

BEYOND AUTOREGRESSION: FAST LLMs VIA SELF-DISTILLATION THROUGH TIME

Anonymous authors

Paper under double-blind review

ABSTRACT

Autoregressive (AR) Large Language Models (LLMs) have demonstrated significant success across numerous tasks. However, the AR modeling paradigm presents certain limitations; for instance, contemporary autoregressive LLMs are trained to generate one token at a time, which can result in noticeable latency. Recent advances have indicated that search and repeated sampling can enhance performance in various applications, such as theorem proving, code generation, and alignment, by utilizing greater computational resources during inference. In this study, we demonstrate that diffusion language models are capable of generating at least 32 tokens simultaneously, while exceeding the performance of AR models in text quality and on the LAMBADA natural language understanding benchmark. This outcome is achieved through a novel distillation method for discrete diffusion models, which reduces the number of inference steps by a factor of 32-64. Practically, at the 1.3B parameters scale, diffusion models, even without caching, can generate tokens at a rate that is up to 8 times faster than AR models employing KV-caching, and we anticipate further improvements with the inclusion of caching. Moreover, we demonstrate the efficacy of our approach for diffusion language models with up to 860M parameters.

1 INTRODUCTION

In recent years, autoregressive (AR) large language models (LLM) have exceeded expectations (Vaswani et al., 2017; Devlin et al., 2018; Radford et al., 2019; Brown et al., 2020b; Kaplan et al., 2020; Raffel et al., 2020; Fedus et al., 2022; Hoffmann et al., 2022; Chowdhery et al., 2023; Google, 2023; Touvron et al., 2023). Importantly, many breakthroughs in coding (Chen et al., 2021), mathematics, and reasoning (Trinh et al., 2024b;a; Romera-Paredes et al., 2024; Hosseini et al., 2024; Wang et al., 2024) were achieved based on decoding large amounts of completions from a base LLM.

Importantly, the benefits of repeated sampling can be so significant that it is often more efficient to use a smaller, faster model rather than a larger, slower one.

More generally, one can improve the performance of a fixed model by scaling up computational resources at inference time (Madaan et al., 2023; Yao et al., 2023; Snell et al., 2024; Wu et al., 2024; Chen et al., 2024; Brown et al., 2024; Goyal et al., 2024), a phenomenon that was previously observed for games (Campbell et al., 2002; Silver et al., 2016; Lerer et al., 2019; Brown et al., 2020a; Jones, 2021). Hence, when tackling reasoning tasks, a major bottleneck is the latency of the model. In this work, we improve the decoding speed of LLMs by moving away from AR modeling. We build on recent breakthroughs in discrete diffusion (Lou et al., 2023; Sahoo et al., 2024; Shi et al., 2024; Ou et al., 2024). Our approach can generate text up to 8 times faster than AR models that use KV caching (Pope et al., 2022). Diffusion models are typically trained to maximize the evidence lower bound (ELBO), which does not consider the desired number of inference steps. Hence, vanilla diffusion models typically require thousands of decoding steps. Fortunately, it is possible to drasti-

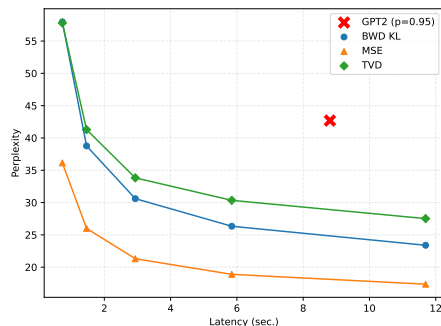



Figure 1: **Perplexity versus latency.** The diffusion models (169M) use 16, 32, 64, 128 and 256 decoding step.

054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107



(a) Accuracy of the correct last word decoded from our model. Distillation with KLD loss leads the student model to outperform the teacher in terms of accuracy on LAMBADA. (b) Perplexity of the last word. The KLD preserves performance best, and even when the student is trained to sample with 16 instead of 1024 steps, the student still matches AR baselines.

Figure 2: **Performance on LAMBADA after multiple rounds of SDTT** with different distillation losses. We pre-train with the masked diffusion language modeling objective (MDLM) (Sahoo et al., 2024) and distill with 7 rounds of SDTT. Note that a single word in the LAMBADA data set often consists of multiple tokens. **We greedily decode all tokens a single forward pass for the diffusion models and decode autoregressively for the AR models.**

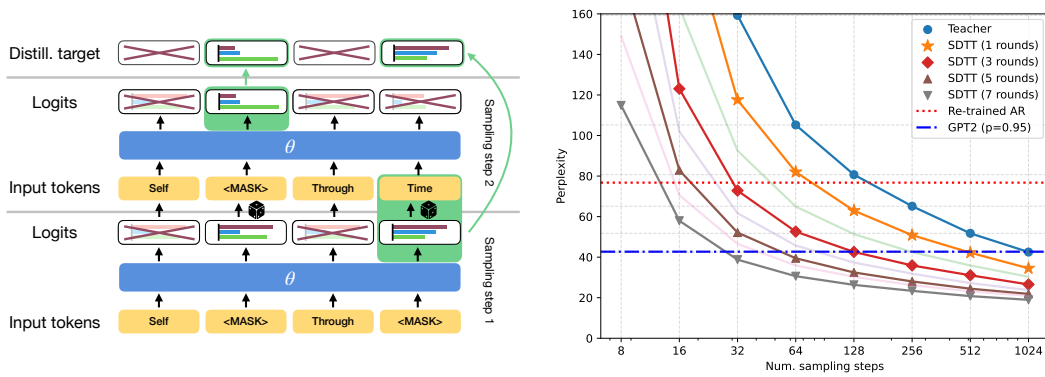
cally reduce the inference costs of *continuous* diffusion models via distillation (Luhman & Luhman, 2021; Salimans & Ho, 2022). Most distillation methods rely on deterministic mappings from noise to data, such as DDIM (Song et al., 2022). The deterministic mappings can be efficiently learned by a student diffusion model to sample in fewer steps. We hypothesize that such deterministic map cannot exist for the diffusion language models studied in this work. Indeed, those models always initialize the denoising process with a sequence of masked token, hence a deterministic algorithm can only generate a single sample. As such, we devise a distillation method that does not depend on deterministic maps. This is a significant finding because faster decoding mechanisms allow exploring a larger search space in applications that require search, planning, and reranking. In summary, our core contributions are as follows:

- We introduce *Self-Distillation Through Time* (SDTT), which allows generating **at least 32** tokens at a time, while achieving better perplexity than GPT-2 with nucleus sampling **for conditional and unconditional generation**. Unlike many distillation methods for continuous diffusion models, SDTT does not rely on deterministic mappings such as DDIM (Song et al., 2022). SDTT is very simple and easy to implement.
- We show that SDTT can generate tokens up to 8 times faster than AR models that use KV caching, **for models with 1.3B parameters, in 16 decoding steps**. Importantly, the discrete diffusion model does not rely on activation caching, suggesting that there is potential for even greater efficiency gains. **The latency gains for smaller models are even greater.**
- We demonstrate the effectiveness of SDTT for models with up to 860M parameters. To the best of our knowledge, this represents the largest publicly available discrete diffusion language model.
- **We evaluate the distilled students on LAMBADA (Paperno et al., 2016) and 6 multiple-choice questions benchmarks from Gao et al. (2021). We find that SDTT preserves the downstream performance of the teacher.**

2 BACKGROUND

2.1 MASKED DIFFUSION LANGUAGE MODELING

We follow the notation of Sahoo et al. (2024) to introduce masked diffusion language modeling (MDLM). Language modeling can be framed as the sequential prediction task of discrete



(a) The distillation targets are the log probabilities that lead to a token being denoised, concatenated with log probabilities of the last step for tokens that remain masked.

(b) SDTT on small models trained for 1M steps. Successive lines correspond to additional SDTT rounds. SDTT can outperform the teacher and GPT-2 with nucleus sampling.

Figure 3: **SDTT**. In figure (a), we illustrate how we prepare the distillation targets. In figure (b), we display the generative perplexity of samples after distillation.

tokens (x_i) coming from a vocabulary $\mathcal{X} = \mathbb{Z}^{<N} = \{0, \dots, N-1\}$ that can take N possible discrete values. A language model would predict sequences of length L , which can be defined as the sequences of x_i 's originating from $\mathcal{X}^L = \{\mathbf{x}^{(i)} = (x_0^{(i)}, \dots, x_{L-1}^{(i)})\}_{i \in \mathbb{Z}^{<\kappa}}$. Let $\mathcal{D} := \{\mathbf{x}^{(0)}, \dots, \mathbf{x}^{(K-1)} : \mathbf{x}^{(i)} \in \mathcal{X}^L\}$ denote the training set. The goal of language modeling is to sample from the unknown distribution $p_0 : \mathcal{X}^L \rightarrow [0, 1]$ that generated the samples in \mathcal{D} .

Similarly to continuous diffusion, we sample from an approximation of p_0 by learning to denoise corrupted examples. One can sample from the model through ancestral sampling, starting from a stationary distribution. The stationary distribution of Sahoo et al. (2024) is such that all tokens of the sentence are replaced with a special MASK token like the MASK token used for pre-training BERT models. However, a key difference between BERT and MDLM is that MDLM is trained on sequences with varying levels of corruption, while BERT uses a fixed ratio.

Discrete absorbing diffusion process MDLM defines a forward process to corrupt data and a backward process to learn to recover data. MDLM uses a continuous-time formulation, with the data distribution denoted as p_0 and the stationary noise distribution as $p_1 = \pi$. The forward process linearly interpolates between the one-hot distribution defined by the original document \mathbf{x} and the stationary distribution π , which places all mass on the MASK token. Mathematically,

$$q(\mathbf{z}_t | \mathbf{x}) := \text{Cat}(\mathbf{z}_t; \alpha_t \mathbf{x} + (1 - \alpha_t) \pi), \quad (1)$$

where the noise injection schedule is defined by α_t , for $t \in [0, 1]$. The constraints on α_t are that $\alpha_t \in [0, 1]$, α_t should be a strictly decreasing function of t , and $\alpha_0 \approx 1, \alpha_1 \approx 0$. The forward process is called absorbing because once a token is assigned to a MASK token, it cannot be reverted to a real token.

We can derive the analytical form of the reverse process $q(\mathbf{z}_s | \mathbf{z}_t, \mathbf{x})$, with $t > s$ and $\alpha_{t|s} = \frac{\alpha_t}{\alpha_s}$ as

$$q(\mathbf{z}_s | \mathbf{z}_t, \mathbf{x}) = \text{Cat} \left(\mathbf{z}_s; \frac{[\alpha_{t|s} \mathbf{z}_t + (1 - \alpha_{t|s}) \mathbf{1} \pi^\top \mathbf{z}_t] \odot [\alpha_s \mathbf{x} + (1 - \alpha_s) \pi]}{\alpha_t \mathbf{z}_t^\top \mathbf{x} + (1 - \alpha_t) \mathbf{z}_t^\top \pi} \right). \quad (2)$$

Objective and parameterization To generate new samples, we can simulate the reverse process from eq. (2). Since the ground-truth sample \mathbf{x} is unknown, Sahoo et al. (2024) learn an approximation \mathbf{x}_θ using a neural network with parameters θ . Sahoo et al. (2024) then use \mathbf{x}_θ instead of \mathbf{x} to simulate the reverse process. The sampling distribution is denoted as $p_\theta(\mathbf{z}_s | \mathbf{z}_t) := q(\mathbf{z}_s | \mathbf{z}_t, \mathbf{x}_\theta(t))$. Sahoo et al. (2024) optimize θ using a continuous version of the negative evidence lower bound (ELBO) of Sohl-Dickstein et al. (2015a). Previous research has shown that continuous-time objectives optimize the data likelihood better (Kingma et al., 2023). Due to the definition of the absorbing

Algorithm 1 Computing the *Self-Distillation Through Time* targets $\tilde{\mathbf{x}}_{\theta}^{\text{teacher}}(\mathbf{z}_t, t, m/k)$

1: **Inputs:** Noisy tensor $\mathbf{x}_t \in \mathbb{R}^{N \times L}$, Starting sampling time $t_{\text{start}} \in [0, 1]^N$, Number of sampling steps $m/k \geq 2$, such that $m/k \in \mathbb{N}_+$, Sampling step size $\Delta \in (0, 1)$, Mask token index $M \in \mathbb{N}$, Minimal sampling time ϵ .

2: **Output:** Distillation targets $\tilde{\mathbf{x}}_{\theta}^{\text{teacher}}(\mathbf{z}_t, t, m/k)$

3:

4: $\text{target} \leftarrow \text{zeros}(N, L, K)$ ▷ Allocate empty tensor for $\tilde{\mathbf{x}}_{\theta}^{\text{teacher}}(\mathbf{z}_t, t, m/k)$

5: $\mathbf{z} \leftarrow \mathbf{x}_t$

6: **for** $i = 0, \dots, m/k - 1$ **do**

7: $t_{\text{curr}} \leftarrow \max(t_{\text{start}} - i \cdot \Delta, \epsilon)$ ▷ Sampling step for the current time

8: $\mathbf{z}_{\text{new}}, \ell_{\text{teacher}} \leftarrow \text{reverse_sample}(\mathbf{z}, t_{\text{curr}}, \Delta)$ ▷ Updated \mathbf{z} & log-probabilities

$\mathbf{x}_{\theta}(\mathbf{z}, t_{\text{curr}})$

9: $\mathbf{U} = \mathbf{z}_{\text{new}} \neq \mathbf{z}$ ▷ Create mask U of tokens that were denoised

10: $\text{target}[\mathbf{U}] \leftarrow \ell_{\text{teacher}}[\mathbf{U}]$ ▷ Extract log-probs for the denoised tokens

11: $\mathbf{z} \leftarrow \mathbf{z}_{\text{new}}$ ▷ Update \mathbf{z} for the next iteration

12: **end for**

13: $\text{target}[\mathbf{z} == M] = \ell_{\text{teacher}}[\mathbf{z} == M]$ ▷ Use log-probs of the last denoising step for masked tokens

14: **return** target ▷ Target log-probs for all masked tokens in \mathbf{x}_t

diffusion process, the ELBO simplifies to a weighted cross-entropy loss between the ground-truth \mathbf{x} and the model predictions \mathbf{x}_{θ} :

$$\mathcal{L}_{\text{NELBO}}^{\infty} = \mathbb{E}_q \int_{t=0}^{t=1} \frac{\alpha'_t}{1 - \alpha_t} \log \langle \mathbf{x}_{\theta}(\mathbf{z}_t, t), \mathbf{x} \rangle dt. \tag{3}$$

To derive eq. (3), Sahoo et al. (2024) impose two properties on $p_{\theta}(\mathbf{z}_s | \mathbf{z}_t)$. First, denoised tokens are never re-masked during sampling. Practically, this is achieved by manipulating the output of the neural network $\mathbf{x}_{\theta}(\mathbf{z}_t, t)$ to ensure that no probability mass is assigned to the MASK token. Secondly, already-denoised tokens are carried-over to the next sampling step. Sahoo et al. (2024) showed that both constraints lead to improved likelihood.

