arXiv:2605.07013v1 [cs.CL] 7 May 2026

Towards Closing the Autoregressive Gap in Language Modeling via Entropy-Gated Continuous Bitstream Diffusion

Georgios Batzolis[1*] **Mark Girolami**[1,2] **Luca Ambrogioni**[3]

Abstract

Diffusion language models (DLMs) promise parallel, order-agnostic generation, but on standard benchmarks they have historically lagged behind autoregressive models in sample quality and diversity. Recent continuous flow and diffusion approaches over token embeddings have narrowed this gap, suggesting continuous state spaces are highly effective for language. In this work, we further close the autoregressive gap by modeling text as a continuous diffusion process over fixed-width binary bitstreams. Our approach represents semantic tokens as analog bit sequences and utilizes a matched-filter residual parameterization to isolate contextual learning from analytic independent-bit posteriors. Crucially, we adopt a stochastic sampler that applies Langevin-type corrections gated by the entropy-rate profile, automatically concentrating stochasticity in high-information regions while remaining nearly deterministic elsewhere. On the One Billion Word Benchmark (LM1B), our 130M-parameter bitstream model reaches a generative perplexity ($\mathrm{Gen.\ PPL}$) of 59.76 at matched real-data entropy (4.31) using 256 neural function evaluations (NFEs), decisively outperforming prior DLM baselines and reaching the autoregressive reference. On OpenWebText (OWT), our stochastic sampler establishes a new continuous-DLM Pareto frontier, achieving $\mathrm{Gen.\ PPL} = 27.06$ at an entropy of 5.26 using $4\times$ fewer steps than previous 1024-NFE baselines. As an additional architectural benefit, bitstream diffusion removes the $\mathcal{O}(V)$ vocabulary scaling bottleneck shared by standard DLMs. By predicting $\mathcal{O}(\log V)$ bitwise logits via semantic bit-patching, our model yields a reduced memory footprint and higher throughput, demonstrating a scalable paradigm for language generation as vocabulary sizes grow.

1 Introduction

Autoregressive language models dominate modern text generation because they define a simple factorization and scale reliably. Their main algorithmic limitation is equally clear: generation is inherently sequential. Continuous diffusion models have become a standard framework for high-dimensional image and video generation [Ho et al., 2020, Song et al., 2021, Karras et al., 2022b]. Diffusion language models (DLMs) offer a compelling alternative compute profile, enabling parallel refinement of all positions, arbitrary infilling, and a tunable compute, quality tradeoff. However, unlocking these advantages has proven difficult. For language, diffusion models have historically suffered from a persistent quality, diversity gap compared to autoregressive baselines: they often yield weak sample quality, or achieve artificially low generative perplexity (GenPPL) only by over-generating safe, frequent tokens, thereby collapsing sample entropy.

*Corresponding author: gb511@cam.ac.uk

Affiliations. [1] Department of Engineering, University of Cambridge. [2] Devotion AI Labs Ltd. [3] Donders Institute for Brain, Cognition, and Behaviour, Radboud University.

Preprint.

A common explanation for this gap has been that language naturally demands discrete diffusion. This view has been challenged by recent continuous flow and diffusion language models over one-hot token embeddings [Roos et al., 2026, Lee et al., 2026, Chen et al., 2026]. These methods demonstrate that continuous models can closely rival strong discrete baselines, suggesting the primary bottleneck is not continuity itself, but the interaction between the state representation, objective, and sampler design.

This paper pushes the continuous diffusion paradigm further towards closing the gap to autoregression and establishing a new state-of-the-art DLM Pareto frontier. Instead of diffusing over token embeddings, following the approach introduced in [Chen et al., 2023] for multimodal generation, we diffuse over *bitstreams*. A sequence of T semantic tokens is encoded as a sequence of fixed-width binary bits, embedded in continuous space, and an EDM-style denoiser is trained to recover the bits from Gaussian corruption. Because the posterior of an isolated bit under Gaussian corruption has a known closed form, we introduce a *matched-filter residual parameterization*: the network analytically computes the independent-bit posterior and focuses its capacity entirely on predicting the contextual residual.

The critical role of stochastic sampling. While the bitstream representation provides a strong foundation, our largest empirical gain comes from the sampler. The deterministic probability-flow sampler is already competitive with recent continuous DLMs, but it is over-contractive: it can obtain good GenPPL by undershooting real-data token entropy. We show that EDM-style stochastic churn corrects this behavior. When applied on an entropy-rate sampling grid, full-band churn improves the GenPPL, entropy frontier without changing the trained model or increasing the NFE budget. Intuitively, the entropy-rate grid concentrates solver resolution where bit uncertainty is resolved, and the same grid also makes the effective stochastic correction strongest in that region. We formalize this continuous-time interpretation in Section D.

Empirical status. As shown in Figure 2, full-band stochasticity is the primary driver of our improved GenPPL–entropy Pareto frontier. On LM1B, our deterministic 256-NFE sampler achieves $\mathrm{Gen.\ PPL} = 82.90$ at entropy 4.30, remaining competitive with prior continuous DLMs. Applying full-band stochasticity shifts the frontier to $\mathrm{Gen.\ PPL} = \mathbf{59.76}$ at matched real-data entropy 4.31. This improves over LangFlow's reported $\mathrm{Gen.\ PPL} = 92.24$ by more than 30 points and reaches the autoregressive quality regime, outperforming the AR baseline's $\mathrm{Gen.\ PPL} = 66.70$ while preserving healthy token entropy. On OpenWebText, our method similarly shifts the high-quality region of the continuous-DLM frontier: the deterministic sampler obtains $\mathrm{Gen.\ PPL} = 46.32$ at entropy 5.13, while full-band stochasticity improves this to $\mathrm{Gen.\ PPL} = \mathbf{27.06}$ at entropy 5.26. This outperforms LangFlow's $\mathrm{Gen.\ PPL} = 36.53$ at slightly lower entropy 5.25, using 256 rather than 1024 NFEs.

Additional computational benefits. As an important structural benefit, bitstream diffusion eliminates the vocabulary-sized output bottleneck shared by almost all DLMs. Simplex, one-hot, and discrete-transition models fundamentally require $\mathcal{O}(V)$ output parameterizations per token. By patching $m = \lceil \log_2 V \rceil$ bits into a single sequence element, our sequence diffusion transformer (SDT) preserves the semantic context length T while replacing the dense vocabulary classifier with a compact $\mathcal{O}(\log V)$ bitwise head. At LM1B scale, this end-to-end simplification yields a $1.6\times$ reduction in peak memory and $2.3\times$ higher training throughput compared to a matched token-space model. Crucially, as context length and vocabulary size grow to OpenWebText scale, these practical gains expand to a $3.3\times$ training throughput speedup and a $19\times$ reduction in generation memory footprint. This demonstrates that the architectural advantage becomes strictly more beneficial as sequences and vocabularies grow.

2 Related Work

Discrete diffusion language models. Discrete DLMs define Markov corruption processes directly on tokens or masks. D3PM [Austin et al., 2021] introduced structured discrete denoising diffusion, and SEDD [Lou et al., 2023] framed discrete diffusion through ratio estimation and score entropy. MDLM [Sahoo et al., 2024] showed that masked diffusion language modeling can be substantially strengthened with a simplified objective and improved training recipe. Duo [Sahoo et al., 2025] further connected uniform-state discrete diffusion to Gaussian diffusion and improved training and sampling. These methods are strong baselines for non-autoregressive language modeling, but they remain tied to categorical transition kernels and often require separate designs for masking, absorbing

states, or uniform corruption. Our approach instead uses a continuous Gaussian process in bit space, while still decoding to valid discrete tokens at the end.

Continuous diffusion for categorical data. Several works have argued that continuous diffusion can be useful for discrete variables when the geometry is handled carefully. Analog Bits [Chen et al., 2023] represents discrete variables as binary bits and trains continuous diffusion models on analog versions of those bits; it also introduced self-conditioning and asymmetric time intervals, both of which influence our design. CDCD [Dieleman et al., 2022] models categorical data with continuous-time and continuous-state diffusion, and emphasizes that categorical geometry should inform the parameterization and objective. Riemannian Diffusion Language Models (RDLM) model categorical distributions using statistical-manifold geometry [Jo and Hwang, 2025]. CANDI [Pynadath et al., 2025] explores hybrid discrete–continuous diffusion. Our contribution is closest in representation to Analog Bits but differs in scale, architecture, language-focused evaluation, matched-filter residual parameterization, entropy-rate scheduling, and entropy-band stochastic sampling. Recently, FLM/FMLM [Lee et al., 2026] and LangFlow [Chen et al., 2026] provide the strongest recent evidence that continuous language models can rival discrete DLMs. FLM performs continuous denoising over one-hot token encodings and trains with cross-entropy objectives; FMLM distills the flow map for few-step generation. LangFlow connects embedding-space DLMs to flow matching via Bregman divergence, introduces an ODE-based NLL bound, proposes an information-uniform noise-scheduling principle, and shows that self-conditioning improves continuous DLMs. Our method validates the same broad thesis, continuous generative modeling can work well for language, but takes a different route. We use fixed-width bitstreams instead of one-hot token embeddings, binary score matching instead of token cross-entropy as the default objective, and an analytic matched-filter decomposition of bit posterior logits.

3 Method

A text example is tokenized into T semantic tokens and encoded as a fixed-width binary sequence

$$x_0 \in \{0,1\}^S, \qquad S = Tm,$$

where m is the number of bits per token or code token. For LM1B we use $T = 128$ and $m = 15$; for OpenWebText we use a fixed 1024-code-token representation with $m = 16$. The model never generates autoregressively: it initializes an analog bit vector from Gaussian noise and denoises all positions in parallel. Figure 1 gives an end-to-end schematic of the architecture.

3.1 Architecture schematic

As forward process, we use a variance-exploding Gaussian corruption model

$$x_\sigma = x_0 + \sigma\epsilon, \qquad \epsilon \sim \mathcal{N}(0, I_S), \qquad \sigma \in [\sigma_{\min}, \sigma_{\max}]. \tag{1}$$

Bits are represented as $0/1$ values, with data center $c = 1/2$ and $\sigma_{\text{data}} = 1/2$. The denoiser predicts bitwise clean probabilities

$$D_\theta(x_\sigma, \sigma) \in (0,1)^S.$$

These probabilities induce a continuous score estimate through the Gaussian posterior-mean identity

$$s_\theta(x_\sigma, \sigma) = \frac{D_\theta(x_\sigma, \sigma) - x_\sigma}{\sigma^2}. \tag{2}$$

Thus the same output defines discrete bit probabilities and provides the score used by the continuous sampler.

3.2 Matched-filter residual parameterization

A central modeling choice is to separate local Gaussian bit denoising from contextual language modeling. For an isolated bit with uniform prior $x_0 \sim \mathrm{Bern}(1/2)$ and observation $x = x_0 + \sigma\epsilon$, the posterior logit is available analytically:

$$\ell_{\text{ind}}(x, \sigma) = \log \frac{p(x_0 = 1 \mid x, \sigma)}{p(x_0 = 0 \mid x, \sigma)} = \frac{x - \frac{1}{2}}{\sigma^2}. \tag{3}$$

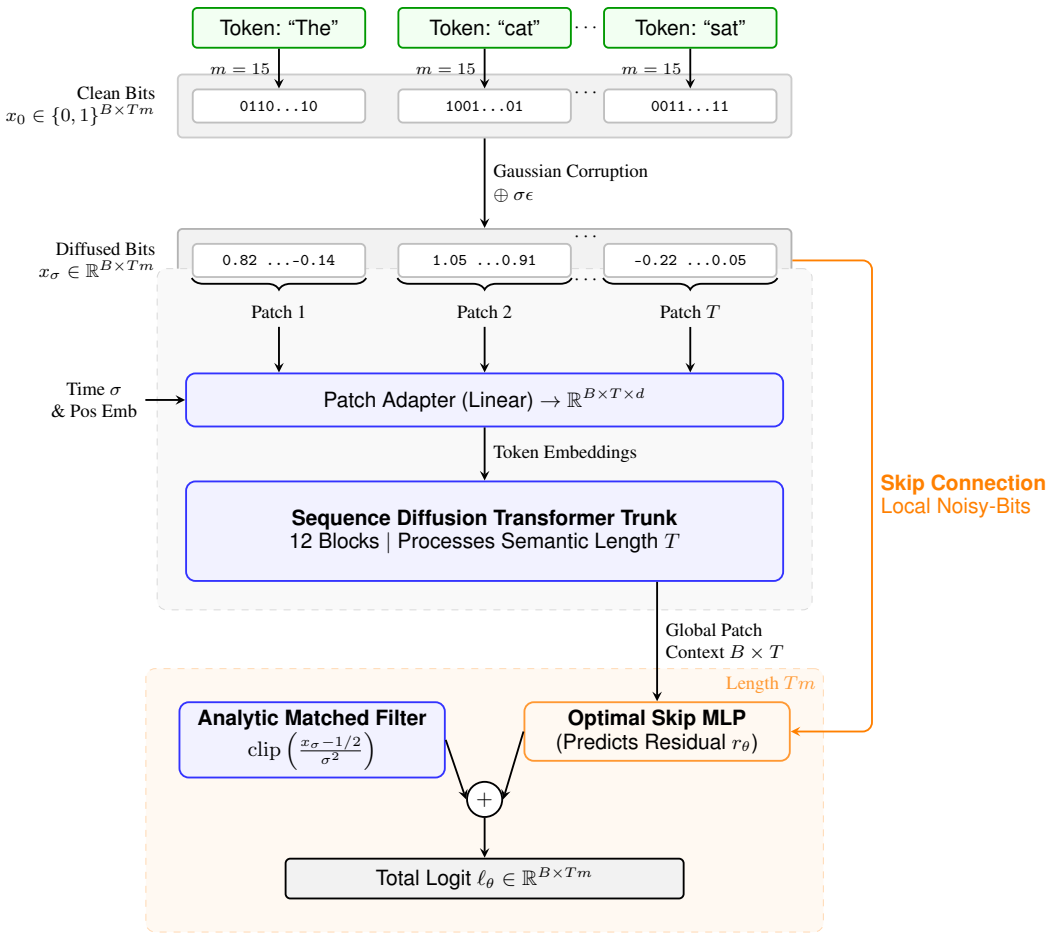

Figure 1: **End-to-end Bitstream Diffusion architecture.** Text is encoded into a Tm-length bitstream and corrupted via Gaussian noise. The diffused bits x_σ are patched into semantic groups and processed by a length-T Transformer trunk. In parallel with the trunk, a skip connection routes the raw noisy bits directly to the head, where a contextual residual is combined with the analytic matched filter to produce bit-level logits.

