusion probabilistic models. In *Proceedings of the International Conference on Machine Learning (ICML), Virtual Event, 18–24 July 2021*; Proceedings of Machine Learning Research: Cambridge, MA, USA, 2021; Volume 139, pp. 8162–8171.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.