2.2 KNOWLEDGE DISTILLATION

Knowledge distillation (Bucila et al., 2006; Hinton et al., 2015) is a technique where a *student* neural network is trained to imitate the predictions of a more complex *teacher* model. One of the main advantages of distillation is the ability to reduce the inference cost associated with sampling from large LLMs while surpassing the performance of smaller models trained without distillation (Gu et al., 2024; Agarwal et al., 2024). The most relevant to our work are the distillation methods that match the predictions of the teacher and the student using a divergence measure δ :

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\delta(\mu_s(\mathbf{x}_t | \mathbf{x}_{<t}); \mu_t(\mathbf{x}_t | \mathbf{x}_{<t}))], \tag{4}$$

Where μ_s, μ_t are the AR distributions of the student and teacher, respectively, and \mathcal{D} represent the training dataset. Common divergence measures include f -divergences (Wen et al., 2023) such as the Kullback-Leibler divergence (KLD) or the total variation distance (TVD).

3 METHOD

3.1 SELF-DISTILLATION THROUGH TIME

As explained in section 2.1, discrete diffusion language models optimize the ELBO over the training examples. Fewer **decoding** steps typically lead to lower sample quality because the approximation of the reverse process is less accurate, as visible in the teacher curve in fig. 4.

To address the issue of low sample quality with fewer **decoding** steps, we propose *Self-Distillation Through Time* (SDTT). SDTT tunes a pre-trained MDLM to allow decoding with significantly fewer steps. Interestingly, our final model decodes samples with lower generative perplexity in 32 steps

Algorithm 2 One training round of *Self-Distillation Through Time*

```

216 1: Inputs: Training set  $\mathcal{D}$ , Teacher  $\mathbf{x}_\theta$ , Divergence measure  $\delta$ , Number of sampling steps  $m/k$ ,
217   Sampling step size  $\Delta$ , Mask token index  $M$ , Total number of training steps  $H$ 
218 2: Output: Distilled student  $\mathbf{x}_\nu$ .
219 3:
220 4:  $\nu \leftarrow \theta$  ▷ Initialize the student with the teacher weights
221 5: for  $i = 0, \dots, H - 1$  do
222 6:    $\mathbf{x}_0 \leftarrow \text{sample\_example}(\mathcal{D})$  ▷ Sample a training example
223 7:    $t_{\text{start}} \sim \mathcal{U}[0, 1]$  ▷ Sample  $t$  uniformly at random
224 8:    $\mathbf{x}_t \sim q_t(\mathbf{x}_t | \mathbf{x}_0)$  ▷ Forward diffusion process. See eq. (1)
225 9:    $\mathbf{x}_{\text{student}} \leftarrow \mathbf{x}_\nu(\mathbf{x}_t, t)$ 
226 10:   $\mathbf{x}_{\text{teacher}} \leftarrow \text{teacher\_SDTT}(\mathbf{x}_t, t_{\text{start}}, m/k, \Delta, M, 1e-5)$  ▷ See algorithm 1
227 11:   $\mathcal{L} \leftarrow \delta(\mathbf{x}_{\text{student}} | \mathbf{x}_{\text{teacher}})$  ▷ Compute divergence between student and SDTT targets.
228 12:   $\nu \leftarrow \text{backprop\_optim}(\mathcal{L}, \nu)$  ▷ Update the parameters of the student with AdamW
229 13: end for
230 14: return  $\mathbf{x}_\nu$ 

```

than the teacher would with 1024 forward passes. In short, SDTT improves the sampling speed by distilling the inference time computation to sample multiple steps into the student.

Let $p_\theta^{(m)}$ be the distribution of samples generated with m steps, using a denoiser with parameters θ . SDTT trains a denoiser with parameters ν to minimize a divergence d between $p_\theta^{(m)}$ and $p_\nu^{(k)}$. Here $k < m$, and k divides m (e.g., $m = 1024$ and $k = 512$):

$$\min_{\nu} d\left(p_\nu^{(k)} || p_\theta^{(m)}\right). \quad (5)$$

Since \mathbf{x}_θ and \mathbf{x}_ν are the only learnable elements of the sampling process, they completely determine the sampling distributions $p_\theta^{(m)}$ and $p_\nu^{(k)}$. As such, training \mathbf{x}_ν to match the predictions of \mathbf{x}_θ with fewer steps minimizes eq. (5). We now present a method for generating targets $\tilde{\mathbf{x}}_\theta^{\text{teacher}}(\mathbf{z}_t, t, m/k)$ to train \mathbf{x}_ν . Mathematically, we optimize the following objective:

$$\min_{\nu} \mathbb{E}_{\mathbf{z}_0 \sim \mathcal{D}, \mathbf{z}_t \sim q_t(\mathbf{z}_t | \mathbf{z}_0)} \left[\delta(\mathbf{x}_\nu(\mathbf{z}_t, t) || \tilde{\mathbf{x}}_\theta^{\text{teacher}}(\mathbf{z}_t, t, m/k)) \right], \quad (6)$$

where δ a divergence measure between the student and the teacher targets $\tilde{\mathbf{x}}_\theta^{\text{teacher}}(\mathbf{z}_t, t, m/k)$. We consider the Kullback-Leibler divergence (KLD), Total Variation Distance (TVD), and Mean-Squared Error (MSE). See appendix B for details on those divergence measures.

Generating the Teacher Targets Following the terminology of knowledge distillation, we call the denoiser \mathbf{x}_θ used for many steps decoding as the *teacher* and the denoiser \mathbf{x}_ν used for a few steps decoding as the *student*. To train \mathbf{x}_ν to match the predictions of \mathbf{x}_θ , we sample from the teacher for m/k steps. Whenever a MASK token is denoised, we collect the log probabilities predicted by the teacher for this MASK token. These log-probabilities become the distillation targets $\tilde{\mathbf{x}}_\theta^{\text{teacher}}(\mathbf{z}_t, t, m/k)$. Algorithm 1 outlines this process and fig. 3a presents it visually. While fig. 3a shows **how to distill two decoding steps in one**, the procedure can be extended to larger values of m/k . The complete SDTT training loop is presented in algorithm 2.

Iterated SDTT SDTT reduces the number of decoding steps by a factor m/k . If we want to reduce the number of decoding steps further, we can apply SDTT with $k' < k$, or alternatively apply SDTT n times, using the newly distilled student as teacher for the next round, **which we refer to as iterated SDTT**. Instead of directly optimizing the divergence in eq. (5), we introduce n intermediate distributions $p_{\nu_i}^{k_i}$ such that m/k_i is an increasing sequence as a function of i . In practice, we choose $m = 2^{10}$ and $k_i = 2^{10-i}$ with $0 \leq i \leq 7$ and sequentially minimize the objective

$$\min_{\nu} d\left(p_{\nu_{j+1}}^{(k_{j+1})} || p_{\nu_j}^{(k_j)}\right), \quad (7)$$

for $0 \leq j < 7$, where ν_j denotes the parameters of the j -th denoiser, with $\nu_0 = \theta$ (teacher). If the minimization procedure was perfect, minimizing eq. (5) or eq. (7) should result in the same solution.

270 However in practice, we observe that it is easier to minimize eq. (7) sequentially for increasing values
 271 of i , in a progressive fashion, similar to Salimans & Ho (2022).

272
 273 As an alternative to iterated SDTT, we tried using a single model and slowly growing the step size
 274 used to generate $\tilde{x}_\theta^{\text{teacher}}(\mathbf{z}_t, t, m/k)$. Unfortunately, this approach was unstable and the loss diverged
 275 after 30-50 steps, irrespective of how small the sampling step size was.

277 4 EXPERIMENTS

278
 279 We distill MDLMs on the OpenWebText dataset (Gokaslan & Cohen, 2019) as it was used to train
 280 recent discrete diffusion language models (Lou et al., 2023; Sahoo et al., 2024). We use the Adam
 281 optimizer with a learning rate of $6e - 5$, a batch size of 128 and no weight decay. We linearly
 282 increase the learning rate for 500 **training steps** and keep it constant afterwards. As a base model,
 283 we reuse the checkpoint released by Sahoo et al. (2024). See appendix C for more details.

284 In section 4.1, we evaluate 3 distillation divergences and show that **iterated SDTT can** reduce the
 285 number of sampling steps by a factor 16-32. In section 4.2, we ablate on the importance of hyperpa-
 286 rameters, including the duration of each round of **iterated SDTT** and the number of sampling steps
 287 to generate the targets $\tilde{x}_\theta^{\text{teacher}}(\mathbf{z}_t, t, m/k)$. In section 4.3, we scale SDTT to models with of up to
 288 860M parameters. Finally, in section 4.4, we compare the latency of SDTT against autoregressive
 289 models that use KV caching.

290
 291 **Generative perplexity** Following prior work (Dieleman et al., 2022; Lou et al., 2023; Sahoo et al.,
 292 2024), we use a larger model to compute the generative perplexity of **unconditional and conditional**
 293 **samples**. We evaluate the smallest students using GPT-2 (large) (Radford et al., 2019). In the scaling
 294 experiments, we use Llama3 8B (Touvron et al., 2023), since we compare models with up to 860M
 295 parameters.

296
 297 **MAUVE** We evaluate conditional generation using the MAUVE score (Pillutla et al., 2021).
 298 MAUVE measures how well a model follows a prompt by comparing multiple generations with
 299 a reference continuation. We use the first 1024 samples with at least 1024 tokens from the WebText
 300 dataset (OpenAI, 2019), take the first 50 tokens as a prompt, and generate 50 tokens of continuation.
 301 For each prompt, we generate 5 continuations, as done in Lou et al. (2023).


302
 303 **Sample diversity** Post-training can drastically reduce the diversity of language models (Kirk et al.,
 304 2024; Agarwal et al., 2024; Li et al., 2024). Hence, we measure the diversity of samples using the
 305 self-BLEU score (Zhu et al., 2018) with the same completions used to compute MAUVE.

306
 307 **Downstream performance** We measure the downstream performance using the **LAMBADA**
 308 **dataset** (Paperno et al., 2016), **as well as 6 multiple-choice question (MCQ) tasks from Gao et al.**
 309 **(2021). On LAMBADA**, we report an upper bound on the perplexity, computed using the ELBO (3).
 310 We also report the suffix accuracy by masking all tokens of the last word and predicting all of them
 311 in a single forward pass, using the **argmax** of the predictions. The diffusion model is correct only
 312 if all the masked tokens are decoded correctly in a single **decoding step**. **The 6 other benchmarks**
 313 **from Gao et al. (2021) evaluate the MCQ accuracy.**

314 4.1 ABLATION ON THE TRAINING DIVERGENCE

315
 316 SDTT requires choosing a divergence δ and we study the Mean-Squared Error (MSE), Total Vari-
 317 ation Distance (TVD) and (reverse) Kullback-Leibler Divergence (KLD). **We apply iterated SDTT**
 318 **for 7 rounds of $10k$ training iterations** and generate $\tilde{x}_\theta^{\text{teacher}}(\mathbf{z}_t, t, m/k)$ with 2 **sampling steps from the**
 319 **teacher** (algorithm 1). We use an exponential moving average (EMA) of the weights with a decay of
 320 0.9999 that we do not reset between rounds.

321 Figure 2 shows that students distilled with the KLD clearly outperform students trained using the
 322 MSE and TVD on LAMBADA. The LAMBADA accuracy of students tuned with the KLD slightly
 323 improves over the teacher, while the perplexity remains better or matches the AR baselines for all
 but the last round of SDTT. The improved accuracy on LAMBADA suggests that the model is better



(a) **Diversity of conditional generation (small scale).** We measure the trade-off between quality and diversity using self-BLEU (Zhu et al., 2018). Deterministic sampling yields a score of 1. The diversity minimally decreases after distillation.

(b) **Scaling SDTT to 860M parameters.** The plot compares the performance of the teacher and final student (7 rounds). The student and teacher have the same size. The small distilled student reaches lower perplexity than the large teacher.

Figure 4: **Sampling step ablations on perplexity.** Perplexity of samples after each round of iterated SDTT. (a): Iterated SDTT on a small model trained for 1M step. (b): Scaling SDTT to larger models trained for 400K steps.

at predicting multiple tokens in parallel after distillation with SDTT, since we evaluate the accuracy by decoding all tokens of the last word simultaneously.