Language modeling, however, requires contextual dependencies between bits, tokens, and distant positions. We therefore ask the network to predict only a contextual residual logit r_θ, and add it to the analytic matched filter:

$$\ell_\theta(x_\sigma, \sigma) = r_\theta(x_\sigma, \sigma, x_{\text{sc}}) + \text{clip}\left(\frac{x_\sigma - \frac{1}{2}}{\sigma^2}, -C, C\right), \qquad D_\theta = \text{sigmoid}(\ell_\theta). \tag{4}$$

Here x_{sc} is the self-conditioning input and $C = 30$ in our runs. The analytic term handles local Gaussian bit denoising, while the network focuses on contextual dependencies between bits and tokens. This is the strongest training-side component in Section 4.4: removing it worsens LM1B GenPPL by over 22 points under a fixed deterministic sampler. A derivation of Equation (3) is given in Section A.1.

3.3 Training objective and self-conditioning

We train with binary score interpolation, implemented as a weighted denoising MSE between posterior probabilities and clean bits:

$$\mathcal{L}_{\text{SM}}(\theta) = \mathbb{E}_{x_0, \epsilon, \sigma}\left[w(\sigma)\, \frac{1}{S} \sum_{i=1}^{S} \left(D_{\theta,i}(x_\sigma, \sigma) - x_{0,i}\right)^2\right], \tag{5}$$

with EDM weighting

$$w(\sigma) = \frac{\sigma^2 + \sigma_{\text{data}}^2}{\sigma^2 \sigma_{\text{data}}^2}. \tag{6}$$

Recent continuous token-space language models such as FLM and LangFlow use token-level objectives based on cross-entropy or related divergences [Lee et al., 2026, Chen et al., 2026]. In our binary setting, score matching is natural because D_θ directly defines the continuous score in Equation (2).

We use self-conditioning by default. With probability $p_{\text{sc}} = 0.5$ during training, we perform an auxiliary no-gradient denoising pass and feed the detached posterior probabilities back into the main denoiser as x_{sc}. At sampling time, we use carry-mode self-conditioning, where the previous denoised prediction is fed into the next denoising step. This follows prior continuous bitstream and language diffusion practice [Chen et al., 2023, 2026].

3.4 Entropy-rate adaptive noise allocation

The standard EDM log-normal noise distribution is a strong generic choice, but bitstream data can have information profiles that depend on the tokenizer, code representation, redundancy, and dataset. We therefore adapt the information-guided noise allocation principle of Stancevic et al. [2025] and Raya et al. [2026]. Let

$$h_{\log}(\sigma) := \frac{d}{d \log \sigma} H(x_0 \mid x_\sigma)$$

denote the conditional entropy rate per unit log-noise. This quantity identifies where the forward process destroys information about the clean bitstream, or equivalently where reverse denoising must resolve the most information.

In practice, we estimate this profile online from the unweighted denoising errors already computed during training, obtaining a proxy $\widehat{h}_{\log}(\sigma)$. This defines a normalized schedule density over $u = \log \sigma$,

$$\pi_\alpha(u) \propto g(e^u; c, n)\, \widehat{h}_{\log}(e^u)^\alpha, \qquad g(\sigma; c, n) = \frac{\sigma^n}{\sigma^n + c^n}. \tag{7}$$

In the final runs we use the sqrt-rate setting $\alpha = 1/2$, with $c = 0.1, n = 3$. Training begins from the EDM log-normal distribution, then transitions to the entropy-rate distribution after warmup.

3.5 Sampling: deterministic flow and stochastic correction

The deterministic sampler integrates the probability-flow ODE induced by Equation (2). Writing $D_\theta(x, \sigma)$ for the denoised bit probabilities,

$$\frac{dx}{d\sigma} = -\sigma s_\theta(x, \sigma) = \frac{x - D_\theta(x, \sigma)}{\sigma}. \tag{8}$$

We use DDIM-style sampling in the main experiments and support Heun correction in the codebase. A key component of our method is the use of entropic time warping, introduced in [Dieleman et al., 2022] for softmax models and generalized in [Stancevic et al., 2025] to arbitrary continuous diffusion models. Unless stated otherwise, reported samples use the entropy-rate grid and carry-mode self-conditioning.

Entropy-gated stochastic churn. We adopt an EDM-style stochastic churn [Karras et al., 2022a]. Before a deterministic step from σ_i to σ_{i+1}, the sampler temporarily moves to a noisier level

$$\hat{\sigma}_i = (1 + \gamma_i)\sigma_i,$$

injects the corresponding Gaussian noise, and then applies the deterministic update from $(\hat{x}_i, \hat{\sigma}_i)$ to σ_{i+1}. For the main stochastic results, we use full-band raw churn, $\gamma_i = S_{\text{churn}}/N$, with $S_{\text{noise}} = 1.003$. The key point is that raw churn and effective continuous-time stochasticity are not the same quantity. In the small-step regime, the churn update behaves like a probability-flow step plus a local Langevin correction with effective reverse-SDE strength

$$\lambda_i \approx \frac{\gamma_i \sigma_i}{\sigma_i - \sigma_{i+1}}. \tag{9}$$

Thus the sigma grid determines where a fixed raw churn budget becomes a large stochastic correction. For a grid uniform in the CDF of $\pi_\alpha(u)$, local spacing satisfies

$$\sigma_i - \sigma_{i+1} \approx \frac{\sigma_i}{N\pi_\alpha(\log \sigma_i)}.$$

Substituting $\gamma_i = S_{\text{churn}}/N$ into Equation (9) gives

$$\lambda_{\text{ent}}(\sigma_i) \approx S_{\text{churn}}\, \pi_\alpha(\log \sigma_i). \tag{10}$$

The entropy-rate grid therefore has a dual role: it allocates deterministic solver points to information-active noise levels and automatically converts constant full-band raw churn into an information-adaptive Langevin correction. This is the mechanism behind the title phrase "entropy-gated." The full derivation, finite-step caveat, and implementation details are given in Section D. We also support the asymmetric time-interval label shift of Analog Bits [Chen et al., 2023], which evaluates the denoiser at a slightly noisier time label and can help in some low-NFE deterministic regimes. Since its effect is weaker once stochastic churn is enabled, all main 256-NFE stochastic results use $\eta = 0$; implementation details and ablations are given in Sections B.3 and E.

3.6 Sequence diffusion transformer and vocabulary-boundary scaling

The denoiser is a sequence diffusion transformer operating at semantic length T, not bit length Tm. The m analog bits of each token are grouped into one patch and projected to a token embedding. The trunk matches the scale of recent continuous DLM baselines: 12 Transformer blocks, hidden width 768, 12 attention heads, feed-forward width 3072, RoPE, AdaLN-zero time conditioning, SwiGLU activations, FlashAttention/SDPA kernels when available, dropout 0.1, and BF16 training. The output head expands each trunk token back to its m bits. Each bit receives both global context from the Transformer patch token and a local skip path carrying the current noisy bit and self-conditioning features. The final bit logit is the sum of the contextual residual and the analytic matched-filter logit in Equation (4). A schematic and full adapter details are given in Figure 1 and Section B.2. This representation removes the vocabulary-sized output boundary used by most DLMs. Simplex, one-hot, embedding-space, and discrete token-state DLMs typically require a $\mathcal{O}(TV)$ output/loss interface. By contrast, Bitstream Diffusion predicts $m = \lceil \log_2 V \rceil$ bit logits per semantic token, replacing this boundary with $\mathcal{O}(T \log V)$ logits while preserving trunk length T. For LM1B, $V = 30{,}522$ and $m = 15$, giving a raw output-logit reduction of $V/m \approx 2035\times$. For OpenWebText, $V = 65{,}536$ and $m = 16$, this reduction grows to $V/m \approx 4096\times$. We measure the practical end-to-end memory and throughput impact in Section 4.3, with structural comparisons in Section A.2.

4 Experiments

We evaluate Bitstream Diffusion on LM1B and OpenWebText (OWT). Because standard DLMs can artificially reduce generative perplexity (GenPPL) by collapsing toward safe, frequent tokens, we evaluate GenPPL jointly with token-frequency entropy. This joint view is critical: we use GenPPL–entropy frontiers to robustly compare configurations and ensure improvements are not artifacts of vocabulary collapse, as suggested in [Pynadath et al., 2026].

4.1 Experimental Setup

Datasets. We evaluate on LM1B packed into 128-token blocks ($m = 15$, `bert-base-uncased`, 1M steps) and OWT using 1024-token blocks (750K steps). For OWT, we retain GPT-2 tokenization for evaluation but train through a reversible `gpt2id_bpe16` codec that maps GPT-2 ID sequences to dense 16-bit code tokens; all metrics are computed only after inverse decoding to GPT-2 IDs/text. Both datasets use a 12-layer SDT trunk (width 768) trained with AdamW, binary score matching, EDM loss weighting, and an entropy-rate noise schedule; details are in Sections A.3 and B.

Metrics & Sampling. GenPPL is computed with external `gpt2-large` on 1024 generated samples. We also report entropy: mean per-sample token unigram entropy in nats, computed on evaluation token IDs (`bert-base-uncased` for LM1B and inverse-decoded GPT-2 IDs for OWT), ensuring comparability with GPT-2-tokenized baselines rather than our intermediate code vocabulary. We sample with either the deterministic probability-flow ODE or full-band stochastic EDM churn ($S_{\text{noise}} = 1.003$); unless stated otherwise, results use 256 NFEs on the entropy-rate grid.

Method	LM1B		OWT	
	GenPPL ↓	Entropy	GenPPL ↓	Entropy
Reference				
Real data	53.06	4.31	15.06	5.44
Autoregressive				
AR Transformer	66.70	4.32	35.90	5.58
Diffusion				
RDLM	268.21	4.33	-	-
CANDI	120.99	4.35	143.13	5.71
MDLM	109.21	4.32	105.15	5.63
Duo	98.14	4.31	77.69	5.55
FLM	96.91	4.29	62.23	5.33
LangFlow (128/1024 NFE)	92.24	4.31	36.53	5.25
Ours (256 NFE)				
Deterministic	82.90 ± 1.11	$4.30 \pm < 0.01$	46.32 ± 0.93	$5.13 \pm < 0.01$
Stochastic	$\mathbf{59.76 \pm 0.57}$	$\mathbf{4.31 \pm < 0.01}$	$\mathbf{27.06 \pm 0.57}$	$\mathbf{5.26 \pm 0.01}$

Table 1: **Main single-point comparison.** We report one representative operating point per external method and compare both generative perplexity and token-frequency entropy. For Bitstream Diffusion, we show both deterministic and full-band stochastic 256-NFE samplers using the same trained checkpoint. To ensure robustness, our results are reported as the mean $\pm$ standard deviation across 10 random sampling seeds. The deterministic row isolates the base probability-flow sampler, while the stochastic row shows the effect of adding full-band churn on the entropy-rate grid at the same NFE budget. The real-data row provides the target values for each dataset. On LM1B, the stochastic sampler matches the real-data entropy regime while substantially improving over reported diffusion baselines. On OWT, it improves over LangFlow at matched or slightly higher reported entropy using four times fewer NFEs; additional points on the stochastic sampler frontier are shown in Figure 2.