Figure 5 shows that the KLD seem to outperform the MSE and TVD objectives on MAUVE. Since we generate sequences of 100 tokens only for MAUVE, following (Lou et al., 2023), we sample with at most 128 steps, and use samples generated with 128 sampling steps from the teacher as a baseline. Note that as observed by Deschenaux & Gulcehre (2024), discrete diffusion models typically achieve slightly lower MAUVE scores than AR models. Nonetheless, distillation with the KLD objective improves the MAUVE score of the students. Similarly fig. 17 shows that continuations from the student distilled with the KLD reaches the lowest perplexity and match GPT-2 with nucleus sampling in 32 forward passes.

In table 1, we compare the downstream performance on the tasks of Gao et al. (2021) before and after distillation. We observe that SDTT minimally affects the results, and that student distilled with the KLD objective reaches higher accuracies than other students in all but one task

Figure 4a measures the diversity of samples using the self-BLEU score (Zhu et al., 2018), for the students distilled with the KLD objective. See appendix A for results with the MSE and TVD. We find that SDTT minimally decreases the diversity. Compared to distilling autoregressive models (Agarwal et al., 2024), SDTT minimally reduces the diversity. For reference, Agarwal et al. (2024) routinely observes an increase of 15 in self-BLEU while we observe a change of at most 2 for the KLD student. See appendix A for more results and details on the self-BLEU score.

Figure 6 shows that students distilled with KLD have higher *unconditional* generative perplexity than those distilled with the MSE. However, KLD is the only objective that preserves performance in the LAMBADA data set while still significantly reducing the generative perplexity compared to the teacher. Therefore, in the remainder of this work, we focus on the KLD.

4.2 ADDITIONAL ABLATIONS

Number of steps in each SDTT round In section 4.1, each round of SDTT consists of $10k$ training iterations. Since the magnitude of the distillation loss does not reliably indicate convergence, we experiment with shorter rounds. We find that reducing the number of training iterations to $5k$ or $2.5k$ negatively impacted conditional generation performance, as shown in fig. 7. However, shorter rounds slightly improved the final generative perplexity (fig. 8) and resulted in marginally better LAMBADA perplexity (fig. 9). Since SDTT does not directly optimize the ELBO, an increase

378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431

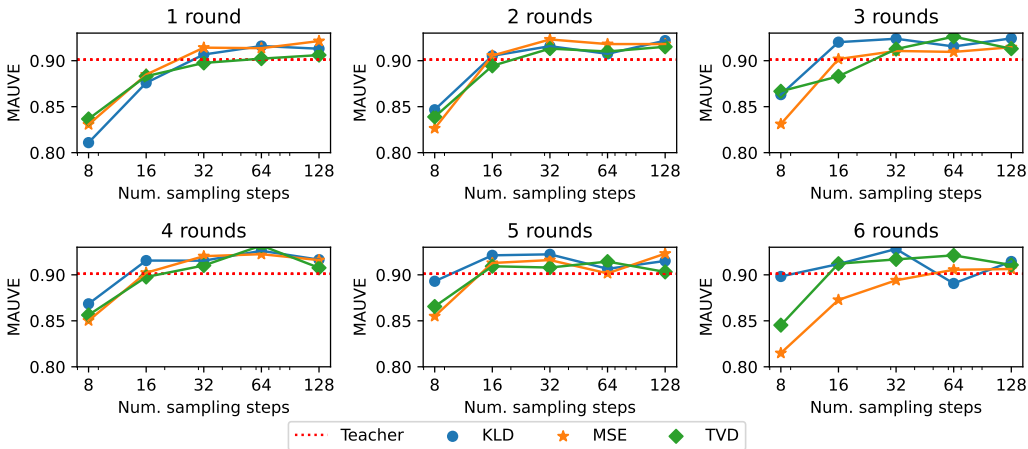


Figure 5: **MAUVE performance** of the student after each round of SDTT. The teacher performance is computed using samples generated with 128 decoding steps.

in perplexity is expected. Interestingly, the LAMBADA accuracy remains unchanged with shorter rounds.

Number of sampling steps to generate the targets In section 4.1, the targets $\tilde{x}_\theta^{\text{teacher}}(z_t, t, m/k)$ are generated using 2 sampling steps from the teacher. We explore distilling a larger number of sampling steps at once (4 or 8), since using more rounds of SDTT may induce more error accumulation in approximating the original teacher. Figure 12 shows that distilling more than two steps at a time is difficult and results in weaker results on LAMBADA. This suggests that the higher stochasticity of the targets generated with four or eight steps makes the task too difficult for the student.

Generating targets with the analytical sampler Lou et al. (2023) observe that using an analytical sampler (Campbell et al., 2022) results in higher quality samples compared to ancestral sampling. However, when generating targets $\tilde{x}_\theta^{\text{teacher}}(z_t, t, m/k)$ with analytical sampling, we observed minimal difference with ancestral sampling, as shown in fig. 10 and 11.

Resetting the optimizer and Exponential Moving Average between rounds Using an Exponential Moving Average (EMA) of the weights is known to improve the quality of samples from diffusion models (Nichol & Dhariwal, 2021). However, when applying SDTT for multiple rounds, it is unclear whether the EMA or current weights should be used as the teacher for successive rounds. Additionally, it could be favorable to reset the optimizer state between rounds as we grow the decoding step size. We experiment with two approaches: either resetting the optimizer state only, or resetting both the EMA and optimizer state. Figure 13 shows the generative perplexity when resetting the optimizer state and using the EMA as the teacher instead of the current weights, while fig. 14 presents the corresponding results for MAUVE. When using the EMA as teacher, since we accumulate updates in the EMA over 10k training iterations only, we use a slightly lower decay rate of 0.999. We find that using the EMA of the weights as the teacher may slightly improve performance.

4.3 SCALING SDTT TO 860M PARAMETERS

We apply SDTT to larger discrete diffusion models with up to 860M parameters. In this experiment, we train the models from scratch for 400k steps with a batch size of 512, a context length of 1024 and the Adam optimizer. We reuse the training configuration of Sahoo et al. (2024) and scale the models to larger sizes. We train 3 model sizes, small (169M), medium (424M) and large (863M). Details of the model architecture for each scale are shown in table 2. As for the other experiments, the models are diffusion transformers (Peebles & Xie, 2023) and we use an EMA with a decay of 0.9999. Although the results in section 4.2 suggest that short distillation rounds might be sufficient, it is unclear whether this result also holds on larger scales. Therefore, we use 10k steps per round of




Figure 6: **Perplexity for different losses and decoding step size.** Generative perplexity over 7 rounds of SDTT with MSE, TVD and KLD. While the KLD leads to a higher perplexity than the MSE; we focus on the KLD because it is the only divergence that retains the performance on the LAMBADA dataset.

SDTT. For simplicity, we generate targets using 2 teacher ancestral decoding steps and do not reset the optimizer state or EMA between rounds.

Since we train larger models, we evaluate the generative perplexity using Llama3 8B (Touvron et al., 2023). The generative perplexity over the 3 model sizes is shown in fig. 4b. Interestingly, the smaller diffusion model (169M) sampled from with 64 steps or more after distillation achieves better generative perplexity than the largest model (863M) when sampling with 1024 steps. In fig. 15, we show that the MAUVE performance also improves after distillation for the medium and larger model. Finally, in fig. 16, we see that the LAMBADA accuracy improves after distillation, similar as in the smaller scale, when using the KLD objective.

4.4 LATENCY WITH SDTT

While SDTT allows sampling from discrete diffusion models with 32-64 times less decoding steps, a quantity of interest to practitioners is the actual latency of text generation. Indeed, while the reduction in the number of sampling steps is large, since discrete diffusion uses a non-causal architecture, we cannot use KV caching (Pope et al., 2022). KV caching improves the inference performance drastically for AR models, hence we compare the latency of SDTT with GPT-2 with KV caching. We successfully reproduce the results of Deschenaux & Gulcehre (2024), which showed a 4x improvement when sampling with 32 steps, **and measure an 8x improvement with 16 decoding steps.** We compute the latency using **untrained** models with around 1.3B parameters, using the same hyperparameters as Deschenaux & Gulcehre (2024). We use a batch size of 8 and time the sampling 10 times after one warm-up step on a single A100 GPU with 80 GiB of RAM. **All models use FlashAttention (Dao et al., 2022). See Appendix A for additional experiments on the latency.**

5 RELATED WORK

Diffusion Models Diffusion models (Sohl-Dickstein et al., 2015b; Ho et al., 2020; Song & Ermon, 2020) are the basis of many state-of-the-art text-to-image models (Ramesh et al., 2022; Rombach et al., 2022; Saharia et al., 2022). After their introduction by Sohl-Dickstein et al. (2015b), Ho et al. (2020) showed that diffusion models can achieve FID scores (Heusel et al., 2017) comparable to GANs (Goodfellow et al., 2014; Arjovsky et al., 2017).

Discrete Diffusion & Diffusion Language Models Prior to Sahoo et al. (2024); Shi et al. (2024); Ou et al. (2024), Lou et al. (2023) introduced a novel discrete diffusion language model called SEDD. When decoding with a large number of steps, SEDD can match or surpass GPT-2 in unconditional text generation. The model of Lou et al. (2023) learn a discrete generalization of the score

Task	GPT-2	Teacher	KLD student	MSE student	TVD student
ARC-Easy	43.81	40.91	40.57	40.45	40.32
ARC-Challenge	19.03	21.08	20.73	19.28	20.05
HellaSwag	28.92	30.50	29.65	29.10	29.18
MathQA	21.21	21.78	21.47	22.28	21.84
PIQA	62.89	59.74	59.85	58.11	58.16
WinoGrande	51.62	50.91	50.75	49.57	50.36

Table 1: **Downstream evaluation results.** We report the accuracy of GPT-2, the teacher and students after 7 rounds of SDTT. Distillation seem to minimally affect the downstream performance.

of continuous diffusion models (Song & Ermon, 2020; Song et al., 2021). Campbell et al. (2022); Zhao et al. (2024) developed the continuous-time discrete diffusion framework. Hooeboom et al. (2021) extended Bernoulli diffusion (Sohl-Dickstein et al., 2015b) to categorical distributions, and Austin et al. (2023) generalized the work of Hooeboom et al. (2021) to more general corruption processes, including absorbing diffusion. **Zheng et al. (2024) develop a family of re-parameterized discrete diffusion models to enhance the training and decoding efficiency.** In parallel, several studies have explored continuous diffusion for language modeling (Li et al., 2022; Dieleman et al., 2022; Han et al., 2023; Chen et al., 2023; Gulrajani & Hashimoto, 2024). Despite recent breakthroughs, diffusion language models still have some drawbacks (Deschenaux & Gulcehre, 2024). **Ye et al. (2024) adapt Chain-of-Thought reasoning (Wei et al., 2023) to diffusion models.**

Distillation of Continuous Diffusion models Distilling continuous diffusion models is a well-studied area. For a comprehensive survey, see Luo (2023). Many distillation methods rely on Denoising Diffusion Implicit Models (DDIM) (Song et al., 2022), which showed that diffusion models can be sampled deterministically. Luhman & Luhman (2021) unroll trajectories sampled with DDIM and train a student to map noise directly to images. Luhman & Luhman (2021) pre-compute a dataset of noise-image pairs. Close to our work, Salimans & Ho (2022) teaches the student to match multiple sampling steps of the teacher, given corrupted training examples. However, unlike Salimans & Ho (2022), we cannot rely on the existence of a deterministic map via DDIM. Consistency distillation (Song et al., 2023) fine-tunes a pre-trained diffusion model to predict the final sample from intermediate points of the sampling trajectory, which enable faster sampling. Luo et al. (2024) distills a pre-trained diffusion model into single-step generator through a novel loss, *Integral Kullback-Leibler divergence*. *SD-XL Turbo* (Sauer et al., 2023) uses an adversarial formulation to sample with 1-4 steps from a latent diffusion model (Rombach et al., 2022).

Masked & Non Auto-Regressive Language Modeling BERT (Devlin et al., 2018) introduced the masked language modeling objective. While BERT focuses on representation learning, discrete diffusion language models are generative. XLNet (Yang et al., 2020) uses a generalized AR pretraining method to model the text distribution over all permutations of the training sequences, outperforming BERT on downstream tasks. Pannatier et al. (2024) adopt a similar objective to XLNet for generative modeling instead of natural language understanding.

6 DISCUSSION

In this work, we introduce *Self-Distillation Through Time* (SDTT), a distillation method for discrete diffusion models. Recent works (Lou et al., 2023; Sahoo et al., 2024; Shi et al., 2024; Ou et al., 2024) suggest that discrete diffusion models can match or outperform autoregressive models in text quality. However, those models require more inference resources than AR models to achieve good performance, because of the non-causal architecture of the neural network that prevents the use of KV caching. We show that SDTT can reduce the number of decoding steps while retaining performance. Our final student is up to 8x faster than AR models that use KV caching and we demonstrate that SDTT is applicable to larger models as well. In future work, we plan to evaluate SDTT on tasks that involve generating a large number of completions from a base language model.

7 REPRODUCIBILITY STATEMENT

We provide details on model architectures, hyperparameters, and provide pseudocode for our algorithm. We built on top of the open source model of Sahoo et al. (2024), which makes it relatively easy for researchers to reproduce our results. Additionally, upon de-anonymization, we will release our code and artifacts.

8 ETHICS STATEMENT

Overall, language models are dual-use technologies, and thus, they can have unethical uses, such as fake content generation, and they can suffer from bias if applied to data sets that are not carefully curated. This paper focuses specifically on speeding up discrete diffusion language models at test time to reduce their computational demands; we do not have specific concerns with regard to this contribution.

REFERENCES

- Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos, Matthieu Geist, and Olivier Bachem. On-policy distillation of language models: Learning from self-generated mistakes, 2024. URL <https://arxiv.org/abs/2306.13649>.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan, 2017. URL <https://arxiv.org/abs/1701.07875>.
- Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces, 2023. URL <https://arxiv.org/abs/2107.03006>.
- Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V. Le, Christopher Ré, and Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling, 2024. URL <https://arxiv.org/abs/2407.21787>.
- Noam Brown, Anton Bakhtin, Adam Lerer, and Qucheng Gong. Combining deep reinforcement learning and search for imperfect-information games, 2020a. URL <https://arxiv.org/abs/2007.13544>.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020b.
- Cristian Bucila, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Knowledge Discovery and Data Mining*, 2006. URL <https://api.semanticscholar.org/CorpusID:11253972>.
- Andrew Campbell, Joe Benton, Valentin De Bortoli, Tom Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models, 2022.
- Murray Campbell, A. Joseph Hoane, and Feng hsiung Hsu. Deep blue. *Artificial Intelligence*, 134(1):57–83, 2002. ISSN 0004-3702. doi: [https://doi.org/10.1016/S0004-3702\(01\)00129-1](https://doi.org/10.1016/S0004-3702(01)00129-1). URL <https://www.sciencedirect.com/science/article/pii/S0004370201001291>.
- Lingjiao Chen, Jared Quincy Davis, Boris Hanin, Peter Bailis, Ion Stoica, Matei Zaharia, and James Zou. Are more llm calls all you need? towards scaling laws of compound inference systems, 2024. URL <https://arxiv.org/abs/2403.02419>.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

- 594 Ting Chen, Ruixiang Zhang, and Geoffrey Hinton. Analog bits: Generating discrete data using
595 diffusion models with self-conditioning, 2023.
596
- 597 Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam
598 Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm:
599 Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):
600 1–113, 2023.
- 601 Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and
602 memory-efficient exact attention with io-awareness, 2022. URL [https://arxiv.org/abs/
603 2205.14135](https://arxiv.org/abs/2205.14135).
604
- 605 Justin Deschenaux and Caglar Gulcehre. Promises, outlooks and challenges of diffusion language
606 modeling, 2024. URL <https://arxiv.org/abs/2406.11473>.
- 607 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep
608 bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
609
- 610 Sander Dieleman, Laurent Sartran, Arman Roshannai, Nikolay Savinov, Yaroslav Ganin, Pierre H.
611 Richemond, Arnaud Doucet, Robin Strudel, Chris Dyer, Conor Durkan, Curtis Hawthorne, Rémi
612 Leblond, Will Grathwohl, and Jonas Adler. Continuous diffusion for categorical data, 2022.
- 613 William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter
614 models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39,
615 2022.
616
- 617 Leo Gao, Jonathan Tow, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence
618 Golding, Jeffrey Hsu, Kyle McDonell, Niklas Muennighoff, Jason Phang, Laria Reynolds, Eric
619 Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. A framework for few-shot lan-
620 guage model evaluation, September 2021. URL [https://doi.org/10.5281/zenodo.
621 5371629](https://doi.org/10.5281/zenodo.5371629).
- 622 Aaron Gokaslan and Vanya Cohen. Openwebtext corpus. [http://Skylion007.github.io/
623 OpenWebTextCorpus](http://Skylion007.github.io/OpenWebTextCorpus), 2019.
624
- 625 Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,
626 Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014. URL <https://arxiv.org/abs/1406.2661>.
627
- 628 Google. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*,
629 2023.
630
- 631 Sachin Goyal, Ziwei Ji, Ankit Singh Rawat, Aditya Krishna Menon, Sanjiv Kumar, and Vaishnavh
632 Nagarajan. Think before you speak: Training language models with pause tokens, 2024. URL
633 <https://arxiv.org/abs/2310.02226>.
- 634 Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. Minillm: Knowledge distillation of large lan-
635 guage models, 2024. URL <https://arxiv.org/abs/2306.08543>.
636
- 637 Ishaan Gulrajani and Tatsunori B Hashimoto. Likelihood-based diffusion language models. *Ad-
638 vances in Neural Information Processing Systems*, 36, 2024.
639
- 640 Xiaochuang Han, Sachin Kumar, and Yulia Tsvetkov. Ssd-lm: Semi-autoregressive simplex-
641 based diffusion language model for text generation and modular control, 2023. URL <https://arxiv.org/abs/2210.17432>.
642
- 643 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochre-
644 iter. Gans trained by a two time-scale update rule converge to a local nash equilibrium.
645 In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and
646 R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran
647 Associates, Inc., 2017. URL [https://proceedings.neurips.cc/paper_files/
paper/2017/file/8ald694707eb0fefe65871369074926d-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/8ald694707eb0fefe65871369074926d-Paper.pdf).

- 648 Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network, 2015.
649 URL <https://arxiv.org/abs/1503.02531>.
- 650 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020.
- 651 Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza
652 Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Train-
653 ing compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- 654 Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows
655 and multinomial diffusion: Learning categorical distributions, 2021. URL <https://arxiv.org/abs/2102.05379>.
- 656 Arian Hosseini, Xingdi Yuan, Nikolay Malkin, Aaron Courville, Alessandro Sordoni, and Rishabh
657 Agarwal. V-star: Training verifiers for self-taught reasoners. *arXiv preprint arXiv:2402.06457*,
658 2024.
- 659 Andy L. Jones. Scaling scaling laws with board games, 2021. URL <https://arxiv.org/abs/2104.03113>.
- 660 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child,
661 Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language
662 models. *arXiv preprint arXiv:2001.08361*, 2020.
- 663 Diederik P. Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models,
664 2023. URL <https://arxiv.org/abs/2107.00630>.
- 665 Robert Kirk, Ishita Mediratta, Christoforos Nalmpantis, Jelena Luketina, Eric Hambro, Edward
666 Grefenstette, and Roberta Raileanu. Understanding the effects of rlhf on llm generalisation and
667 diversity, 2024. URL <https://arxiv.org/abs/2310.06452>.
- 668 Adam Lerer, Hengyuan Hu, Jakob Foerster, and Noam Brown. Improving policies via search
669 in cooperative partially observable games, 2019. URL <https://arxiv.org/abs/1912.02318>.
- 670 Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B. Hashimoto.
671 Diffusion-lm improves controllable text generation, 2022. URL <https://arxiv.org/abs/2205.14217>.
- 672 Ziniu Li, Congliang Chen, Tian Xu, Zeyu Qin, Jiancong Xiao, Ruoyu Sun, and Zhi-Quan Luo. En-
673 tropic distribution matching in supervised fine-tuning of llms: Less overfitting and better diversity,
674 2024. URL <https://arxiv.org/abs/2408.16673>.
- 675 Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion language modeling by estimating
676 the ratios of the data distribution. *arXiv preprint arXiv:2310.16834*, 2023.
- 677 Eric Luhman and Troy Luhman. Knowledge distillation in iterative generative models for improved
678 sampling speed, 2021. URL <https://arxiv.org/abs/2101.02388>.
- 679 Weijian Luo. A comprehensive survey on knowledge distillation of diffusion models, 2023. URL
680 <https://arxiv.org/abs/2304.04262>.
- 681 Weijian Luo, Tianyang Hu, Shifeng Zhang, Jiacheng Sun, Zhenguo Li, and Zhihua Zhang. Diff-
682 instruct: A universal approach for transferring knowledge from pre-trained diffusion models,
683 2024. URL <https://arxiv.org/abs/2305.18455>.
- 684 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri
685 Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad
686 Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. Self-
687 refine: Iterative refinement with self-feedback, 2023. URL <https://arxiv.org/abs/2303.17651>.
- 688 Alex Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models, 2021.

- 702 OpenAI. Gpt-2 output dataset. <https://github.com/openai/gpt-2-output-dataset>, 2019. Accessed: 2024-09-30.
703
704
- 705 Jingyang Ou, Shen Nie, Kaiwen Xue, Fengqi Zhu, Jiacheng Sun, Zhenguo Li, and Chongxuan Li.
706 Your absorbing discrete diffusion secretly models the conditional distributions of clean data, 2024.
707 URL <https://arxiv.org/abs/2406.03736>.
- 708 Arnaud Pannatier, Evann Courdier, and Francois Fleuret. Sigma-gpts: A new approach to autore-
709 gressive models, 2024. URL <https://arxiv.org/abs/2404.09562>.
710
- 711 Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Quan Ngoc Pham, Raffaella Bernardi,
712 Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. The lambada dataset:
713 Word prediction requiring a broad discourse context, 2016. URL <https://arxiv.org/abs/1606.06031>.
714
- 715 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of*
716 *the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.
717
- 718 Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi,
719 and Zaid Harchaoui. Mauve: Measuring the gap between neural text and human text using diver-
720 gence frontiers, 2021.
721
- 722 Reiner Pope, Sholto Douglas, Aakanksha Chowdhery, Jacob Devlin, James Bradbury, Anselm Lev-
723 skaya, Jonathan Heek, Kefan Xiao, Shivani Agrawal, and Jeff Dean. Efficiently scaling trans-
724 former inference, 2022. URL <https://arxiv.org/abs/2211.05102>.
- 725 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language
726 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
727
- 728 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi
729 Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text
730 transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
731
- 732 Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-
733 conditional image generation with clip latents, 2022. URL <https://arxiv.org/abs/2204.06125>.
734
- 735 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
736 resolution image synthesis with latent diffusion models, 2022. URL <https://arxiv.org/abs/2112.10752>.
737
738
- 739 Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog,
740 M Pawan Kumar, Emilien Dupont, Francisco JR Ruiz, Jordan S Ellenberg, Pengming Wang,
741 Omar Fawzi, et al. Mathematical discoveries from program search with large language models.
742 *Nature*, 625(7995):468–475, 2024.
- 743 Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kam-
744 yar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Sal-
745 imans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image dif-
746 fusion models with deep language understanding, 2022. URL <https://arxiv.org/abs/2205.11487>.
747
748
- 749 Subham Sekhar Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T
750 Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language
751 models, 2024. URL <https://arxiv.org/abs/2406.07524>.
- 752 Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models, 2022.
753 URL <https://arxiv.org/abs/2202.00512>.
754
- 755 Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion dis-
tillation, 2023. URL <https://arxiv.org/abs/2311.17042>.

- 756 Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis K. Titsias. Simplified and gen-
757 eralized masked diffusion for discrete data, 2024. URL [https://arxiv.org/abs/2406.](https://arxiv.org/abs/2406.04329)
758 04329.
- 759 David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche,
760 Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman,
761 Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine
762 Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of go with
763 deep neural networks and tree search. *Nature*, 529(7587):484–489, Jan 2016. ISSN 1476-4687.
764 doi: 10.1038/nature16961. URL <https://doi.org/10.1038/nature16961>.
- 765 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally
766 can be more effective than scaling model parameters, 2024. URL [https://arxiv.org/](https://arxiv.org/abs/2408.03314)
767 [abs/2408.03314](https://arxiv.org/abs/2408.03314).
- 768 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised
769 learning using nonequilibrium thermodynamics. In Francis Bach and David Blei (eds.), *Pro-*
770 *ceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings*
771 *of Machine Learning Research*, pp. 2256–2265, Lille, France, 07–09 Jul 2015a. PMLR. URL
772 <https://proceedings.mlr.press/v37/sohl-dickstein15.html>.
- 773 Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsuper-
774 vised learning using nonequilibrium thermodynamics, 2015b. URL [https://arxiv.org/](https://arxiv.org/abs/1503.03585)
775 [abs/1503.03585](https://arxiv.org/abs/1503.03585).
- 776 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models, 2022. URL
777 <https://arxiv.org/abs/2010.02502>.
- 778 Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution,
779 2020. URL <https://arxiv.org/abs/1907.05600>.
- 780 Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben
781 Poole. Score-based generative modeling through stochastic differential equations, 2021.
- 782 Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models, 2023. URL
783 <https://arxiv.org/abs/2303.01469>.
- 784 Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. Roformer: En-
785 hanced transformer with rotary position embedding, 2023. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2104.09864)
786 [2104.09864](https://arxiv.org/abs/2104.09864).
- 787 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
788 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
789 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- 790 Trieu H. Trinh, Yuhuai Wu, Quoc V. Le, He He, and Thang Luong. Solving olympiad ge-
791 ometry without human demonstrations. *Nature*, 625(7995):476–482, Jan 2024a. ISSN
792 1476-4687. doi: 10.1038/s41586-023-06747-5. URL [https://doi.org/10.1038/](https://doi.org/10.1038/s41586-023-06747-5)
793 [s41586-023-06747-5](https://doi.org/10.1038/s41586-023-06747-5).
- 794 Trieu H Trinh, Yuhuai Wu, Quoc V Le, He He, and Thang Luong. Solving olympiad geometry
795 without human demonstrations. *Nature*, 625(7995):476–482, 2024b.
- 796 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
797 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural informa-*
798 *tion processing systems*, 30, 2017.
- 799 Peiyi Wang, Lei Li, Zhihong Shao, R. X. Xu, Damai Dai, Yifei Li, Deli Chen, Y. Wu, and Zhifang
800 Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations, 2024.
801 URL <https://arxiv.org/abs/2312.08935>.
- 802 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc
803 Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models,
804 2023. URL <https://arxiv.org/abs/2201.11903>.

810 Yuqiao Wen, Zichao Li, Wenyu Du, and Lili Mou. f-divergence minimization for sequence-level
811 knowledge distillation, 2023. URL <https://arxiv.org/abs/2307.15190>.
812

813 Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. An empirical analysis
814 of compute-optimal inference for problem-solving with language models, 2024. URL <https://arxiv.org/abs/2408.00724>.
815

816 Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V. Le.
817 Xlnet: Generalized autoregressive pretraining for language understanding, 2020. URL <https://arxiv.org/abs/1906.08237>.
818

819 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik
820 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models, 2023.
821 URL <https://arxiv.org/abs/2305.10601>.
822

823 Jiacheng Ye, Shansan Gong, Liheng Chen, Lin Zheng, Jiahui Gao, Han Shi, Chuan Wu, Xin Jiang,
824 Zhenguo Li, Wei Bi, and Lingpeng Kong. Diffusion of thoughts: Chain-of-thought reasoning in
825 diffusion language models, 2024. URL <https://arxiv.org/abs/2402.07754>.
826

827 Lingxiao Zhao, Xueying Ding, Lijun Yu, and Leman Akoglu. Unified discrete diffusion for categor-
828 ical data, 2024. URL <https://arxiv.org/abs/2402.03701>.