4.2 Main results

Table 1 and Figure 2 summarize the main comparison. To keep the table focused, we report one representative operating point per external method. For Bitstream Diffusion, we report both deterministic and stochastic 256-NFE samplers using the same trained checkpoint and architecture. The deterministic row isolates the base probability-flow sampler, while the stochastic row shows the effect of adding full-band churn on the entropy-rate grid at the same NFE budget. This distinction is central to the result: the deterministic sampler is already competitive, but it is consistently more contractive, whereas stochastic churn on the entropy-rate grid improves the GenPPL–entropy trade-off without changing the model or increasing NFE.

LM1B. As shown in Table 1 and Figure 2, the deterministic 256-NFE sampler is already competitive, but adding full-band stochasticity fundamentally shifts the Pareto frontier. Our stochastic sampler obtains GenPPL 59.76 at entropy 4.31, exactly matching the real-data entropy regime. This substantially outperforms all prior diffusion baselines reported (e.g., improving > 30 GenPPL points over LangFlow) and successfully reaches the text quality regime of the autoregressive (AR) baseline ($\mathrm{Gen.\ PPL} = 66.70$) without sacrificing token diversity.

OpenWebText. On OWT, the deterministic 256-NFE sampler is strong in GenPPL but visibly under-entropic relative to the real data (5.44), indicating over-contraction. Adding full-band stochasticity corrects this: our selected 256-NFE sampler obtains GenPPL 27.06 at entropy 5.26, directly improving over LangFlow's 1024-NFE reported point using four times fewer denoiser evaluations. Increasing the churn budget exposes a useful quality–diversity frontier (Figure 2), reaching $\mathrm{Gen.\ PPL} = 34.35$ at entropy 5.32, matching AR baseline quality while pushing closer to real-data entropy. *Uncurated generated samples for both datasets are provided in Section G.*

4.3 Computational and memory efficiency

Standard continuous token-space diffusion models require a vocabulary-wide output boundary, forming dense $B \times T \times V$ logits at every denoising evaluation. As context length T and vocabulary

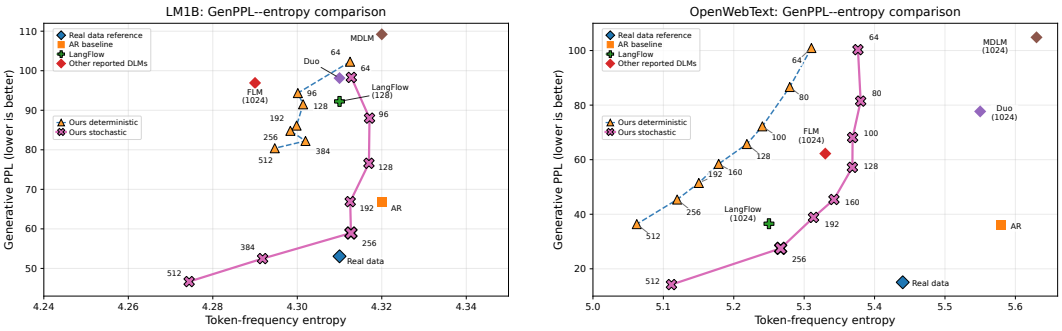

(a) LM1B. (b) OWT.

Figure 2: **GenPPL–entropy comparison.** Generative perplexity (lower is better) versus token-frequency entropy on LM1B and OWT. The deterministic curve isolates our base probability-flow sampler. The highlighted stochastic trajectory represents our optimized Pareto frontier: for each NFE budget, we swept the stochastic churn parameter across eight values and selected the operating point that gives a good balance between generative perplexity and entropy. On both datasets, Bitstream Diffusion establishes a strictly superior operating frontier. On LM1B, our 256-NFE sampler decisively outperforms all prior diffusion baselines while operating exactly in the real-data entropy regime. On OWT, the stochastic frontier consistently dominates previous models, achieving better generative perplexity and entropy than prior 1024-NFE baselines using only a fraction of the compute.

size V grow, this becomes a substantial memory and throughput bottleneck. Bitstream Diffusion replaces this boundary with a patched bitstream head that forms only $B \times T \times \lceil \log_2 V \rceil$ logits, while preserving the semantic sequence length T. To isolate this systems effect, we profile matched token-space and bitstream models using synthetic batches, removing dataloader, tokenization, disk-I/O, decoding, and host-device transfer effects. Both models use the same 12-layer Sequence Diffusion Transformer trunk with width $d = 768$, BF16 mixed precision, and identical profiling settings. Full protocols and batch-size sweeps are given in Section F.

Dataset scale	Mode	Batch	Peak VRAM reduction	Throughput speedup
LM1B ($T = 128$, $V = 30$k)	Train	512	1.64×	2.31×
LM1B ($T = 128$, $V = 30$k)	Generate	1024	10.02×	2.00×
OWT ($T = 1024$, $V = 65$k)	Train	64	2.46×	3.29×
OWT ($T = 1024$, $V = 65$k)	Generate	64	19.25×	2.65×

Table 2: **End-to-end systems impact of the bitstream boundary.** We compare matched SDT trunks using either a vocabulary-wide token boundary or the patched bitstream boundary. For each dataset and mode, we report metrics at the largest common batch size that fits both models on the profiling GPU. For training, the batch column denotes the per-step optimization batch size; for generation, it denotes the generation batch. Generation throughput is measured in completed semantic tokens per second and includes the full 128-NFE denoising trajectory.

Table 2 shows that the compact bitstream boundary produces substantial end-to-end gains. During training, the shared Transformer trunk, optimizer state, and backward activations dilute the raw output-logit reduction, but the practical gains remain large: Bitstream Diffusion reduces peak memory by 1.64× and improves throughput by 2.31× on LM1B, and by 2.46× and 3.29× on OWT. This has a direct practical consequence: at larger contexts and vocabularies, the token-space baseline reaches the single-GPU memory limit at much smaller per-device batches, forcing either smaller effective batches or additional GPUs to distribute the same workload. By contrast, the bitstream boundary remains feasible at substantially larger per-device batches, improving hardware utilization and reducing the need for memory-driven parallelization.

The advantage is larger during generation, where no optimizer state or backward activations are stored, making the output boundary a much larger fraction of the active memory footprint. At LM1B scale, Bitstream Diffusion reduces generation memory from 52.55 GiB to 5.25 GiB at batch size 1024, while doubling throughput. At OpenWebText scale, where the token boundary has $T = 1024$ and

$V = 65{,}536$, Bitstream Diffusion reduces generation memory by $19.25\times$ and improves throughput by $2.65\times$ at the largest common generation batch, allowing substantially larger batches before reaching the same hardware memory limit.

4.4 Ablations

Table 3 isolates the main training-side choices on LM1B. Self-conditioning is enabled in all variants [Chen et al., 2023, 2026]. All variants use the same 256-NFE deterministic sampler and the same 256-NFE stochastic churn sweep, so differences reflect the learned model rather than per-variant sampler tuning.

Variant	Schedule + Objective + Filter	Deterministic		Stochastic	
		GenPPL ↓	Entropy	GenPPL ↓	Entropy
Reference					
Real data	–	53.06	4.31	53.06	4.31
Ablations					
Config A (Base)	entropy-rate + SM + matched filter	85.37 ± 0.60	4.30	59.93 ± 0.90	4.31
Config B	EDM schedule + SM + matched filter	81.95 ± 0.75	4.29	54.04 ± 0.65	4.29
Config C	entropy-rate + CE + matched filter	91.93 ± 1.25	4.29	54.68 ± 0.74	4.27
Config D	entropy-rate + SM + no matched-filter prior	105.96 ± 1.32	4.34	66.18 ± 1.13	4.33

Table 3: **Training ablations on LM1B.** All variants use 256 NFEs. Stochastic results are aggregated over 10 sampling seeds and the same target-churn sweep $\gamma \in [0.11, 0.20]$. Entropy standard deviations are omitted because they are negligible (< 0.01). Config A is the default protocol used in the main experiments.

The matched-filter residual is the clearest training-side component. Removing the analytic independent-bit posterior substantially worsens both deterministic and stochastic GenPPL, confirming that the closed-form local Gaussian denoiser is a useful inductive bias: the network need not relearn isolated-bit denoising and can instead focus on contextual dependencies.

The objective comparison also supports binary score matching. Replacing score matching with cross-entropy worsens the deterministic sampler and, under stochastic sampling, reaches lower GenPPL only by undershooting real-data entropy more strongly. This is undesirable for our evaluation setting, where GenPPL must be interpreted jointly with token-frequency entropy.

The schedule comparison is a quality–diversity trade-off rather than a clear win for either schedule. EDM gives lower GenPPL, but consistently operates in a more contractive, lower-entropy regime. The entropy-rate schedule gives strong GenPPL while matching the real-data entropy under stochastic sampling. We therefore use it as the default because it is both safer empirically and data-driven: it adapts noise allocation to the measured bitstream information profile across datasets, tokenizers, encodings, or redundant representations.

Why not discrete diffusion on bits? We also tested a SEDD-style discrete diffusion baseline [Lou et al., 2024] directly on LM1B bitstreams, preserving the same $\mathcal{O}(\log V)$ boundary. It substantially underperforms the matched token-level SEDD baseline: at 256 NFEs, GenPPL is 285.09 versus 126.28, at comparable entropy (Section C). This suggests that the compact bit boundary alone is not enough to make bit-level discrete diffusion competitive in this setting.

Churn controls the quality–diversity frontier. Figure 3 isolates stochastic churn on OWT at fixed NFE $= 256$ and $\eta = 0$. The deterministic sampler is over-contractive: it undershoots real-data entropy and gives worse GenPPL than the best stochastic settings. Moderate full-band churn improves both GenPPL and entropy, moving samples into a healthier quality–diversity regime. Increasing churn further continues to raise entropy, but eventually worsens GenPPL. Thus S_{churn} is best understood as a test-time frontier knob: moderate stochasticity corrects premature contraction, while excessive stochasticity injects too much uncertainty.

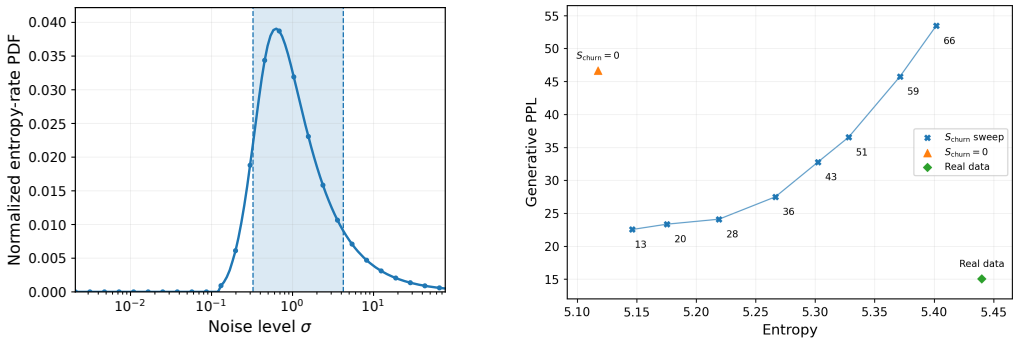

(a) Entropy-rate profile.

(b) Churn sweep at fixed 256 NFEs.

Figure 3: **Stochastic churn controls the OWT quality–diversity frontier. Left:** Normalized entropy-rate profile over noise level σ; the shaded region marks the broad entropy-active region used for this diagnostic. **Right:** At fixed NFE $= 256$ and $\eta = 0$, moderate full-band churn improves both GenPPL and entropy relative to deterministic sampling, while excessive churn increases entropy at the cost of GenPPL.

Sampler-design ablations. Additional ablations in Section E support the default sampler choices. The entropy-rate grid improves deterministic sampling relative to the Karras grid and makes stochastic churn substantially more effective (Figure 4). This indicates that stochasticity is not a generic benefit by itself: it is most useful when the sampling grid aligns both solver resolution and stochastic correction with the information-active part of the trajectory. The continuous-time interpretation in Section D explains this behavior, since the effective Langevin strength depends on local grid spacing as well as the raw churn parameter. The stochasticity-window ablation further shows that the correction should be broad rather than hand-localized (Figure 5). Narrow entropy-CDF windows are sensitive to their location, whereas broad windows are consistently stronger; the full entropy-supported range $q \in [0, 1]$ gives the best GenPPL in this sweep, with $q \in [0.1, 0.9]$ close behind. Thus the useful stochastic correction appears to act over a broad information-active region, rather than through a single narrowly localized entropy band.

5 Discussion

This work demonstrates that continuous diffusion over semantic bitstreams can substantially narrow the gap between diffusion language models and autoregressive generation while also improving computational scaling properties. Our results suggest that the key challenge for continuous language diffusion is not continuity itself, but the interaction between representation, parameterization, and sampling. By replacing token-space diffusion with fixed-width semantic bitstreams, the model benefits from both structured inductive biases and a more favorable $O(\log V)$ output scaling. A central finding is the importance of entropy-gated stochastic sampling. Deterministic probability-flow sampling is already competitive, but tends to be over-contractive and under-entropic. Adding stochastic churn on the entropy-rate grid consistently improves the quality–diversity trade-off, suggesting that stochasticity is most effective when concentrated in information-active regions of the reverse process. The bitstream formulation may become especially valuable for future multilingual and multimodal models. Large unified vocabularies spanning many languages, image tokens, audio units, or multimodal semantic codes can become prohibitively expensive for standard diffusion models with $O(V)$ output heads. In contrast, bitstream diffusion scales only logarithmically with vocabulary size, potentially enabling much larger and more flexible vocabularies without the same output-boundary bottleneck. This suggests a promising direction toward unified generative architectures operating across heterogeneous modalities and languages.

6 Limitations

The experiments in this work are limited to approximately 130M-parameter models on LM1B and OpenWebText. Further research is needed to show whether the same gains persist at substantially

larger scales comparable to modern frontier models. Finally, diffusion language models remain more computationally expensive at inference time than autoregressive models due to iterative denoising. Further advances in caching, distillation, accelerated samplers, or consistency-based generation may be necessary to make diffusion-based language generation practical in latency-sensitive settings.

References

Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. In *Advances in Neural Information Processing Systems*, 2021.

Ting Chen, Ruixiang Zhang, and Geoffrey Hinton. Analog bits: Generating discrete data using diffusion models with self-conditioning. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=3itjR9QxFw.

Yuxin Chen, Chumeng Liang, Hangke Sui, Ruihan Guo, Chaoran Cheng, Jiaxuan You, and Ge Liu. Langflow: Continuous diffusion rivals discrete in language modeling, 2026. URL https://arxiv.org/abs/2604.11748.

Sander Dieleman, Laurent Sartran, Arman Roshannai, Nikolay Savinov, Yaroslav Ganin, Pierre H. Richemond, Arnaud Doucet, Robin Strudel, Chris Dyer, Conor Durkan, Curtis Hawthorne, Rémi Leblond, Will Grathwohl, and Jonas Adler. Continuous diffusion for categorical data, 2022. URL https://arxiv.org/abs/2211.15089.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020.

Jaehyeong Jo and Sung Ju Hwang. Continuous diffusion model for language modeling, 2025. URL https://arxiv.org/abs/2502.11564.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *arXiv*, 2022a. doi: 10.48550/arxiv.2206.00364.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Advances in Neural Information Processing Systems*, 2022b.

Chanhyuk Lee, Jaehoon Yoo, Manan Agarwal, Sheel Shah, Jerry Huang, Aditi Raghunathan, Seunghoon Hong, Nicholas M. Boffi, and Jinwoo Kim. One-step language modeling via continuous denoising, 2026. URL https://arxiv.org/abs/2602.16813.

Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution, 2023. URL https://arxiv.org/abs/2310.16834.

Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 32819–32848. PMLR, 2024.

Patrick Pynadath, Jiaxin Shi, and Ruqi Zhang. Candi: Hybrid discrete-continuous diffusion models, 2025. URL https://arxiv.org/abs/2510.22510.

Patrick Pynadath, Jiaxin Shi, and Ruqi Zhang. Generative frontiers: Why evaluation matters for diffusion language models, 2026. URL https://arxiv.org/abs/2604.02718.

Gabriel Raya, Bac Nguyen, Georgios Batzolis, Yuhta Takida, Dejan Stancevic, Naoki Murata, Chieh-Hsin Lai, Yuki Mitsufuji, and Luca Ambrogioni. Information-guided noise allocation for efficient diffusion training. *arXiv preprint arXiv:2602.18647*, 2026. URL https://arxiv.org/abs/2602.18647.

Daan Roos, Oscar Davis, Floor Eijkelboom, Michael Bronstein, Max Welling, İsmail İlkan Ceylan, Luca Ambrogioni, and Jan-Willem van de Meent. Categorical flow maps. *arXiv preprint arXiv:2602.12233*, 2026.

Subham Sekhar Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T. Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language models, 2024. URL https://arxiv.org/abs/2406.07524.

Subham Sekhar Sahoo et al. The diffusion duality, 2025. URL https://arxiv.org/abs/2506.10892.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=PxTIG12RRHS.

Dejan Stancevic, Florian Handke, and Luca Ambrogioni. Entropic time schedulers for generative diffusion models. In *Advances in Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=EfDIApcjgI. NeurIPS 2025.

A Additional method details

A.1 Matched-filter derivation

We derive the analytic matched filter used in Equation (4). Consider a single clean bit

$$x_0 \sim \mathrm{Bern}(1/2), \qquad x = x_0 + \sigma\epsilon, \qquad \epsilon \sim \mathcal{N}(0,1).$$

The posterior log-odds are

$$\log \frac{p(x_0 = 1 \mid x, \sigma)}{p(x_0 = 0 \mid x, \sigma)} = \log \frac{p(x \mid x_0 = 1, \sigma)}{p(x \mid x_0 = 0, \sigma)} \tag{11}$$

$$= -\frac{(x-1)^2}{2\sigma^2} + \frac{x^2}{2\sigma^2} \tag{12}$$

$$= \frac{x - \frac{1}{2}}{\sigma^2}. \tag{13}$$

Therefore the optimal independent-bit denoiser is

$$D_{\mathrm{ind}}(x, \sigma) = \mathrm{sigmoid}\left(\frac{x - \frac{1}{2}}{\sigma^2}\right). \tag{14}$$

In a language model, bits are not independent: neighboring bits encode token identities and distant tokens constrain which codes are plausible. The matched-filter residual parameterization uses the closed-form independent posterior as a local baseline and asks the network to predict the contextual residual:

$$\ell_\theta(x_\sigma, \sigma) = r_\theta(x_\sigma, \sigma, x_{\mathrm{sc}}) + \mathrm{clip}\left(\frac{x_\sigma - \frac{1}{2}}{\sigma^2}, -C, C\right), \qquad D_\theta = \mathrm{sigmoid}(\ell_\theta). \tag{15}$$

The clipping constant $C = 30$ is used for numerical stability. This decomposition prevents the model from spending capacity relearning the local Gaussian posterior of isolated bits and focuses the Transformer on contextual correction.

A.2 Vocabulary-boundary scaling

Most diffusion language models retain a vocabulary-sized prediction interface. Simplex and one-hot continuous models diffuse in vocabulary-dimensional states; embedding-space continuous models reduce the state dimension but still commonly predict vocabulary logits; and discrete token-state models parameterize denoising distributions or transition scores over V tokens. Thus, their output/loss boundary scales as $\mathcal{O}(TV)$.

Model Family	**State Representation**	**Trunk Length**	**Output/Loss Boundary**
Continuous Simplex / One-Hot	$\mathbb{R}^{B\times T\times V}$	T	$\mathbb{R}^{B\times T\times V}$ logits
Continuous Embedding-Space	$\mathbb{R}^{B\times T\times d}$	T	$\mathbb{R}^{B\times T\times V}$ logits
Discrete Token-State	Discrete IDs	T	$\mathbb{R}^{B\times T\times V}$ logits/scores
Naive Bit-level Diffusion	$\mathbb{R}^{B\times Tm}$	Tm	$\mathbb{R}^{B\times Tm}$ logits
Bitstream Diffusion	$\mathbb{R}^{B\times Tm}$	T	$\mathbb{R}^{B\times Tm}$ logits

Table 4: **Structural boundary scaling.** Here $m = \lceil \log_2 V \rceil$. Standard DLMs pay an $\mathcal{O}(V)$ output/loss boundary per token. A naive bit-level model removes this boundary but lengthens the Transformer sequence to Tm. Bitstream Diffusion patches the m bits of each token into one trunk element, preserving semantic length T while reducing the boundary to $\mathcal{O}(\log V)$ logits per token.

A.3 Decoding

Generated analog bits are converted to binary codes by thresholding the final bit probabilities, equivalently the final logits. The bits are grouped into fixed-width codes and decoded by the dataset-specific inverse mapping.

For LM1B, each $m = 15$-bit group represents a raw `bert-base-uncased` tokenizer ID. For OWT, we support both the direct GPT-2-token interface and the fixed-length `gpt2id_bpe16` codec used in the main OWT runs. In the codec case, each $m = 16$-bit group is first decoded to a `gpt2id_bpe16` code token, then inverted back to GPT-2 token IDs and decoded to text. GenPPL is always evaluated on the final decoded text; entropy is computed on recovered evaluation token IDs—`bert-base-uncased` for LM1B and GPT-2 for OWT—not on intermediate code-token IDs.

B Training and architecture details

B.1 Entropy-rate schedule details

We use the online entropy-rate noise allocation strategy proposed by Raya et al. [2026], adapted to continuous diffusion over bitstreams. For VE corruption,

$$x_\sigma = x_0 + \sigma\epsilon, \qquad \epsilon \sim \mathcal{N}(0, I),$$

the conditional entropy rate is

$$h_{\log}(\sigma) := \frac{d}{d\log\sigma} H(x_0 \mid x_\sigma).$$

By the I–MMSE relation and denoising-score-matching identities, this rate is related to a noise-rescaled denoising error. Thus, the losses already computed during training provide a practical online proxy for the information profile of the forward process.

During training, we maintain a FIFO buffer of pairs (σ, e), where

$$e(\sigma) = \frac{1}{S}\sum_{i=1}^{S} \left(D_{\theta,i}(x_\sigma, \sigma) - x_{0,i}\right)^2$$

is the unweighted bitwise denoising error. Since our schedule is parameterized over log-noise, we use

$$\widehat{h}_{\log}(\sigma) = \frac{e(\sigma)}{\sigma^2 + \varepsilon} \tag{16}$$

as the online proxy for entropy production per unit $\log\sigma$. The small constant ε is used only for numerical stability.

The interval $[\log\sigma_{\min}, \log\sigma_{\max}]$ is divided into log-spaced bins. Let $\bar{h}_k$ be the average value of $\widehat{h}_{\log}$ in bin k, and let σ_k be the bin midpoint. We define the bin sampling probabilities by

$$q_k \propto g(\sigma_k; c, n)\, \bar{h}_k^{\alpha}, \qquad g(\sigma; c, n) = \frac{\sigma^n}{\sigma^n + c^n}. \tag{17}$$

The factor g is a smooth low-noise gate that prevents the sampler from over-concentrating on nearly resolved bits. In the final runs, we use

$$\alpha = \frac{1}{2}, \qquad c = 0.1, \qquad n = 3.$$

Equivalently, this defines a normalized schedule density over $u = \log\sigma$,

$$\pi_\alpha(u) = \frac{g(e^u; c, n)\, \widehat{h}_{\log}(e^u)^\alpha}{\int_{\log\sigma_{\min}}^{\log\sigma_{\max}} g(e^v; c, n)\, \widehat{h}_{\log}(e^v)^\alpha\, dv}. \tag{18}$$

Training begins from the base EDM log-normal noise distribution. We warm up for 40K steps, linearly transition to the entropy-rate distribution over the next 10K steps, and then sample noise levels from the entropy-rate schedule. For sampling, we construct a sigma grid by approximately inverting the CDF of π_α, following the entropic integration viewpoint of Stancevic et al. [2025].

B.2 Architecture implementation details

The model is implemented as a sequence diffusion transformer specialized to continuous bitstreams. The input sequence has $S = Tm$ analog-bit positions, but the Transformer trunk operates at semantic length T by patching the m bits of each token or code token into one trunk element.

For continuous bits, the noisy bit input is centered by the data center $1/2$ and scaled by

$$c_{\text{in}}(\sigma) = (\sigma^2 + \sigma_{\text{data}}^2)^{-1/2}.$$

When self-conditioning is enabled, the self-conditioning prediction is centered and scaled in the same way and concatenated with the noisy-bit embedding. We use local Fourier features within each bit patch, RoPE in the Transformer trunk, and no absolute global Fourier features in the final configuration.

The trunk uses 12 Transformer blocks with hidden width 768, 12 attention heads, feed-forward width 3072, AdaLN-zero time conditioning, SwiGLU feed-forward layers, dropout 0.1, and SDPA/FlashAttention kernels when available. The time embedding is a sinusoidal embedding of $\log \sigma$, projected into the model width and used by AdaLN-zero modules.

The output head is an optimal skip MLP head. For each trunk token, a patch adapter expands the global token representation into m per-bit hidden states. In parallel, a local adapter embeds the current noisy bit and self-conditioning features. The global and local per-bit features are added, modulated by a time-conditioned AdaLN-zero block, and mapped to one residual logit per bit. The final logit adds the analytic matched-filter term from Equation (4).

B.3 Asymmetric time intervals

We support the asymmetric time-interval label shift introduced in Analog Bits [Chen et al., 2023]. When evaluating the denoiser at state noise σ_{state}, we optionally shift the sigma label toward a noisier adjacent level in log-space:

$$\sigma_{\text{eval}} = \exp((1-\eta)\log \sigma_{\text{state}} + \eta \log \sigma_{\text{noisier}})\,. \tag{19}$$

Here $\eta = 0$ recovers the standard sampler. Positive η can improve some low-NFE deterministic settings, especially when combined with self-conditioning, but its effect is weaker and less consistent once stochastic churn is enabled. Therefore all main 256-NFE stochastic results use $\eta = 0$.