829 Lin Zheng, Jianbo Yuan, Lei Yu, and Lingpeng Kong. A reparameterized discrete diffusion model
830 for text generation, 2024. URL <https://arxiv.org/abs/2302.05737>.

831 Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. Texygen:
832 A benchmarking platform for text generation models, 2018. URL <https://arxiv.org/abs/1802.01886>.
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

864 A ADDITIONAL ABLATION RESULTS

865
866 In this section, we show additional plots on the ablations we conducted. Because the KLD was best
867 in retaining the performance on the LAMBADA dataset, we used it in most the ablations. Hence,
868 unless specified, the following experiments distill using the KLD.

869
870 **Ablations on the number of steps per round of SDTT** In fig. 7 we show the MAUVE perfor-
871 mance. In fig. 8 we show the generative perplexity, and in fig. 9, we show results on LAMBADA.
872

873 **Ablation on the analytic sampler** In fig. 10 we show results on LAMBADA, and on fig. 11 the
874 MAUVE score.
875

876 **Distilling more than 2 steps at once** In fig. 12, we show the generative perplexity.
877

878 **Ablation on the optimizer state and exponential moving average of the weights** In fig. 13 we
879 show the generative perplexity when resetting the EMA and optimizer state. In fig. 13, we compare
880 the generative perplexity when resetting the optimizer state only, and when resetting the EMA state.
881 Finally, in fig. 14, we show the MAUVE score.
882

883 **Plots for scaled SDTT** In fig. 15 we show the MAUVE score and in fig. 16, we show results on
884 LAMBADA.
885

886 **Conditional perplexity with TVD** In fig. 17c, we show the conditional perplexity (prompt ex-
887 cluded) on the small scale, for models trained for 1M steps. Empirically, the TVD performs worse
888 than the KLD and MSE.
889

890 **Measuring the diversity** We evaluate the generation diversity using the self-BLEU score (Zhu
891 et al., 2018). The self-BLEU score averages the BLEU score between one completion and the
892 others. Therefore, when the sampling algorithm is deterministic, the self-BLEU score is 1, and a
893 lower self-BLEU score denotes a more diverse set of samples. Formally, let $X = \{x_1, \dots, x_n\}$ be
894 conditionally-generated sequences, starting with the same prompt. The self-BLEU score can be
895 computed as
896

$$897 \text{self-BLEU} := \frac{1}{n} \sum_i \text{BLEU}(x_i, X \setminus \{x_i\}). \quad (8)$$

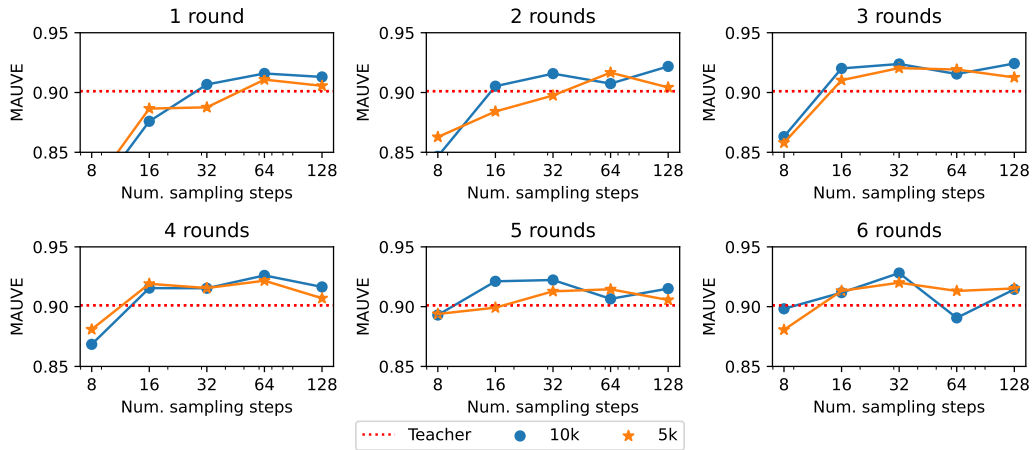
900 We compute the self-BLEU score using 1000 prompts, as for MAUVE, and generate 5 continua-
901 tions per prompt. Figure 4a, fig. 18a and fig. 18b show the self-bleu score after distillation with
902 the KLD, MSE and TVD objectives. Each objective only minimally decrease the diversity after
903 distillation. Compared to on-policy distillation of autoregressive models (Agarwal et al., 2024), the
904 decrease is marginal, as Agarwal et al. (2024) observe an increase of self-BLEU of the order of
905 10-20, demonstrating a more significant decrease in diversity.
906

907 **Decoding latency** In addition to the results on the 1.3B scale, we report the latency for models
908 with 169M, 424M, 863M, 3B and 8B parameters. We compute the latency with a batch size of 8 and
909 4. Figure 19 shows the latency with a batch size of 8 and fig. 20 using a batch size of 4. Figure 21
910 shows the trade-off between latency and perplexity. We measure the latency at the small model size
911 and compare GPT-2 with the final students after 7 rounds of distillation.
912

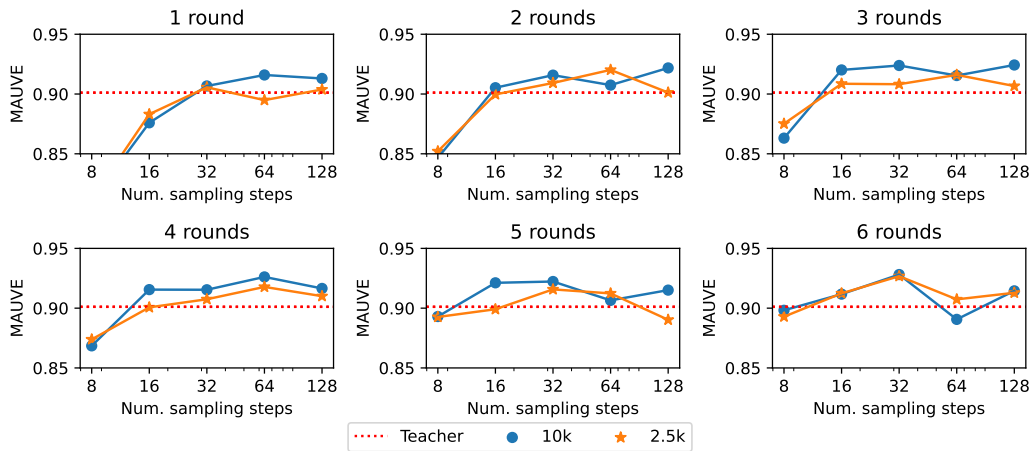
913 B ADDITIONAL DETAILS ON THE DIVERGENCE MEASURES

914
915 In this work, we teach the student to match the teacher targets $\tilde{\mathbf{x}}_{\theta}^{\text{teacher}}(\mathbf{z}_t, t, m/k)$ generated by algo-
916 rithm 1. We penalize the student deviating from the targets using one of three divergence measure:
917 the Kullback-Leibler Divergence (KLD), the Total Variation Distance (TVD), and the Mean-Squared
Error (MSE). We now describe each of them.

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

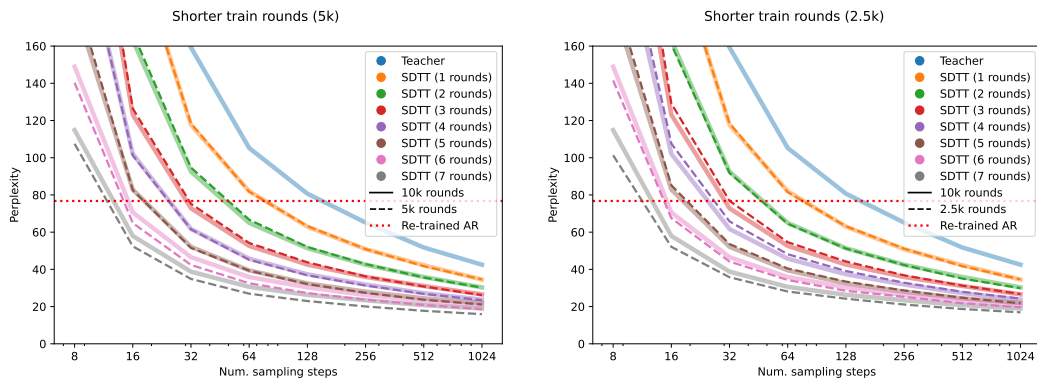


(a) 10k vs 5k iter/round.



(b) 10k vs 2.5k iter/round.

Figure 7: MAUVE performance with fewer steps per distillation round. It seems that using 5k or 2.5k distillation steps instead of 10k per round is detrimental to the MAUVE performance.



(a) 10k vs 5k iter/round.

(b) 10k vs 2.5k iter/round.

Figure 8: Generative perplexity with fewer steps per distillation round. Using 5k or 2.5k steps per round yields slightly improved perplexity after the latest distillation rounds while being a slightly worse in intermediate ones.

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025






Figure 9: Performance on LAMBADA when distilling with fewer steps per distillation round.



(a) Generative perplexity.



(b) LAMBADA.

Figure 10: Generative perplexity and performance on the LAMBADA dataset when using the analytical sampler. We find no clear benefit over the ancestral sampler.

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

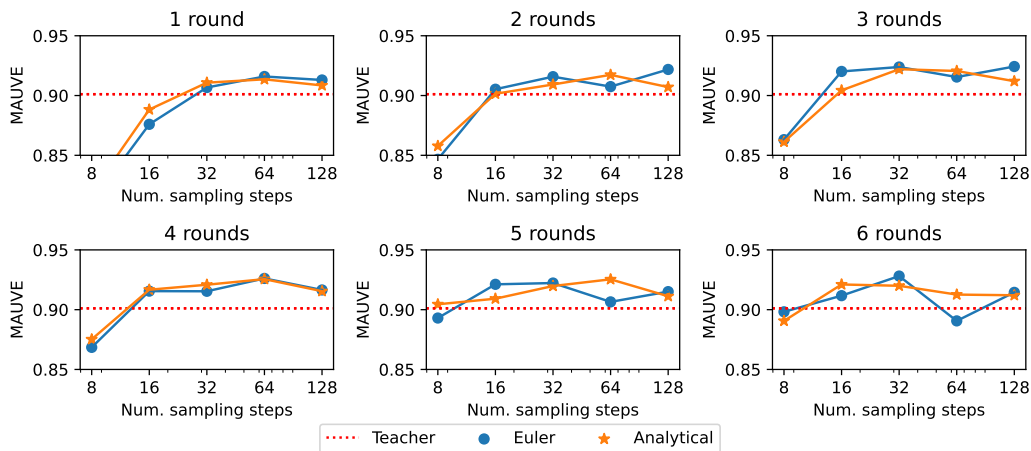
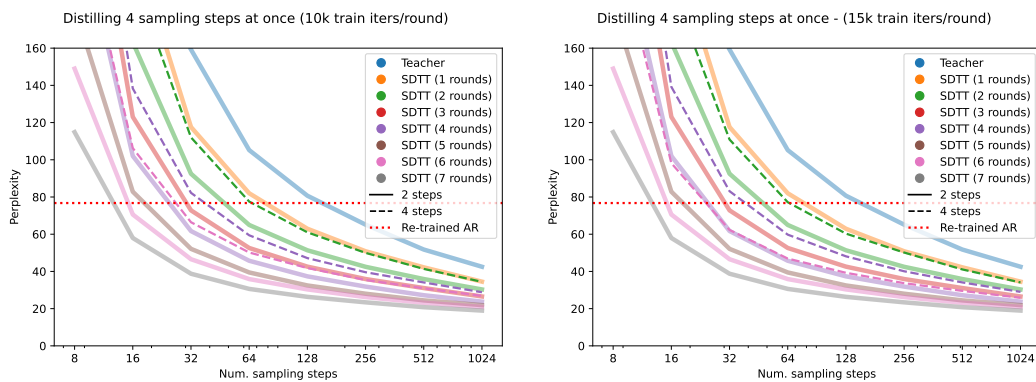
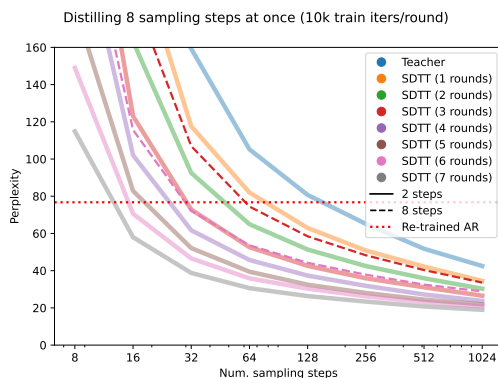


Figure 11: MAUVE performance when distilling using the ancestral sampler used by Lou et al. (2023). We find no clear benefit over the ancestral sampler.



(a) 4 steps.

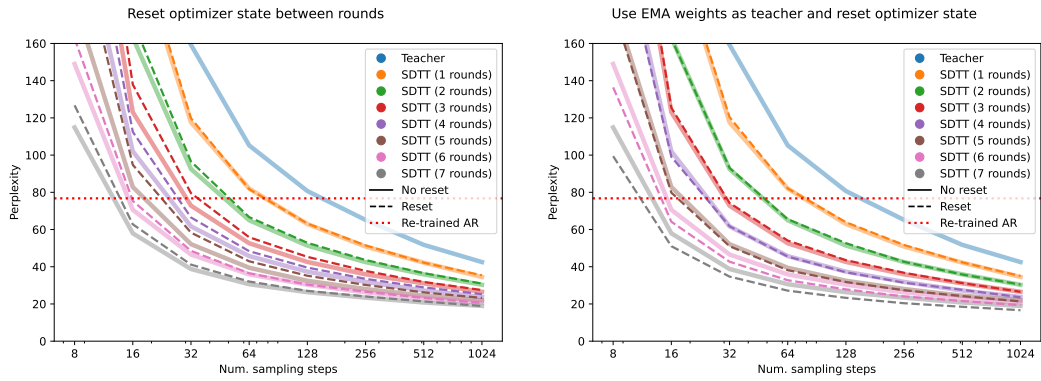
(b) 4 steps and 15k iter/round.