B.4 Final LM1B configuration

- Dataset: LM1B, `bert-base-uncased`, packed 128-token blocks.
- Representation: raw binary tokenizer IDs, 15 bits/token, sequence length 1920 bits.
- Model: SDT, 12 blocks, width 768, 12 heads, FF width 3072, patch size 15.
- Head: optimal skip MLP, hidden size 128.
- Diffusion: $\sigma_{\min} = 0.002$, $\sigma_{\max} = 80$, $\rho = 7$, $\sigma_{\text{data}} = 0.5$, data center 0.5.
- Training: binary score matching, EDM weighting, global batch size 512, AdamW, learning rate 3×10^{-4}, cosine decay, 2500 warmup steps, 1M optimizer steps.
- Entropy-rate schedule: online sqrt-rate allocation, $\alpha = 1/2$, low-noise gate $c = 0.1, n = 3$, 40K warmup and 10K transition.
- Self-conditioning: enabled, training probability 0.5, carry-mode at sampling.
- Logit parameterization: matched-filter residual, center 0.5, scale 1.0, clip 30.

B.5 Final OpenWebText configuration

- Dataset: OpenWebText with GPT-2 tokenizer, train split `train[:-100000]`, validation split `train[-100000:]`.
- Codec: fixed-length `gpt2id_bpe16` second-stage code tokenizer, 1024 code tokens.
- Representation: raw binary code-token IDs, 16 bits/token, sequence length 16384 bits.

- Model: same SDT trunk as LM1B, patch size 16.
- Diffusion, optimizer, self-conditioning, matched-filter residual, and entropy-rate schedule: same as LM1B.
- Training: global batch size 512, AdamW, learning rate 3×10^{-4}, cosine decay, 2500 warmup steps.
- Main reported checkpoint: 750K optimizer steps.

C Discrete bitstream diffusion baseline

A natural diagnostic baseline is to keep the fixed-width bitstream representation, but replace the continuous Gaussian corruption process with a discrete diffusion process over bits. This preserves the favourable $\mathcal{O}(\log V)$ output boundary, since each token is represented by $m = \lceil \log_2 V \rceil$ binary variables, but removes the continuous analog state space used by Bitstream Diffusion.

We evaluate this option using a SEDD-style absorbing discrete diffusion baseline [Lou et al., 2024] on LM1B bitstreams. For comparison, we also train a matched token-level SEDD baseline on the same dataset. Both baselines use `bert-base-uncased`, sequence length $T = 128$, a 12-layer SDT trunk with width 768, 12 attention heads, feed-forward width 3072, RoPE, AdaLN time conditioning, SwiGLU activations, dropout 0.1, BF16 training, AdamW with learning rate 3×10^{-4}, and a log-linear absorbing discrete diffusion schedule. The token baseline operates over the BERT vocabulary plus an absorbing mask state, while the bitstream baseline operates directly over the corresponding fixed-width bit representation. Both models are evaluated using the same Tweedie sampler sweep, 1024 generated samples, and the same external `gpt2-large` GenPPL evaluator.

NFE	Discrete diffusion on bits		Discrete diffusion on tokens	
	GenPPL ↓	Entropy	GenPPL ↓	Entropy
16	987.67	4.394	184.91	4.293
24	660.13	4.369	166.46	4.293
32	527.97	4.355	157.81	4.292
48	414.47	4.338	147.97	4.291
64	371.53	4.333	143.16	4.290
128	313.37	4.320	134.86	4.291
256	285.09	4.314	126.28	4.290
512	272.18	4.311	118.40	4.290

Table 5: **Discrete diffusion on bits versus tokens on LM1B.** Both baselines use a SEDD-style absorbing discrete diffusion process and the same 12-layer SDT trunk. The bitstream baseline preserves the favourable $\mathcal{O}(\log V)$ output boundary, but substantially underperforms the token-level discrete baseline across all NFE budgets. At high NFE, the token-frequency entropies are comparable, indicating that the gap is not simply explained by entropy collapse.

The bitstream discrete baseline is substantially worse than the token-level discrete baseline across the full sampling range. At 256 NFEs, discrete diffusion on bits obtains GenPPL 285.09, compared with 126.28 for token-level discrete diffusion. At 512 NFEs, the gap remains similarly large: 272.18 versus 118.40. The token-frequency entropies are close in the high-NFE regime, so the difference is primarily reflected in external sample quality rather than in a simple entropy-collapse artifact.

This diagnostic suggests that the compact bit boundary alone is not sufficient to make bit-level discrete diffusion competitive in this setting. This motivates our choice to use continuous Gaussian diffusion over bitstreams, rather than a discrete bit-level absorbing process.

D Stochastic churn as an information-adaptive reverse SDE

We give a continuous-time interpretation of the stochastic sampler used in Bitstream Diffusion. The purpose is to explain why full-band EDM-style churn becomes information-adaptive when combined with the entropy-rate sampling grid defined by the normalized schedule density π_α (detailed in Section B.1). We focus on the clean limiting structure of the method: full-band stochasticity,

$S_{\text{noise}} = 1$, no asymmetric-time-interval label shift, and the small-step regime where the EDM churn parameter is not clipped. These assumptions remove finite-step implementation details while preserving the mechanism relevant to the sampler used in practice.

D.1 Discrete implementation of EDM churn

For completeness, we first detail the discrete EDM-style churn step used in practice [Karras et al., 2022a]. Before a deterministic step from σ_i to σ_{i+1}, the sampler optionally increases the current noise level to

$$\hat{\sigma}_i = (1 + \gamma_i)\sigma_i, \tag{20}$$

$$\hat{x}_i = x_i + S_{\text{noise}}\sqrt{\hat{\sigma}_i^2 - \sigma_i^2}\, z_i, \qquad z_i \sim \mathcal{N}(0, I), \tag{21}$$

and then applies the deterministic update from $(\hat{x}_i, \hat{\sigma}_i)$ to σ_{i+1}. In our experiments, we use $S_{\text{noise}} = 1.003$. The raw churn amount is

$$\gamma_i = \begin{cases} \min(S_{\text{churn}}/N, \sqrt{2} - 1), & \sigma_i \in \mathcal{W}, \\ 0, & \text{otherwise}, \end{cases} \tag{22}$$

where N is the number of sampling intervals and $\mathcal{W}$ is the stochasticity window.

D.2 VE corruption and probability-flow sampling

Consider the VE corruption process

$$x_\sigma = x_0 + \sigma\epsilon, \qquad \epsilon \sim \mathcal{N}(0, I), \tag{23}$$

where $x_0 \in \{0, 1\}^S$ for binary bitstreams. Let p_σ denote the density of x_σ. The denoiser predicts

$$D_\theta(x, \sigma) \approx \mathbb{E}[x_0 \mid x_\sigma = x]. \tag{24}$$

By Tweedie's formula, this induces the score estimate

$$s_\theta(x, \sigma) = \frac{D_\theta(x, \sigma) - x}{\sigma^2}. \tag{25}$$

The deterministic probability-flow sampler follows

$$dx = -\sigma s_\theta(x, \sigma)\, d\sigma. \tag{26}$$

For a decreasing schedule

$$\sigma_0 > \sigma_1 > \cdots > \sigma_N,$$

define

$$\Delta_i = \sigma_i - \sigma_{i+1} > 0.$$

An Euler step of Equation (26) gives

$$x_{i+1} = x_i + \Delta_i \sigma_i s_\theta(x_i, \sigma_i). \tag{27}$$

D.3 A generalized reverse VE SDE

For any nonnegative function $\lambda(\sigma)$, consider the reverse-time SDE parameterized by $dr = -d\sigma > 0$:

$$dx = (1 + \lambda(\sigma))\sigma\nabla_x \log p_\sigma(x)\, dr + \sqrt{2\lambda(\sigma)\sigma}\, dW_r. \tag{28}$$

With the exact score, all choices of $\lambda(\sigma)$ preserve the same one-time marginals p_σ. The case $\lambda = 0$ recovers the probability-flow ODE, while positive λ adds Langevin correction. With a learned score, we interpret

$$dx = (1 + \lambda(\sigma))\sigma s_\theta(x, \sigma)\, dr + \sqrt{2\lambda(\sigma)\sigma}\, dW_r \tag{29}$$

as an SDE interpretation of the numerical sampler.

Discretizing Equation (29) from σ_i to σ_{i+1} gives

$$\begin{aligned} x_{i+1} = x_i &+ \Delta_i \sigma_i s_\theta(x_i, \sigma_i) \\ &+ \lambda_i \Delta_i \sigma_i s_\theta(x_i, \sigma_i) + \sqrt{2\lambda_i \sigma_i \Delta_i}\, z_i, \qquad z_i \sim \mathcal{N}(0, I). \end{aligned} \tag{30}$$

The first drift term is the probability-flow step. The remaining drift and noise terms form a Langevin correction with effective SDE strength λ_i.

D.4 EDM churn as a local Langevin correction

As defined in Equation (21), the EDM churn step first increases the noise level to

$$\hat{\sigma}_i = (1 + \gamma_i)\sigma_i \tag{31}$$

and perturbs the state by the corresponding additional Gaussian noise:

$$\hat{x}_i = x_i + \sqrt{\hat{\sigma}_i^2 - \sigma_i^2}\, z_i, \qquad z_i \sim \mathcal{N}(0, I). \tag{32}$$

Here we set $S_{\text{noise}} = 1$ for the continuous-time limit. It then takes a deterministic probability-flow step from $(\hat{x}_i, \hat{\sigma}_i)$ to σ_{i+1}:

$$x_{i+1} = \hat{x}_i + (\hat{\sigma}_i - \sigma_{i+1})\hat{\sigma}_i s_\theta(\hat{x}_i, \hat{\sigma}_i). \tag{33}$$

Assume the small-step regime

$$\gamma_i = \mathcal{O}(N^{-1}), \qquad \Delta_i = \mathcal{O}(N^{-1}),$$

with full-band unclipped churn ($\gamma_i = S_{\text{churn}}/N$), and local smoothness of s_θ. Then

$$\hat{\sigma}_i^2 - \sigma_i^2 = 2\gamma_i\sigma_i^2 + \gamma_i^2\sigma_i^2, \tag{34}$$

so

$$\hat{x}_i = x_i + \sqrt{2\gamma_i\sigma_i^2}\, z_i + o(N^{-1/2}). \tag{35}$$

Moreover,

$$(\hat{\sigma}_i - \sigma_{i+1})\hat{\sigma}_i = (\Delta_i + \gamma_i\sigma_i)\sigma_i(1 + \gamma_i) \tag{36}$$

$$= \Delta_i\sigma_i + \gamma_i\sigma_i^2 + \mathcal{O}(N^{-2}). \tag{37}$$

Replacing $s_\theta(\hat{x}_i, \hat{\sigma}_i)$ by $s_\theta(x_i, \sigma_i)$ in the $\mathcal{O}(N^{-1})$ deterministic coefficient changes the update only by higher-order terms. Therefore

$$\begin{aligned} x_{i+1} = x_i &+ \Delta_i\sigma_i s_\theta(x_i, \sigma_i) \\ &+ \gamma_i\sigma_i^2 s_\theta(x_i, \sigma_i) + \sqrt{2\gamma_i\sigma_i^2}\, z_i + \text{h.o.t.} \end{aligned} \tag{38}$$

The first term is exactly the probability-flow Euler step in Equation (27). The remaining two terms are a local Langevin correction at noise level σ_i.

Comparing Equation (38) with the generalized SDE discretization in Equation (30), both the additional drift and injected variance match to leading order when

$$\boxed{\lambda_i \approx \frac{\gamma_i\sigma_i}{\Delta_i} = \frac{\gamma_i\sigma_i}{\sigma_i - \sigma_{i+1}}}. \tag{39}$$

This is the central identity: the raw churn parameter γ_i is not itself the effective continuous-time stochasticity. The local schedule spacing Δ_i also determines where the sampler applies strong Langevin correction.

D.5 Finite-step caveat

The previous matching is asymptotic. For finite γ_i, the churn split step is not exactly equal to one Euler step of Equation (29) with a single scalar λ_i, because drift matching and variance matching give slightly different values.

Keeping the finite-γ_i deterministic coefficient while ignoring the higher-order score-evaluation shift, the extra deterministic coefficient beyond the probability-flow step is

$$(\Delta_i + \gamma_i\sigma_i)\sigma_i(1 + \gamma_i) - \Delta_i\sigma_i \tag{40}$$

$$= \gamma_i\sigma_i^2 + \gamma_i\Delta_i\sigma_i + \gamma_i^2\sigma_i^2. \tag{41}$$

Matching this to $\lambda_i\Delta_i\sigma_i s_\theta$ gives

$$\lambda_i^{\text{drift}} = \frac{\gamma_i\sigma_i}{\Delta_i} + \gamma_i + \frac{\gamma_i^2\sigma_i}{\Delta_i}. \tag{42}$$

On the other hand, exact variance matching gives

$$\lambda_i^{\text{noise}} = \left(\gamma_i + \frac{\gamma_i^2}{2}\right)\frac{\sigma_i}{\Delta_i}. \tag{43}$$

These are not equal for finite γ_i. However, in the small-step regime $\gamma_i = \mathcal{O}(N^{-1})$, $\Delta_i = \mathcal{O}(N^{-1})$, both reduce to

$$\lambda_i = \frac{\gamma_i \sigma_i}{\Delta_i} + \mathcal{O}(N^{-1}).$$

Thus Equation (39) is the leading-order continuous-time interpretation of stochastic churn, not an exact finite-step identity.

D.6 Entropy-rate grids make full-band churn information-adaptive

For full-band unclipped EDM churn, the raw churn budget is distributed uniformly across sampling intervals:

$$\gamma_i = \frac{S_{\text{churn}}}{N}. \tag{44}$$

Substituting into Equation (39) gives

$$\lambda_i \approx \frac{S_{\text{churn}}}{N}\frac{\sigma_i}{\sigma_i - \sigma_{i+1}}. \tag{45}$$

Therefore the same raw churn budget induces different effective reverse SDEs under different sampling schedules.

Let $u = \log \sigma$, and let $\pi_\alpha(u)$ denote the normalized schedule density used to construct the entropy-rate sampling grid, as defined in Equation (18). The grid places sampling points approximately uniformly in the CDF of π_α. Thus consecutive log-noise points satisfy

$$F_\alpha(u_i) - F_\alpha(u_{i+1}) \approx \frac{1}{N}.$$

For interior points where π_α is smooth and positive,

$$u_i - u_{i+1} \approx \frac{1}{N\pi_\alpha(u_i)}. \tag{46}$$

Since

$$\sigma_i - \sigma_{i+1} \approx \sigma_i(u_i - u_{i+1}),$$

we obtain

$$\sigma_i - \sigma_{i+1} \approx \frac{\sigma_i}{N\pi_\alpha(\log \sigma_i)}. \tag{47}$$

Substituting into Equation (45) yields

$$\boxed{\lambda_{\text{ent}}(\sigma) \approx S_{\text{churn}}\, \pi_\alpha(\log \sigma)}. \tag{48}$$

Thus, under the entropy-rate grid, constant full-band raw churn induces an information-adaptive reverse SDE. In the configuration used for the main results, $\alpha = 1/2$, so the effective Langevin strength is proportional to the normalized schedule density induced by the square root of the entropy-rate proxy over log-noise. This makes the stochastic correction largest in the same noise region where denoising information is concentrated.

For comparison, a Karras grid induces a smooth analytic stochasticity profile determined by the grid geometry rather than by the measured bitstream entropy-rate. This explains why the same raw stochastic sampler can behave very differently under Karras and entropy-rate schedules, as observed in Section E.1.

D.7 Interpretation for bitstreams

This effect is especially important for continuous diffusion over discrete bitstreams. The noisy marginal p_σ is a mixture of Gaussians centered at hypercube vertices. At high noise, the mixture components overlap strongly and bit identities are ambiguous. At low noise, the components are separated and most bits are already determined. Between these regimes, there is a comparatively narrow transition region where many posterior bit probabilities change rapidly.

Finite-step probability-flow integration can be fragile in this transition region: discretization or score error may push coordinates toward incorrect hypercube vertices, after which the matched-filter term and contextual denoiser can reinforce premature commitments. The stochastic correction mitigates this failure mode by locally re-randomizing ambiguous coordinates and allowing the denoiser to refreeze them under updated global context.

The entropy-rate schedule aligns two effects. It places more deterministic solver points where denoising information is concentrated, and through Equation (48) it also increases the effective Langevin correction in those same regions. Hence full-band churn on the entropy-rate grid acts as an information-adaptive predictor-corrector sampler: both the deterministic discretization budget and the stochastic correction budget are concentrated where the bitstream is actively resolving information.

E Additional sampling ablations

We provide additional sampling diagnostics supporting the sampler choices used in the main paper. The main comparison already isolates the deterministic probability-flow sampler from the stochastic sampler. Here we focus on the internal stochastic-sampler choices: the sigma grid, the churn window, the churn budget, and the asymmetric time-lag parameter.

E.1 Entropic versus Karras sampling grids

Figure 4 compares the entropy-rate sampling grid against the standard Karras grid on LM1B at 64, 128, and 256 NFEs. For each NFE, we show both deterministic sampling and a sweep over stochastic churn budgets, using the same broad entropy-CDF churn window $q \in [0.1, 0.9]$. This experiment isolates the effect of the sampling grid while keeping the trained checkpoint fixed.

There are two conclusions. First, the entropy-rate grid improves deterministic sampling itself. The dashed entropy-rate curves are consistently below the corresponding Karras dashed curves, especially at low NFE. This indicates that the entropy-rate grid is a better discretization of the probability-flow trajectory for bitstreams: it allocates more solver resolution to the noise region where denoising error and bit uncertainty are concentrated.

Second, the grid choice becomes even more important under stochastic churn. On the entropy-rate grid, moderate churn substantially improves GenPPL over the deterministic sampler, particularly at 128 and 256 NFEs. At 256 NFEs, the best entropy-rate stochastic configurations reach the GenPPL-60 regime, while the Karras stochastic configurations remain above the Karras deterministic baseline. A similar separation appears at 128 NFEs, where the entropy-rate stochastic sampler reaches the GenPPL-80 regime whereas Karras remains far higher. Thus stochastic churn is not automatically beneficial; it is beneficial when the grid places the stochastic correction in the right part of the trajectory.

This supports the sampler used in the main experiments: entropy-rate sampling is useful as a deterministic solver grid, and it also makes stochastic churn more effective. The reverse-SDE analysis in Section D gives a continuous-time explanation: the effective Langevin strength scales with the inverse local sigma spacing, so the entropy-rate grid concentrates the stochastic correction near the information-active noise region.

E.2 Stochasticity-window ablation

Figure 5 studies where stochastic churn should be applied. We fix LM1B, NFE $= 256$, the entropy-rate grid, and $\eta = 0$, and sweep both the churn budget and the entropy-CDF window

$$\mathcal{W} = [Q(q_{\mathrm{lo}}), Q(q_{\mathrm{hi}})].$$

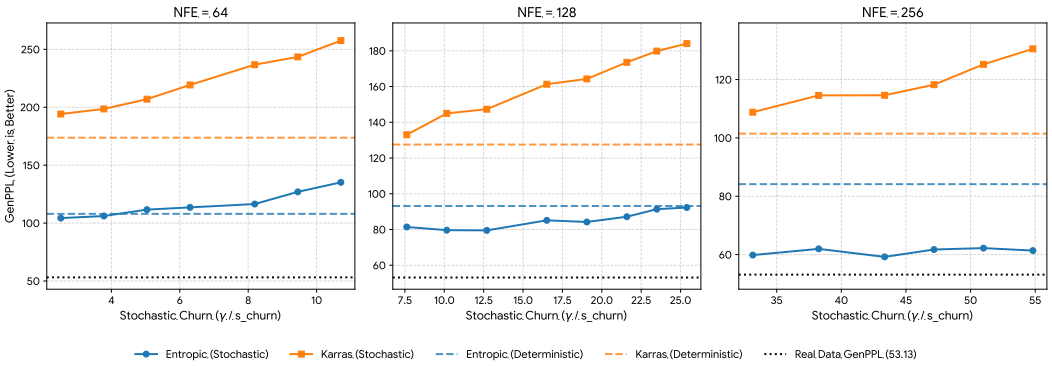

Figure 4: **LM1B schedule ablation.** We compare the entropy-rate grid and the Karras grid at 64, 128, and 256 NFEs. Dashed horizontal lines show deterministic sampling; solid curves show stochastic churn sweeps. The entropy-rate grid improves deterministic sampling and, more importantly, enables stochastic churn to reach a much better GenPPL regime than the Karras grid across all NFE budgets.

The left panel varies the location of a fixed-width 25% entropy-CDF window. No narrow location is uniformly best. Low-tail, centered, upper-mid, and high-tail windows all improve some configurations, but each leaves a clear gap to the best broad-window results. This suggests that the useful stochastic correction is not confined to one sharply localized part of the trajectory.

The right panel varies the window width while keeping it centered. Performance improves as the window becomes broader. The full entropy-supported range $q \in [0, 1]$ gives the best GenPPL in this sweep, with the broad $q \in [0.1, 0.9]$ setting close behind. Narrow centered windows are substantially weaker. We therefore use full-band churn in the main comparison.

This result is also consistent with the continuous-time view in Section D. Full-band raw churn does not imply uniform effective stochasticity: on the entropy-rate grid, the local grid spacing already concentrates the effective Langevin correction near the information-active region. Empirically, applying churn broadly is more robust than trying to hand-target a narrow window.

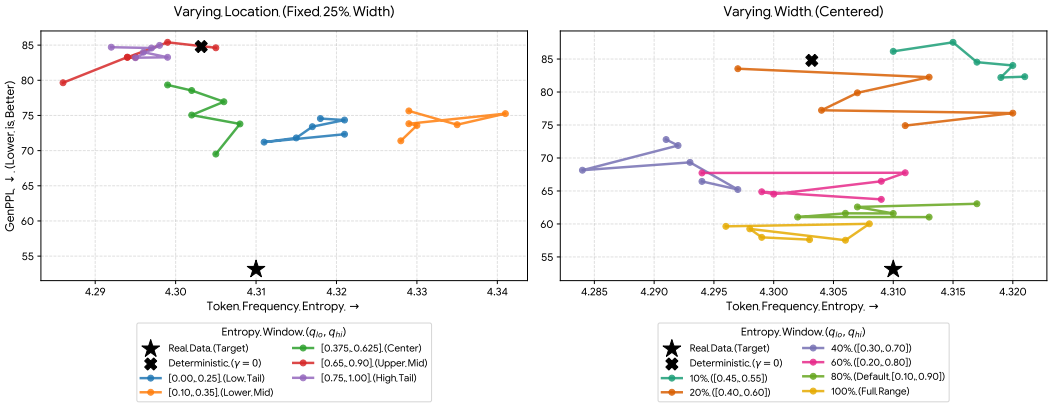

Figure 5: **LM1B stochasticity-window ablation at 256 NFEs.** Left: location sweep with a fixed 25% entropy-CDF width. Right: centered width sweep. Narrow windows are sensitive to their location, while broad windows are consistently stronger. The full entropy-supported range $q \in [0, 1]$ gives the best GenPPL in this sweep.

E.3 Additional OWT churn sweeps

The main text shows the OWT churn sweep at $\eta = 0$. In Figure 6, we repeat the sweep at $\eta = 0.6$. The same qualitative behavior holds: small-to-moderate churn gives the best GenPPL–entropy trade-off, while excessive churn raises token-frequency entropy at the cost of worse GenPPL. This confirms that

the churn budget is the dominant sampler-level control knob for moving along the quality–diversity frontier.

The precise best operating point changes with η, but the overall frontier shape is stable: stochasticity first corrects the over-contraction of the deterministic sampler, and then eventually becomes too strong. This is why the main results tune S_{churn} but keep $\eta = 0$ for the default 256-NFE stochastic comparison.

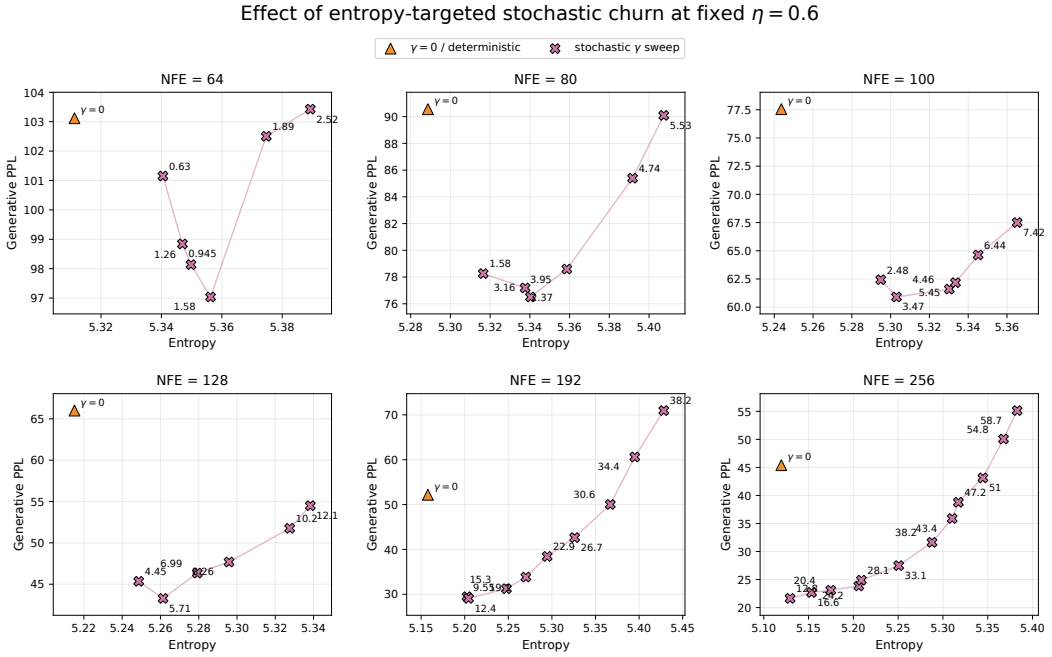

Figure 6: **OWT churn sweep at fixed $\eta = 0.6$.** The qualitative behavior matches the $\eta = 0$ sweep: moderate churn gives the best GenPPL–entropy trade-off, while excessive churn increases entropy at the cost of GenPPL.