(c) 8 steps.

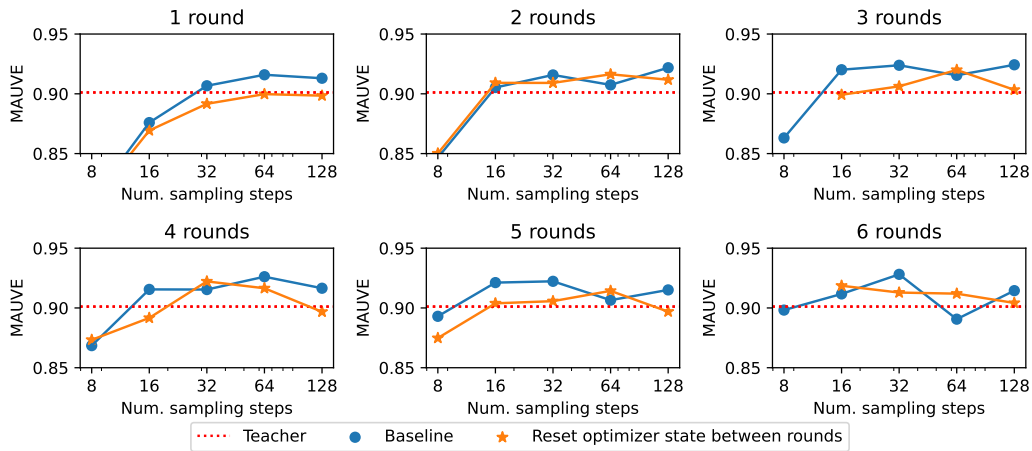
Figure 12: Trying to distill more than 2 teacher steps at once. (a): Distilling 4 steps at once. (b): Distilling 4 teacher sampling steps at once with more training iterations per round (15k). (c): Distilling 8 sampling steps per iteration. Overall, distilling more than 2 steps at a time seem to hurt performance. One could expect that distilling more than 2 steps at once would require longer rounds to train, hence we tried growing the round to 15k steps per round, which hurt the performance of the student.

1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

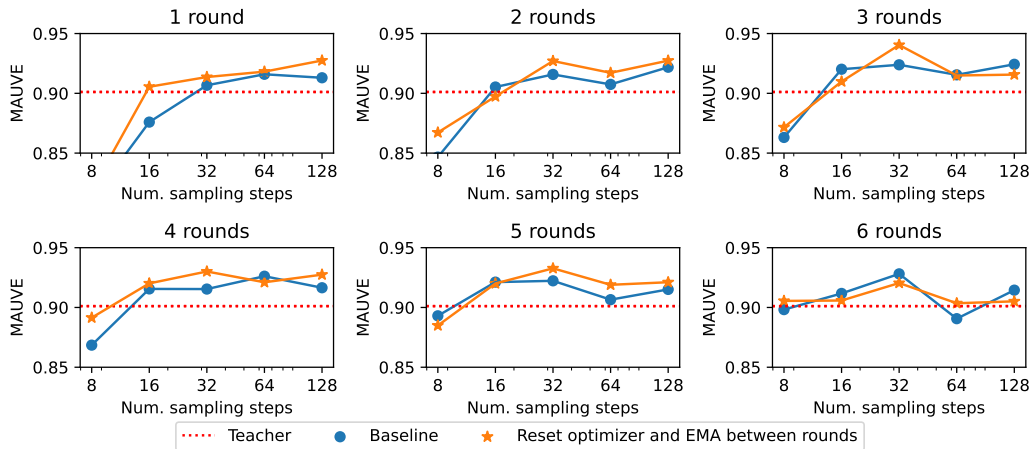


(a) Resetting the optimizer state between rounds. (b) Reset the optimizer state and use EMA of weights as teacher.

Figure 13: Generative perplexity when resetting optimizer or EMA state between rounds of SDTT.



(a) Resetting the optimizer state only.



(b) Reset the optimizer state and use EMA of weights as teacher.

Figure 14: MAUVE performance when resetting optimizer or EMA state between rounds of SDTT.

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

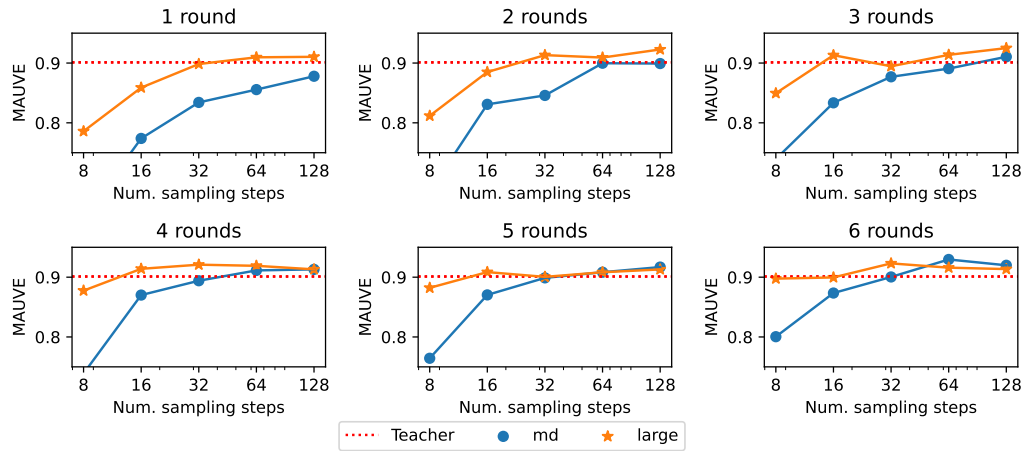


Figure 15: MAUVE performance of medium and large models pretrained for 400k steps. This experiment supports our claims that SDTT helps the final models to approach the performance of the teacher with less sampling steps.

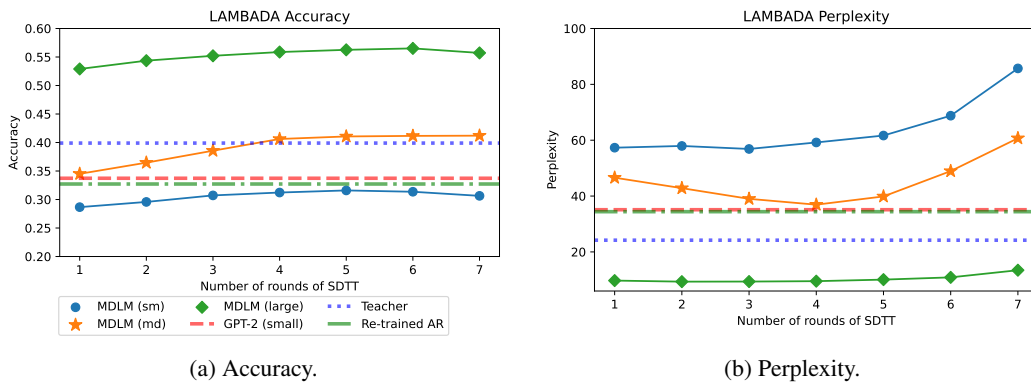
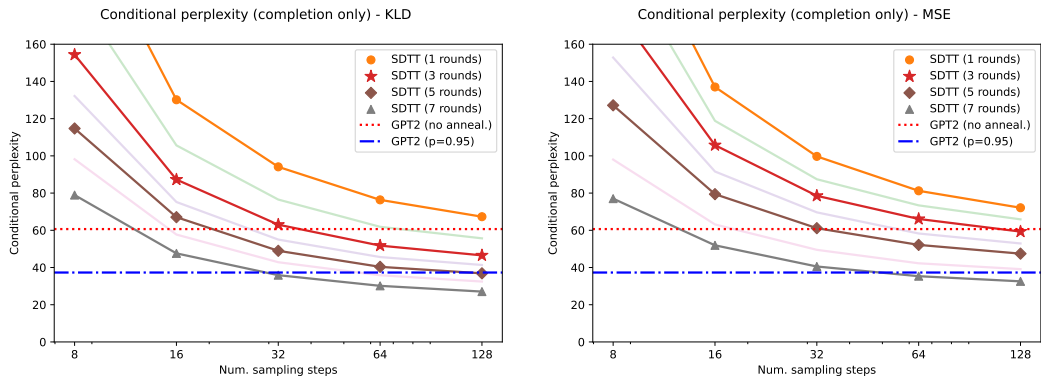
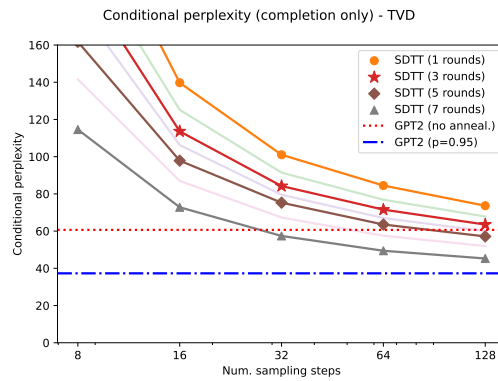


Figure 16: Accuracy and perplexity on LAMBADA when scaling SDTT to larger models. All models are trained for 400k steps before distillation. On the small scale, training for 400k steps instead of 1M yields a weaker model. Interestingly, the perplexity can improve after distillation when the models are undertrained.

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




(a) Perplexity of completions when distilling with the **KLD** objective. (b) Perplexity of completions when distilling with the **MSE** objective.



(c) Perplexity of completions when distilling with the **TVD** objective.

Figure 17: **Conditional perplexity.** Perplexity of the completions using GPT-2 large, excluding the prompt. SDTT with TVD performs worse. The final student distilled with KLD matches GPT-2 with nucleus sampling. Ground-truth continuations have a perplexity ≈ 13.11 .



(a) Distillation with the MSE loss. (b) Distillation with the TVD loss.

Figure 18: **Diversity of conditional generation (small scale).** We measure the trade-off between quality and diversity using Self-BLEU (Zhu et al., 2018). Deterministic sampling yields a score of 1. The diversity minimally decreases after distillation.

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

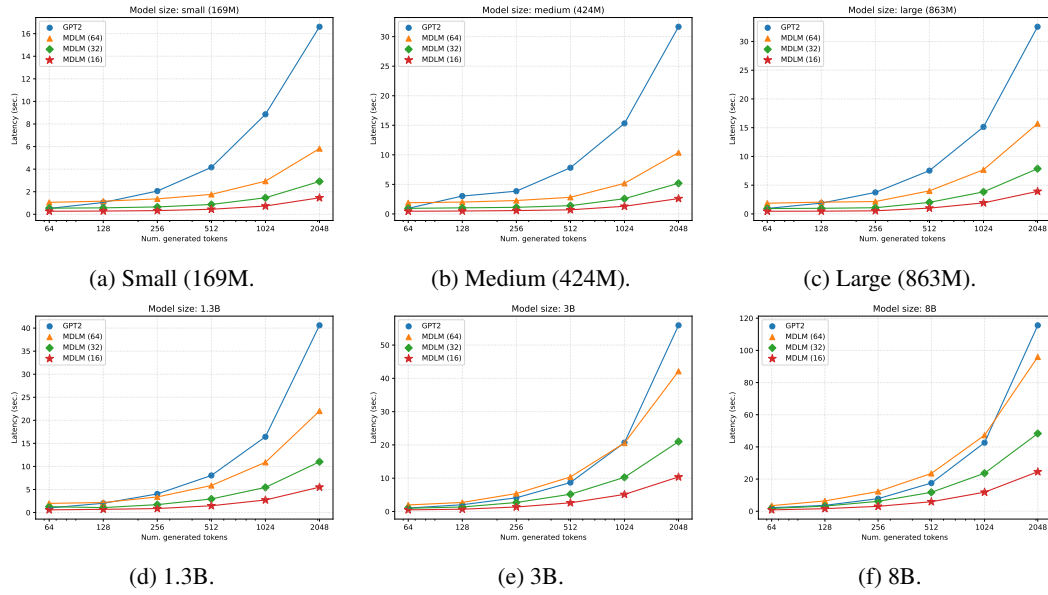


Figure 19: Additional latency experiments with a batch size of 8.

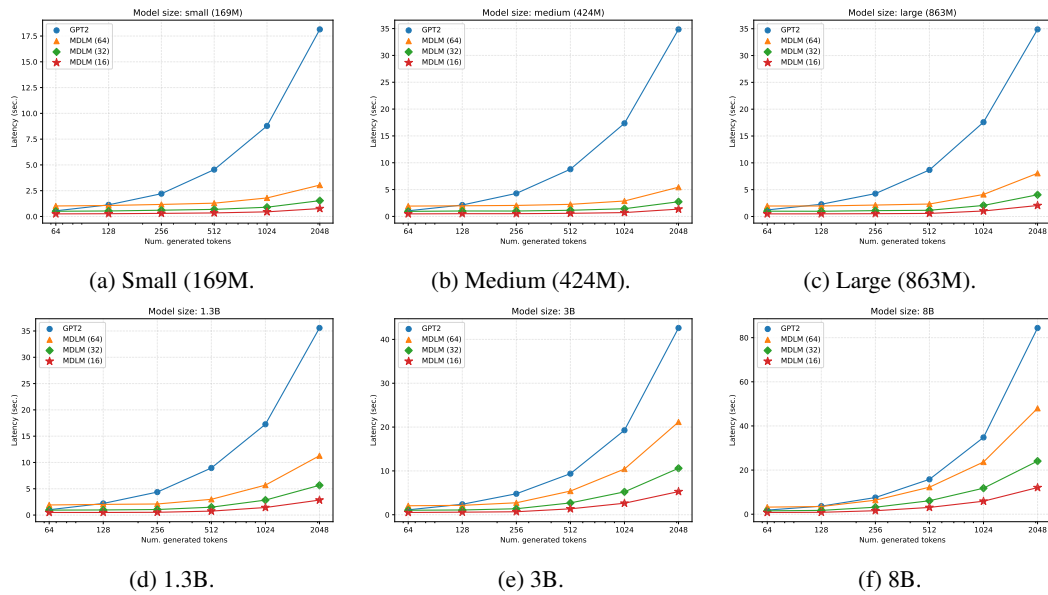


Figure 20: Additional latency experiments with a batch size of 4.