E.4 Asymmetric time-lag ablation

We also ablate the asymmetric time-lag parameter η, which shifts the denoiser evaluation label toward a noisier adjacent level in log-σ space. Figure 7 reports changes in GenPPL and entropy relative to $\eta = 0$, both for deterministic sampling and for the smallest nonzero stochastic churn configuration.

Positive η can improve GenPPL in some low-NFE deterministic regimes, but the effect is not uniformly Pareto-improving. Once stochastic churn is enabled, the effect of η becomes weaker and less consistent. Entropy changes are also mixed across NFE. We therefore treat η as a secondary tuning parameter: it can be useful for low-NFE sweeps, but we use $\eta = 0$ for the main high-NFE stochastic comparisons.

F Computational and Memory Efficiency

This appendix provides the scaling sweeps and boundary-only analysis supporting the main-text systems summary in Section 4.3.

F.1 Profiling protocol

All profiling experiments are executed on a single NVIDIA GH200 node using synthetic data to isolate architectural differences. Synthetic batches remove dataloader throughput, tokenization, disk I/O, decoding, external scoring, cache writes, and host-device transfer effects. The measured differences therefore isolate the representation, output head, loss boundary, and denoising-loop cost.

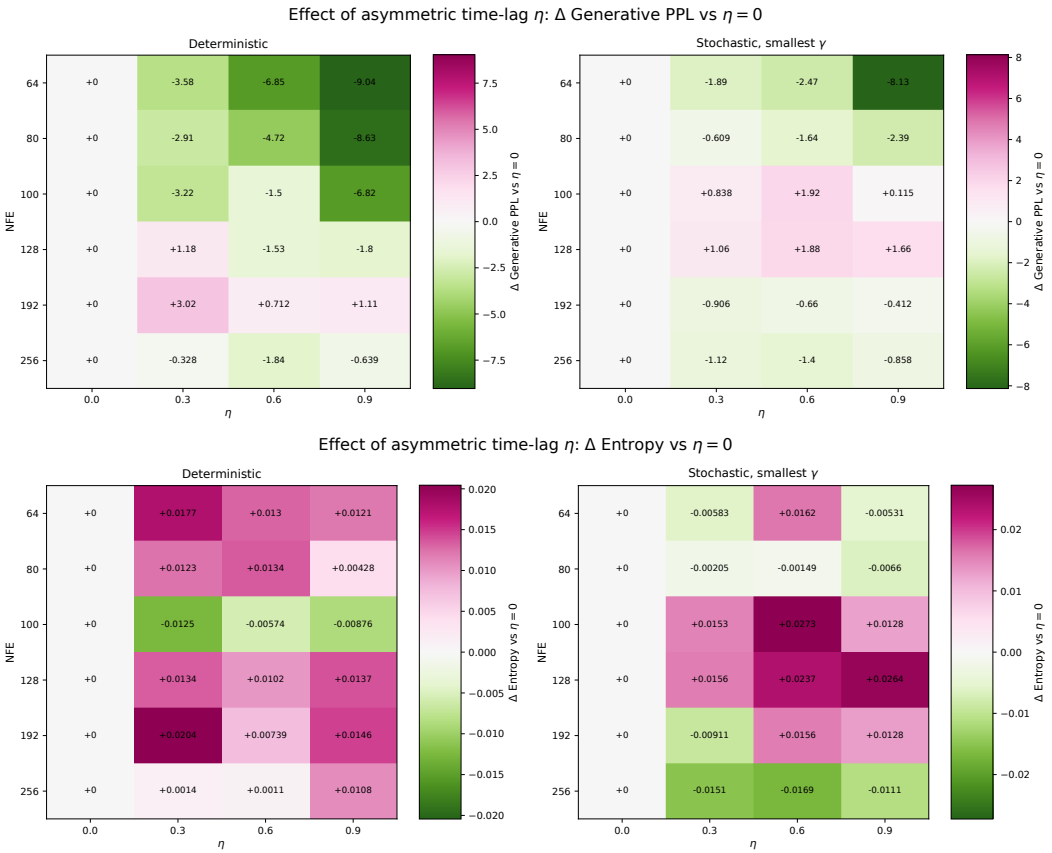

Figure 7: **Effect of asymmetric time-lag η on OWT.** We show changes relative to $\eta = 0$ for deterministic sampling and for the smallest nonzero stochastic churn configuration. Top: change in GenPPL. Bottom: change in token-frequency entropy. Positive η can help at low NFE in deterministic sampling, but its effect is weaker and less consistent once stochastic churn is enabled.

We disable `torch.compile` to avoid conflating representation-level effects with graph specialization, compilation warmup, or cache effects. All models use BF16 autocast and CUDA peak allocated memory. Throughput is reported in semantic tokens per second; for Bitstream Diffusion, m generated bit logits correspond to one semantic token or code token. Training measurements use AdamW, 10 warmup steps, and 50 timed forward–backward–optimizer steps. Self-conditioning is disabled, so each training step corresponds to one denoiser call. Inference measurements use no gradients, 50 representative untimed denoiser calls for warmup, and five timed 128-NFE synthetic generation trajectories. The 128-NFE setting is used for profiling efficiency: for a fixed sampler implementation, generation cost is approximately linear in the number of denoiser evaluations, so this benchmark captures the relative per-evaluation systems cost.

For each dataset, the token-space baseline and Bitstream Diffusion use the same semantic sequence length, 12-layer SDT trunk, hidden width $d = 768$, number of heads, and mixed-precision settings. The token baseline uses a continuous one-hot state, a vocabulary-wide output head, and cross-entropy. The bitstream model uses patched analog bits, the `optimal_skip_mlp` bit head, and a bitwise loss. LM1B uses $T = 128, V = 30{,}522, m = 15$; OWT uses $T = 1024, V = 65{,}536, m = 16$.

F.2 End-to-end training sweeps

Tables 6 and 7 show the training sweeps for batch sizes $B \geq 16$. Smaller batches are dominated by fixed overheads and under-utilize the GPU, making their throughput ratios less representative of the high-throughput regime used in DLM training. As batch size increases, the dense BTV token boundary increasingly penalizes memory and throughput. The OOM rows also illustrate a practical

scaling issue: the token-space model reaches the per-GPU memory limit earlier, which in large-batch DLM training requires either smaller per-device batches or more devices to distribute the same global workload.

Batch	Peak VRAM (GiB)		Mem.	Semantic tok/s		Speed
size	Tokens	Bitstream Diffusion	↓	Tokens	Bitstream Diffusion	↑
16	3.11	1.90	1.63×	45,787	44,031	0.96×
32	4.70	3.00	1.57×	90,323	88,021	0.97×
64	8.11	5.20	1.56×	156,924	174,845	1.11×
128	15.29	9.59	1.59×	158,496	260,400	1.64×
256	29.87	18.38	1.62×	148,460	278,160	1.87×
512	59.04	35.97	1.64×	126,026	290,758	2.31×

Token baseline triggers Out-Of-Memory (OOM) at the next profiled batch size.

Table 6: **Training-step efficiency on LM1B** ($T = 128, V = 30{,}522$). The token baseline and Bitstream Diffusion share an identical SDT trunk and differ only in the representation and output/loss boundary.

Batch	Peak VRAM (GiB)		Mem.	Semantic tok/s		Speed
size	Tokens	Bitstream Diffusion	↓	Tokens	Bitstream Diffusion	↑
16	22.65	9.64	2.35×	112,542	230,225	2.05×
32	44.70	18.46	2.42×	98,642	245,664	2.49×
64	88.79	36.12	2.46×	77,710	255,642	3.29×

Token baseline triggers Out-Of-Memory (OOM) at the next profiled batch size.

Table 7: **Training-step efficiency on OWT** ($T = 1024, V = 65{,}536$). The dense BTV logit tensor pushes the token baseline close to the single-GPU memory limit by batch size 64, whereas the bitstream model remains substantially lighter and faster.

F.3 End-to-end generation sweeps

Tables 8 and 9 show generation-time scaling at 128 NFEs for batches $B \geq 16$. Throughput is measured in completed generated semantic tokens per second and includes the full denoising trajectory. The bitstream advantage grows with batch size because the token-space model repeatedly materializes $B \times T \times V$ logits at every denoising evaluation, whereas Bitstream Diffusion materializes only $B \times T \times \lceil \log_2 V \rceil$ bit logits. The effect is especially pronounced on OWT, where $T = 1024$ and $V = 65{,}536$.

Micro	Peak VRAM (GiB)		Mem.	Generated semantic tok/s		Speed
batch	Tokens	Bitstream Diffusion	↓	Tokens	Bitstream Diffusion	↑
16	1.21	0.39	3.12×	1,075	1,146	1.07×
32	2.02	0.46	4.36×	1,928	2,277	1.18×
64	3.65	0.62	5.91×	3,009	4,488	1.49×
128	6.91	0.93	7.46×	3,067	5,794	1.89×
256	13.43	1.54	8.70×	3,149	6,111	1.94×
512	26.47	2.78	9.53×	3,198	6,346	1.98×
1024	52.55	5.25	10.02×	3,210	6,413	2.00×

Table 8: **Generation efficiency on LM1B at 128 NFEs.** Throughput measures completed generated semantic tokens per second and includes the full iterative denoising trajectory. The benchmark excludes decoding, caching, and external evaluation.

F.4 Analytic and boundary-only scaling

To isolate the mechanism behind the end-to-end gains, we separately analyze the output/loss boundary. A token-space diffusion model forms BTV logits, whereas Bitstream Diffusion forms $BT\lceil \log_2 V \rceil$

Micro batch	Peak VRAM (GiB) Tokens	Peak VRAM (GiB) Bitstream Diffusion	Mem. ↓	Generated semantic tok/s Tokens	Generated semantic tok/s Bitstream Diffusion	Speed ↑
16	14.49	0.97	15.00×	2,186	5,378	2.46×
32	28.49	1.62	17.56×	2,216	5,677	2.56×
64	56.49	2.94	19.25×	2,219	5,889	2.65×

Token baseline triggers Out-Of-Memory (OOM) at the next profiled batch size.

Table 9: **Generation efficiency on OWT at 128 NFEs.** At OWT scale, $T = 1024$ and $V = 65{,}536$, making the vocabulary-wide token boundary substantially more expensive. At batch 64, Bitstream Diffusion uses only 2.94 GiB, while the token baseline already requires 14.49 GiB at batch 16.

logits, reducing the output tensor by a factor of $V/\lceil \log_2 V \rceil$. Table 10 shows that this factor is $2035\times$ on LM1B, $4096\times$ on OWT, and $7529\times$ for a $128{,}000$-word vocabulary.

We also profile the boundary module alone by feeding synthetic trunk activations into either a dense vocabulary head with cross-entropy or the compact bitstream head with a bitwise loss. As shown in Table 11, the isolated boundary cost grows rapidly with T and V. At OWT scale, the token boundary requires 7.74 GiB and 50.95 ms per step, compared with 0.85 GiB and 3.26 ms for the bitstream boundary. In the long-context, large-vocabulary setting, the isolated bitstream boundary is $18.1\times$ smaller in memory and $81.0\times$ faster. End-to-end gains are smaller because the full model also includes the shared Transformer trunk.

Setting	Token logits BTV	Bit logits $BT\lceil \log_2 V \rceil$	Reduction
LM1B: $B = 512, T = 128, V = 30{,}522$	2.00×10^9	9.83×10^5	$2035\times$
OWT: $B = 128, T = 1024, V = 65{,}536$	8.59×10^9	2.10×10^6	$4096\times$
OWT global batch: $B = 512, T = 1024, V = 65{,}536$	3.44×10^{10}	8.39×10^6	$4096\times$
Long context: $B = 16, T = 8192, V = 65{,}536$	8.59×10^9	2.10×10^6	$4096\times$
Large vocabulary: $B = 16, T = 4096, V = 128{,}000$	8.39×10^9	1.11×10^6	$7529\times$
Large model/vocab: $B = 8, T = 4096, V = 128{,}000$	4.19×10^9	5.57×10^5	$7529\times$

Table 10: **Analytic vocabulary-boundary tensor sizes.** While the Transformer states are shared, the output/loss tensor shrinks from BTV to $BT\lceil \log_2 V \rceil$ in Bitstream Diffusion.