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




Figure 21: Perplexity vs wall-time latency (in seconds) for small models. We use 16, 32, 64, 128 ans 256 decoding step for the diffusion models.

Model size	small	medium	large	1.3B	3B	8B
# params	169M	424M	863M	1.3B	3B	8B
Num Layers	12	24	24	24	26	40
Embedding dim.	768	1024	1536	2048	3072	4096
Num. heads	12	16	16	32	32	32

Table 2: Hyperparameters of the diffusion models at different scales. All models use RoPE positional encoding (Su et al., 2023).

B.1 KULLBACK-LEIBLER DIVERGENCE

The *Kullback-Leibler Divergence* (KLD) between two discrete distributions p and q defined on the same finite sample space Ω is computed as

$$D_{KL}(p||q) := \sum_{x \in \Omega} p(x) \log \frac{p(x)}{q(x)}. \tag{9}$$

The KLD has a unique minimum when p and q are equal, however the KLD is not symmetric, meaning that $D_{KL}(p||q) \neq D_{KL}(q||p)$ in general. In this work, we train the student with the reverse KLD $D_{KL}(p_{\theta}||p_{teacher})$. In the next paragraphs, we present differences between $D_{KL}(p_{teacher}||p_{\theta})$ (forward KLD) and $D_{KL}(p_{\theta}||p_{teacher})$ (reverse KLD).

The Forward KLD The forward KLD is called zero-avoiding because if $p_{target}(x)$ is non-zero but $p_{\theta}(x)$ is close to zero, then $p_{target}(x) \frac{p_{target}(x)}{p_{\theta}(x)}$ will be large. To minimize the forward KLD, p_{θ} will try to assign non-zero probability to all points where p_{target} is non-zero.

The Reverse KLD The reverse KLD is called zero-forcing because if $p_{target}(x)$ is close to zero but $p_{\theta}(x)$ is not, $p_{\theta}(x) \frac{p_{\theta}(x)}{p_{target}}$ will be large. To minimize the reverse KLD, p_{θ} will try to assign zero probability to points where p_{target} is close to zero.

B.2 TOTAL VARIATION DISTANCE

The *total variation distance* (TVD) is a metric used to compare two probability distributions. For two discrete probability distributions p and q defined on the same finite sample space Ω , the TVD is computed as:

$$d_{TV}(p, q) = \frac{1}{2} \sum_{x \in \Omega} |p(x) - q(x)|. \tag{10}$$

1350 The factor of $1/2$ ensures that the TVD ranges between 0 and 1, where $d_{\text{TV}}(p, q) = 0$ if and only if
 1351 $p = q$.

1352

1353 B.3 MEAN-SQUARED ERROR

1354

1355 Unlike the Kullback-Leibler divergence (KLD) and Total Variation Distance (TVD), the MSE can
 1356 be used to compare any scalar quantities, not just probability distributions. For numerical stability,
 1357 we compute the MSE in log space:

1358

$$1359 \text{MSE}(p, q) = \frac{1}{|\Omega|} \sum_{x \in \Omega} (\log p(x) - \log q(x))^2. \quad (11)$$

1360

1361

1362 B.4 χ^2 DIVERGENCE

1363

1364 The χ^2 divergence can be used to compare two probability distributions. For two discrete probability
 1365 distributions p and q defined on the same sample space Ω the χ^2 divergence is computed as:

1366

$$1367 d_{\chi^2}(p, q) = \sum_{x \in \Omega} q(x) \left(\frac{p(x)}{q(x)} - 1 \right)^2 = \sum_{x \in \Omega} \frac{1}{q(x)} (p(x) - q(x))^2. \quad (12)$$

1368

1369 As such, we see that the χ^2 divergence is related to the MSE. Note that when using the MSE for
 1370 distillation, we penalize the error in log space, while the χ^2 penalizes error in probability space.
 1371 Additionally, the MSE uses a uniform weight factor $\frac{1}{|\Omega|}$ for each term of the sum, while the χ^2
 1372 divergence uses a weight of $\frac{1}{q(x)}$.

1373

1374

1375 C IMPLEMENTATION DETAILS

1376

1377 **Architecture** To compare with Sahoo et al. (2024), we trained the diffusion models using their
 1378 code and pre-processing steps on the OpenWebText dataset (Gokaslan & Cohen, 2019). As Sahoo
 1379 et al. (2024), our models are not conditioned on the noise level. Nonetheless, Sahoo et al. (2024)
 1380 kept the architecture of Lou et al. (2023) unchanged and makes the model unconditional by feeding
 1381 it a zero tensor instead of the noise level. Removing the adaptive layers could improve the sampling
 1382 speed further, but we avoided modifying the architecture to prevent potential problems. See table 2
 1383 for the hyperparameters of our models.

1384

1385

1386 D TEXT EXAMPLES

1387

1388 We include non-cherry picked text generated from the small distilled model with KLD loss from
 1389 the last round of distillation via unconditional sampling with varying number of steps. We show the
 1390 first 512 tokens to so that the text fits on one page. Remember that those models are small and not
 1391 fine-tuned for text quality. They can also start generating in the middle of sentences, since they are
 trained on a concatenated corpus of documents.

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403

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

Text generated with 16 steps (1/3)

invite to the gathering, because he was invited in 2008, probably on a regular basis thereafter. But to become a scientist, to verify those cred veracity is important," inlamali said.

CNN is thus creating a monster that has supporting cascade of other grand jury investigations, he said.

"In the case of Mr. Eliaschis, I wrote in a today to a number of everyone involved in consideration of this matter; these people are invited, named and considered 'committed' to the process and trust of the Nation," he said.

There have been no complaints or formal complaints and this will directly no longer be CNN's standard and indepth coverage.<|endoftext|>because my office cherish diversity , this is the approach we have come across. their subscription model is great, and right now there are folks in our office that want to help to promote diversity. so we're looking forward to hearing those responses from them. although we realize it is a different place than we run it. but outside of this, I think we've never had a lot of conversations (especially here) about community-based leadership being happier than market-based leadership: the leader is fantastic, the person is valuable and talented.in the UK things don't that way. it has a leaderless culture which has not been well with a hierarchical planning and organizing process. we have a very specific image of this kind of organization. but one of such qualities is the image of someone accomplished like everybody else does, which interests us as do the talks of one of our public figures .

as part of what we'd love to do here on our social impact endeavors. recognize that most of the work here, we're in the midst of the first day of the interview (which Prof. Garry had posted to the blog). Garry was kind enough to come participate in the interview as well as conducting and perusing on his and the next few competitors, in order to get valuable feedback. I wanted to feel enthusiastic about the process, eager to share feedback, and expect to have a very professional experience.

but, broadly speaking, more than a lot of the things that we struggle with the capacity to report, we just made a draft, and then got the post called. there is one example of things throughout the draft that made me the most-- . The media is a small part in Far Left.

it is really important to have a relationship with your employer. At a high

1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511

Text generated with 16 steps (2/3)

kids not being in our schools as a result. The kids put their families in Florida schools in this district, this is Georgia school, and not only do they have enough time to work for a Florida firm, it's not desirable for our kids to be in Florida school anymore."

The brothers turned to public education and the governors quickly asked them if they wanted to. Bonding Aid then was contacted once asked for a special order from the administration setting aside \$524,000, but they were denied despite the requestBy Bill Othello, according to government spokesman. The brothers wrote letters to the governors numerous times claiming the information was false, including one letter, which suggested that they funnel \$2,000 to the Slothouse Clearing House schools through the University of Florida. However, one of the brothers, Chris Yates, told Bloomberg News that he still was shocked and horrified by the correspondence, saying, "Obviously, I felt like a coward to be in this of a very uncomfortable situation."

Instead however, Yates said, he reacted very much like he pissed off at one of his favorite politicians, Bush.

Florida's two Republican leaders have strained relations, in particular with recent governors Doug Ducey, a Republican , and Jeb Bush, a Republican, addressing his concerns. "I do not think the other leaders will do this, but we do have to work to make sure that is how we have to do it, something that's something that we need to be doing, and that there will be always a need for better quality education, too," Bush also said. "We do not want Florida to stop funding education and essentially contradict the fact that what has a provenance in this planet is that killing was at least 10 percentage and many people got dead. We I beginning to have problem with that and I think that's the first where we know, for sure, when we're going to eventually share that information with the public, and we feel in order to deal with that we have to agree to the efforts that we are engaged in and also [Haley and I] are ultimately going to have to provide that . We will have the heavy lifting to do whatever we do. I think this is irresponsible but also that we are not going to start having a real conversation because we've got a lot to do, but that was a eye opener to us."

When asked on Tuesday if he was dazed and that he regretted attending the meeting

1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565

Text generated with 16 steps (3/3)

was subjected to will then acidize in the form of foam. The pH and pitch of foam create a, too substantial a gap between an egg and tissue, which will drive it to accumulate faster, and can thus cause irritation to the sensitive parts of the body. It's penalties are well as analgesia, and shortens chances of learning how the material works.

Mr. Segal, who was involved in the study, did not acknowledge the limitations of the study to treat his own specifically painful dental fracture, but added that it succeeds in all aspects of the process. "This is the first time that it might make a significant contribution to improving dental health. Until then we will have to get better at adjusting what have changed to make sure that it is effective." Yet, in June 2013, Laquan Phillips, 20, a promising medical student at the Jackson State University School of Medicine and the son of a cyclist Philando Castile, killed two or four weeks earlier, was forced to give a the dozens of officers. officer just about three seconds to pull up on the vehicle trying to lock the black car into silos, and Fewell, the neighborhood managed by the officer who arrested Phillips, was refusing to go all the way. The officers were concerned about the direction of bullets so that they could hit a casting bit.

While police believed the bullet hit a man on the left side - and a belief that had been consistent since police had been in deep denial, it failed to hit the man on the right side of the table. When you interjected a shot into the man's first body, the face mostly rolled down the throat - a crushing moment of motherhood and last sweat for a father of two - and thankfully the other one was stopped in his tracks.

"We believe him," said John Milliken, a police officer at the time of his shooting, who could not confirm other deaths but said that Oli was firing gun. Others would say it was the product of head trauma.

The policeants in Roswell could rely on a variety of lethal weapons, they said, including ammunition that police had accessed. I want to thank the people for the first- aid kit for the family, and the people and the people for Justice Jesse James also, at hand members of the Black community the 71st.

An anonymous person was having a phone conversation with the Buckeyes's interim president, who was to take part in the participants of

1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619

Text generated with 32 steps (1/3)

Wilkins was he committed the acts of vandalism as a juvenile

Woodward said he had "since confessed the crime based on information and an explanation for what he did." But a police spokesman told the news outlet that investigators work for the government and it is the responsibility of whether it is the individual with knowledge of the crime or that should be punished as well as those involved with the system.

"Liberals can withhold confidential information from the public on the basis of any reason or whether such information is a public interest," the spokesman said on condition of the anonymity because of the investigation.

Fox News reporter Brian Stelter, reporting the agency's look at the alleged Hammond scheme, took a high-profile line on political campaigning in Russia and current affairs on Fox News Channel's propaganda channel, The Which?," for instance.

The Rasmussen Reports cited concerns of a spike in voter fraud, citing the divide between the Democrats and voters , many of whom Republicans outside the party vowed to lose in the election.

The Washington-based advocacy group We Are America, which publishes figures from the polls, said the report was based on the promises of conservative news outlets, independent organization, and the use of pollsoddy reporting.

"The communication and reporting of American media has been critical," We Are Russia said in a statement following the report.

The United States media has been biased against the election, which Russian officials linked to hacking on Democratic computers and other promotional efforts that seem aimed at raising security fears in Moscow and toppling many of the Kremlin's market allies in the U.S, which is moving closer to winning the election before a referendum begin taking place.

Russian officials have attributed their bias in opinion polls toward the 2016 elections to an uptick in unemployment statistics, with figures they're publishing being forecast on the eve of their presidential elections.

Russian officials say they would like to prevent fraud, and polls show that they may have aggressive on potentially maintaining a narrow five-point lead.

The Rasmussen poll report came days after the state Attorney General filed a lawsuit in general against the Rasmussen Reports' report, and asked the state to look to the opinion polls and determine if the Rasmussen poll did a " good job."

Democrats announced plans to use a federal court case to

1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

Text generated with 32 steps (2/3)

Sherman was a city historian for 28 years before 1988 and former relations adviser spoke with more than a dozen Asheville staff members during an interview outside the hotel early Monday morning. In a brief interview, Sherman gave employees an overview of what they seerved from the two Highlanders - including a handwritten piece of electric paper and pictures of the cut and logings, perhaps sweeter than the compensation Yancey and Davis did compensation settlement for.

Photo: Courtesy U.S. BAG via Jan. 30; The documents provided in the report stem from a link to a minor change in the law, according to the report.

The law officially establishes an enforcement and oversight process (PDF) for transferring payments to the other committee, or the City or State Board, that has the mandate to require the cash transfers from the other side ;

The legislation, upon meeting each committee, instead directs the legislature to set its own financial policy in the other committee;

The method suggests payments to the other committee can form the legislature's own committee to pass enforcement and oversight legislation, and that each of the committees may vote on motions for a resolution and give the public a vote or motion to approve legislation by the Joint Finance Committee.

That, the city and state will follow the same process as the New Hampshire City and Lodging Act. However, in the Nov. election the Legislature prepares to decide whether to modify the rules and procedures surrounding such a Senate bill, so it would be unclear how long the interim Legislature policy to be implemented once it's adoption is in place.

The terms of the report are not necessarily overarching, some of the other ways are specifically saying that a proposed legislation that would transfer payments to the other committee must be amended to include a general bill , before it would be recognized as a legislative body.

in Khan at editors@time.com

Twitter: @in_khan<|endoftext|> independent party New Greens has urged Parliament to declare that there was no reason for the war. Many Australians who believe that it would be better to their children going to war are handed local police officers the chance only for training next week. The leader of South's Republic of the African state, Christopher Robertson, said that the government's refusal to give any status to the force is the worst for Australia's history. In a statement to the Liberal party, Mr Robertson said it is important to keep a police force on ground as it sees fit. This has caused panic among all 16 states, in particular the growing

1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

Text generated with 32 steps (3/3)

, the police, media and public

International Council of the Community of Europe (IECE) is preparing the next steps, being taken on Monday - that could the potential to make firm changes to Britain and the EU, including the introduction of modern regulations and procedures on trans financial transactions from ordinary individuals to private corporations.

However, senior mainstream EU members told the Telegraph that some of the changes in the legislation at the moment " are not clearly relevant to that role" for regulatory oversight, as well as legal procedures, for the provision of business and financial industry, social, welfare, police, promotions, public health and economic activities - and finally, to give the EU the right to make the legal and regulatory decision that is needed by the wider market.

A senior Council official in the UK, speaking on condition of anonymity for the sensitivity of the meeting with his staff, said there had been little progress in negotiation negotiations and how to implement the changes because the legislation was still on the way. The senior official said both countries had worked fully with each other so the legislation could meet the basic existing EU regulations.

"We agree then that we will need to amend the legislation in a way that is appropriate for EU law in the parts of the UK. So review and consultation is something we are looking very closely to to ensure that our new amendments to the legislation give additional leadership and framework ... to verify we can implement those measures and will provide appropriate support of expertise to the UK in relation to getting them."

However, the official said it could take some time still to implement the legislation and that any changes put in place require a complete clean review of the wording of legislation, and part with EU institutions in that way.

"We understand the process that is underway on how to implement those measures - and implementing those measures is going rapidly, so we can't yet start to assess the situation," he said.

The Council has agreed on January 25 to pressure the UK government on the decision to leave the EU to good effect

"The previous government had helped reform the law around the UK and contributed to that change," said Matthew Kennedy , an official for the Council.

The Council will work on the new amendments as part of two of the terms of engagement with the British government, signed by David Cameron.

"They are part of a broader effort and to do so has not been determined, so the Prime Minister and those

1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781

Text generated with 64 steps (1/3)

create big gas storage wells locally. BP has proposed investing billions of dollars to build new big storage wells, which will raise gas production to new levels of production.

The research has been co-sponsored by Congress lobbying to oppose oil and gas drilling (FAPL), but the administration has yet to outline how far oil and gas drilling and the Eagle Ford expansion will take its new forms, with new risks rising.

SUS energy industry officials say changes in the measures ordered and new funding levels have been little bearing on oil production, capped in about 60% last year, but some proposals are still receiving orders - still others have yet to be overcome - by the Department of Energy to produce a package on a proposed Dakota Access oil pipeline to the West in the West Coast.

The distinction between fossil fuel and gas is also blurred on energy measures that have been laid out to Congress about a year; most are now in the process to agree on the proposed Keystone XL pipeline, and the Dakota Access pipeline.

The US energy market is fast moving, so the focus of the Obama administration has been to push jobs to the US and to encourage Americans to have better financial prospects for creating their new jobs

Gans to set low-carbon renewable energy targets will soon launch at the White House. Since 2009, infrastructure projects have established several significant initiatives : restructuring the coal sector in the West of Europe, reducing the tariffs on renewable energy in Britain, and implementing zero-carbon policies in South Africa, despite high difficulties in succeeding in implementing EU emissions standards.

The infrastructure initiatives, announced and unveiled by the president on Wednesday, are focused on the aim of finding effective ways to build on competitiveness at long-term scales and to help around the world create a sustainable energy portfolio in the areas needed to mitigate the risks in local economies.

But institutions as big and large as central banks infrastructure plans to support renewables have been especially critical in recent years. A number of EU states such as Finland and Sweden were taking initiative with their plans. The states, which will also launch a number of the other initiatives at the White House ceremony, said that "the energy policy environment will receive high profile development in the UK and the economy of the EU in most of them".

The far-reaching strategy has been to place low-carbon targets in the UK's renewable energy sector, of course, with a push of the private sector - much of which is being led by governments and businesses - to scale their renewable energy

1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835

Text generated with 64 steps (2/3)

exchange Bitcoin, it should be able to deliver such services.

Basically, a service that hinges on Coinbase make payments for those bitcoin users who use this unique means of purchasing and storing bitcoin, and that it's able to reach retailers as wide as possible, which makes it so attractive for merchants for transaction speeds. Once they've taken the step of overcoming a long wait for an WePay service, they're permitting merchants to accept Coinbase to begin. What explains the delay is that even if they decide moving forward with an offer to pay for maintaining the price of bitcoin and the impressive growth of bitcoin, it remains a high barrier and additional cost for some users. What's sure for certain is that merchants will accept the service, but that it will deliver a version of bitcoin as a method of payment.

They also plan to allow Coinbase to maintain control of the new address, to ensure the security of the integrity of Coinbase, as well to further expand the service available to Mac users. And besides, Coinbase is still investing in BitPay, a leading bitcoin exchange offering for free an alternative to fiat value, which would be a perfect fit for a marketplace for all that bitcoin in a heavily regulated world.

We have reached to WePay Bitcoin for our response and for their views on our WePay service.<|endoftext|>How do you describe the typical user of marijuana. Your questions are always open to the cipscoop.net email team. Today is Day of Cannabis, an event celebrating the global outside of the marijuana industry. Let us know your favorite questions.

K: Let's talk about the for-profit group that supports marijuana, Gives of Health, which established the first nationwide standard price system for marijuana. Can you name a name?

In its first statement, Gives of Health defined itself as committed to supporting a national free market system of marijuana, promoting policies to reduce the supply of legal prescription drugs, reduce tobacco prices, reduce wasteful government spending, and taxing and regulating marijuana. We believe a marijuana marketbased tax system is good for our public health and for the economy by helping people circumvent drug laws and purchase marijuana, reducing tax costs associated with prescription analgesics and avoiding tax evasion.

K: Leon and several advocacy groups that try to crack down on marijuana, as the Gang of Eight, introduced bills to make cannabis and medical marijuana for adults in more than a dozen states. What is that to say about addressing a lot of the problems?

1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889

Text generated with 64 steps (3/3)

by the military. The army (which elected its President) saw the most important coup in its history.

The broader American Catholic Church claim that one of the first indications of the twentieth century thought of "liberated Christians" not earlier than July, 1960 was that the Protestant church would be leveled within their territory. Protestant mainline hierarchies in the late 1950, and early 1960, which had witnessed the life of a Calvinist Protestant movement, built the Protestant image of a new creed that preached tolerance and loyalty to the Church, desperately looking for a new denomination. Although the political movements had been building up for a long time, there was a Protestant political movement that thwarted the church and introduced a sense of impending doom specifically to a certain movement, a belief that while the Reagan government's attack on the Church made it Americans personally and at the center of society, and eventually dissolved and destroyed Protestant institutions. The church, with the exception of the United Methodist Episcopal Church (MCCA), heavily engaged in Latin American, Australian, and indigenous communities, is recognized as having one of the first Protestant hegemonies in the United States, following the formation of St. Mary Catholic Church in Louisville in 1963.

The circumstances of the resurgence of radicalism have led to a political movement that introduced a distinctly Protestant identity and spread feelings of insecurity and alienation that today plague converts to Christianity. This is unusually widespread within the broader American Catholic Church. Often associated with radical movements, such trends may signal a significant shift in the Catholic Church from the separation techniques throughout the history of the United States Church. The Southern Baptist Convention Conference, established its Center for the Restoration of Christ of God, in America. Later, it established an Adventist church in Texas. The Maclean Church professes and operates another Adventist church, the Church of America, which began forming one of the most centrist and largest denominations in the country in 1960. As a result of this, the largest denomination would soon be Holy Cross Church in North Carolina. Nick Sheymia, president of of Faith in America, which maintains a hard line between conservative and fundamentalist Christianity on the left, told the New Apostolic Arrangement Network that the church has always "walked in the overall shadow of the much broader church that attempted its rule and control of the nation."

Ol Howard, an educator and the author of the American Catholic Church: A Study in Truth, said, "The church seems inept at the end of the time. But to most Americans it

1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943

Text generated with 128 steps (1/3)

most countries, but in the UK, where computers are run by the as many as 15 people involved is more complicated. Many experts think that governments may need to overhaul their machines because they are unable to retain their influence because budgets are stretched too thin. In such a complicated situation, they say, public sector efforts to move ahead in technology, particularly climate change, would have to be put in real action.

Another major concern was a government effort to undermine the private sector in the construction sector and financial services industry. Ben Gove, chief executive of the House of Commons, said the government needed to address the role of the public service sector, and made explicit that the role of the labour force has become an effective vehicle to cut income from the poorest working people to the richest.

20%: the economic consequences of mass industrialisation Read more

In a joint statement with several coalition partners, the Conservative party also called for changes in the public sector to ensure the return on public investment, as well as the budget. Labour pushed back against some reforms, which involve modernising the benefits system for older people to make it easier for businesses to compete to win jobs. However, these reforms critics say apply only to the people who need them.

Ben Smith, the shadow environment minister with a record of climate change on the Labour side, expressed "general disappointment that this report shows significant shortcomings". In practice, the Commons report is not meant to advocate for such reforms. Labour would have hoped to represent a party with a strong alliance with existing policies and a government that has committed to investing in more workers overall.

Labour responded by saying: "We're pleased the MPs voted to put in place further policies to help the workforce and benefit everyone."

Labour paid for the narrow, decisive vote with a single vote. They say the findings show the government has missed the importance of its work for helping the working class within the economy, so that the enormous impact of automation benefits is linked with the exploitation and marginalisation of the working class.

The government has failed to recognise the fact that energy profits have been rising. Since 2008, bosses have tripled the ranks of the British electricity and gas companies, helping their annual profits rise by 3.1bn over the previous decade.

The biggest companies have made more profits over the decade before. This has created an incentive to make for more secure and sustainable energy security. But as well as bosses, they must collectively recognise what is happening,

1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997

Text generated with 128 steps (2/3)

many really exciting, non-independent artists, and all these indie bands, jazz and folk-rock, and rock-and-roll jazz bands, and you see a lot of new artists making these kinds of fantastic music. What do you think about these artists in Seattle?

M: I my whole life have never really wanted to do anything outside indie music, so I wanted to go way mainstream. It's much of a role model to try and spread the word. I think everyone gets invested in it so that people can use their music, but even back then, nobody had tried to bring these kinds of music to that community, and I saw that it was a good way you should find someone who's going to do that in the right way. This city can be an extremely helpful tool for developers to come up with the right media resources to develop quality music, and now there are so many indie studios over the country just attracting artists, and I think it's very helpful, and so there are a lot more open tools available.

S. Norman, what is probably the biggest impediment you put in terms of the music scene, and what's really important for you and I to bring your music to audiences in Portland or New York?

M: I would say the only really obstacle to me, it's the fact that I have to work with people in Seattle, and I try to involve them in it as much as possible because of where they actually come from. Many of the people we were talking to are basically part timers in the industry, and that I and I have had to work with, so many people are also a part of the community, among all of the audiences we've been meeting. And I really came out of this scene with a lot of excitement about it. And because I grew up growing up here in Seattle that was crowded and full of loud noise, it had just turned out to this place with so many people in this scene, that I found myself moving more to the music that I come from, at least to me, in a small, vibrant, and growing community, and it makes that sound more accessible to everyone. And what is really so encouraging is that while most of the people that we were talking to at events, they are also taking part in actions within the community, supporting our music through connections and sharing it with each other in our community and the city. Right now, there's really not a

1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051

Text generated with 128 steps (3/3)

to write this because it blew me out of my seat. I've been getting a few negative reviews in the past couple of days and I realize that people are much harder than ever to read them. Well, everything I've heard to this is overwhelming, and I'm happy to say that there am so much that I know that has really blown me out of my seat that this situation is much worse than anything I have told you. I want to be very direct to you, regarding even some of the things that you have heard. Feel free to let me know what you feel like - whether I have a personal concern or a sincere thing you would like to hear about this topic.

The response to this story is overwhelming support and love! I just want to reach out to my Facebook page to tell you what I I wish to do for you and really appreciate. I don't want to talk to anyone just to tell you how I'm very sorry. I want to tell you about all the things that I've been through so far and what I've done for you. I'm really sorry for everything that I have done to this community, and I am really appreciative for my support. Not to even go into detail the issue too much, please do head over to the next page in the post.

I ultimately don't want to tell you that this is your biggest issue - I don't assume that it help you much. Instead, I want to say that you should just get some things out of the way. I've encountered many other people just coming forward to this, and the fear of losing anything that I try to say about these issues can only interfere with such a process. However, now I like being out there trying to help the world, and I am so thankful that the community might help me too. Although I try to read this as about once a day, I just appreciate the amount of your feedback and hopefully I can build more positive ones to read by the evening.

Lastly, I want to inform everyone regarding this topic in case that other readers want to buy my blog here....So hopefully when I write something to my readers in order to help them out, I like your support. I want to support every reader so that I can talk to other readers for any information that can help me, to further tell my story.

C: I still have the horrible feeling that I