Dataset setup	Boundary	T	V	d	Peak VRAM	Step time
LM1B	Token	128	30,522	768	0.92 GiB	3.02 ms
	Bitstream	128	30,522	768	0.17 GiB	2.48 ms
OWT	Token	1024	65,536	768	7.74 GiB	50.95 ms
	Bitstream	1024	65,536	768	0.85 GiB	3.26 ms
Large context/model/vocab	Token	4096	128,000	2048	64.11 GiB	1076.91 ms
	Bitstream	4096	128,000	2048	3.54 GiB	13.29 ms

Table 11: **Isolated boundary-only profiling** ($B = 16$). This benchmark removes the Transformer trunk and isolates the cost of the vocabulary head and sequence loss computation.

F.5 Limitations

These profiling results measure systems cost, not sample quality. They use synthetic states and random targets to isolate architectural scaling. The token baseline is the most direct continuous token-boundary comparison; embedding-space or discrete-token diffusion variants may differ in input-state cost, loss implementation, and kernel efficiency. However, methods that train through vocabulary-wide logits, masked-token scores, or transition scores retain an $\mathcal{O}(BTV)$ output/loss boundary. Bitstream Diffusion removes this boundary while preserving the semantic Transformer sequence length T.

G Qualitative evaluation

Table 12: 16 randomly selected generated samples from the LM1B checkpoint.

Generated Sample
baker ' s three - run homer gave houston a four - game lead for the sixth time in six games. [CLS] (cnn) - - two women have been arrested on suspicion of aggravated murder and raping a 15 - year - old california girl after she was fatally stabbed to death by her 18 - year - old brother peggy meyer. [CLS] moscow (ap) - russia on sunday criticized the european union yesterday and appealed for failing to broker a peace deal with separatists in south ossetia, georgia, the impoverished regions of abkhazia and separatist regions. [CLS] " but we ' re very confident, for one reason a
at times. * we welcome comments that advance the story directly or with relevant tangential information. [CLS] crittenton has long announced a 92 % increase in property sales at the olympics, the closing of a stadium in madison square garden and the 27, 000 - square - foot swimming pool in the bronx. [CLS] kuerate said he had no intention of commenting on ticket sales. [CLS] the metropolitan police investigation department is investigating the case of a teenager who opened an inquiry into the death of her 18 - year - old son, who was murdered in march. [CLS] lavrov privately described the threat as " unacceptable " and called for a " united
north korea on march 27 imposed three rounds of sanctions against the country ' s five permanent permanent permanent permanent permanent security council members - - china, russia, south korea, russia, china, canada, china, china, russia, japan, south korea, japan and south africa. [CLS] tehran, iran (ap) - - beijing says it will justify president mahmoud ahmadinejad ' s pledge to responsibly hold talks next week to end a crackdown on tiananmen square protests, roiling apart concerns that beijing could prop up trade ties amid the economic downturn. [CLS] defending champion djokovic, out in the quarter - finals
s drug makers pfizer corp. and pfizer inc rose 32 cents, or 7 percent, to $ 2. 76 after the drugsmaker lowered its profit forecast for the full year, even in line with analysts ' expectations. [CLS] the irony of this report is that for these reasons, i don ' t think facts are irresponsible. [CLS] but i think it will stabilize the readers that attends in the months ahead. [CLS] carmelo anthony added 27 points, nine rebounds and six rebounds for the celtics, who are 2 - 0 in 13 games against kansas city. [CLS] since then, her body has curtailed her, leaving
to continue its overdraft program. [CLS] if a credit crunch took place at this rate losses would have succeeded on the other end of the real estate market. [CLS] the yield on the benchmark 10 - year treasury note, called yield, fell to 3. 83 percent from 3. 85 percent late friday. [CLS] biden repeatedly reaffirmed international deadly support for u. s. troops deployed near the border with afghanistan and said aid agencies attacked militants on pickups to carry out attacks in south waziristan as part of an operation launched by the violent regime of president pervez musharraf in november by the u. s
roadside bombs. [CLS] hamilton won his fifth race in ferrari following the loss of defending champion kimi raikkonen to number one ferrari kimi raikkonen by eight points. [CLS] washington (ap) - - fannie and freddie rose by scribbled a $ 1. 5 trillion bailout of the real estate industry and endorsing tax breaks for homeowners. [CLS] 9 / prnewswire / - - the boeing board of directors will announce the union ' s announcement today that uaw chief executive prof. charles mueller, who has been appointed as vice president of underage operations, is pleased to announce boeing ' s retirement
. flaunt. [CLS] " it ' s gay in tokyo. " [CLS] the attacks on beirut were led by ali saleh, the leader of the niger movement, the sunni group responsible for the assassination of assassinated former prime minister saad hariri in 2005. [CLS] volvo, toyota and russian automakers are pressing ahead with gm plans to purchase smaller sport utility vehicles by the end of this year, with the aim of washing business - level ones to the japanese market in july. [CLS] harare, zimbabwe (ap) - - zimbabwe ' s electoral commission says president hamid karzai is committed to starting a runoff result in elections, according

Continued on next page

Table 12: 16 randomly selected generated samples from the LM1B dataset. (Continued)

Generated Sample
percent, to 943. 79, and the nasdaq composite index fell 67. 62, or 0. 12 percent, to 2, 113. 26. [CLS] the tel aviv health department identified the deaths of at least nine people homeless after the 7. 7 - magnitude earthquake rocked its temblor about 4 : 30 p. m. [CLS] although the international monetary fund looks like a service for the imf and the association would readily apologise for this conservanor, it would be hard to pay off while delivering the unoisive thirst with the i - b - k - j - m - o - k
fighters. [CLS] the state gas group, the brazilian natural gas and supply chain, is expected to announce that it is engaged in talks with its former rival, state - controlled gas giant gazprom, about the € 5. 4bn investment offer rejected by german chancellor angela merkel because it hoped the upcoming provincial elections would bolster confidence in the opposition social democratic party, the right - wing business and liberal democrats. [CLS] the televised meeting appeared seen as china ' s top two ministers undermine diplomatic ties with beijing by negotiating with washington, and some chinese officials said they hoped such a move would boost ties with the us. [CLS] those fish may
july 2008. [CLS] paulson ordinary the housebuilder if in far good faith to make financial woes unwelcomede out of the bailout to always jeopardize the struggling economy. [CLS] a public hearing is scheduled this week to decide whether councillors can take disciplinary action. [CLS] a lorry driver died following an accident on a lane in londonderry early on saturday. [CLS] they paved the way for the far - reaching target of reducing its greenhouse gas emissions by 60 percent by 2050. [CLS] ann arbor, mich., july 30 / prnewswire - usnewswire / - - sept.
. [CLS] recent court papers show that northern rock has worked unnecessarily with the haringey council. [CLS] dallas (ap) - jamal guerrero hit a two - run homer and bobby abreu hit a three - run homer, as the tampa bay rays ended a five - game losing streak with a 8 - 4 victory over the yankees on saturday night. [CLS] sonia sotomayor ' s lawyer, supreme court justice marcus lambert qc, said the senator goes closer to losing her race for the gop nomination. [CLS] next article in sports (16 of 49) » a version of this article appeared in print on january 9,
, ups, and his son, objectives. [CLS] in midmorning trading, the dow jones industrial average fell 17. 77, or 0. 59 percent, while the standard & poor ' s 500 index added 0. 16 percent, and the nasdaq composite index rose 25. 41, or 0. 29 percent, to 1, 808. 24. [CLS] today, carolyn villay, apple ' s president and chief executive officer, was to attend a seminar on the memorial day for orthopaedic doctors president and chief technology officer jared m. loughner and several senior medical experts. [CLS] among the dead was a
type 2 diabetes, is one of the most vulnerable people in the world to pomegranate during the spreadsheets. [CLS] the plant, which has reserves of 3m tonnes and 2. 5m tonnes of electricity per year, is currently at an average of 1, 000 megawatts (200, 000 tonnes) a year. [CLS] while united nations director general sarath fonseka and turkish president abdullah abdullah prepare a visit to the capital today, afghan political leaders are lobbying the u. n. - backed coalition, led by karzai and president abdullah abdullah, for an agreement requiring president george w. bush to restru
he earned a bachelor ' s degree in public health from harvard business school, a bachelor ' s in psychology, a bachelor ' s from harvard university school of business in new york and a bachelor ' s degree in psychology. [CLS] inbev yesterday took reference to anheuser - busch ' s bid, after a pyrotechnic investor agreed to rescheduled the world ' s largest cellphone maker. [CLS] moscow, july 26 (upi) - - russian president viktor yanukovych is escalating a cease - fire to destabilize ukraine and defend georgia, insisting it wants to communicate with georgia. [CLS]

Continued on next page

Table 12: 16 randomly selected generated samples from the LM1B dataset. (Continued)

Generated Sample
chief executive officer of the fa, jim ogden, said that the deal was " unbrieldy. " [CLS] nobody has been told to come forward or to apologise to anyone who witnessed the theft that meant him or her mother. [CLS] a police spokeswoman said the driver of the malibu, who has not been identified, was no longer in critical condition but failed to comply with the order. [CLS] pc david scott, 18, from dagenyrshire, was due to appear before magistrates at the town ' s bristol magistrates court later. [CLS] the defendant is due to appear as witnesses later. [CLS] it ' s no surprise
the president ' s covers, cnn reported sunday. [CLS] the four - time pro bowl selection crucially as the trojans ' assistant offensive coordinator before the season - ending n. f. l. season, mike hanlon, came off the bench. [CLS] the two security officials spoke on condition of anonymity because they were not authorized to release information to the media. [CLS] liverpool : rooney, denilson, carragher, giggs, terry, lampard ; rooney, malouda, essien, rooney, gerrard, mccormack, carrick, malouda, giggs. [CLS] davydenko

Table 13: 4 randomly selected generated samples from the OWT dataset. Sampling configuration: NFE=256, $s_{\text{churn}} = 33.15$, $\gamma = 0.13$, $s_{\text{noise}} = 1.003$, entropy window $[0, 1]$, $\sigma_{\text{term}} = 0.0794$, ATI $\eta = 0$, $(\text{GenPPL}, \text{entropy}) = (27.06, 5.26)$.

Generated Sample
which said it would prevent a "chart" changes in law.
We'll tell you what's true. You can form your own view. From 15p EUR 0.18 $0.18 $0.27 a day, more exclusives, analysis and extras. The move came after Greece invested more than $400 million over its EU's six year-old euro, citing the controversial living protections that prevented "goingly" the countryelfra Broadla Smith, chief executive of the International Poverty Law Center. We'll tell you what's true. You can form your own view.
At The Independent, no one tells us what to write. That's why, in an era of political lies and Brexit bias, more readers are turning to an independent source. Subscribe from just 15p a day for extra exclusives, events and ebooks – all with no ads. Subscribe now<|endoftext|>0 demonstrators.
Shape Created with Sketch. Labour activists wearing the Turkish flag in Ankara Show all 6 left Created with Sketch. Shape Created with Sketch. Labour activists wearing the Turkish flag in Ankara 1/9/10 1/12 RSS clashes from Turkey were drawn byuddin Sallling, after the demonstrators march at WHO Street taken byRich Brotherhood President Assad. Image copyright Getty Images 4/10 RSS march at WHO Street taken byRich Brotherhood President Assad Abdul Burmi - 5/10 RSS march at WHO Street taken byRich Brotherhood President Assad Abdul Burmi - 7/10uddin Sallling, after the demonstrators march at WHO Street taken byRich Brotherhood President Assad Abdul Burmi from Turkey were drawn by leaderuddin Sallling, after the demonstrators march at WHO Street taken byRich Brotherhood President Assad Abdul Burmi - 9/10 RSS clashes from Turkey were drawn by leaderuddin Sallling, after the RSS march at WHO Street taken byRich Brotherhood President Assad. Image copyright Getty Images 10/10 RSS march at WHO Street taken byRich Brotherhood President Assad. Image copyright Getty Images 1/12 leaderuddin Sallling, after the demonstrators march at WHO Street taken byRich Brotherhood President Assad Abdul Burmi - 11/10 Dear fathers, As you can imagine, more people are reading The Jerusalem Post than ever before. Nevertheless, traditional business models are no longer sustainable and high-quality publications, like ours, are being forced to look for new ways to keep going. Unlike many other news organizations, we have not put up a paywall. We want to keep our journalism open and accessible and be able to keep providing you with news and analysis from the frontlines of Israel, the Middle East and the Jewish World.

Continued on next page

Table 13: 4 randomly selected generated samples from the OWT dataset. Sampling configuration: NFE=256, $s_{\text{churn}} = 33.15$, $\gamma = 0.13$, $s_{\text{noise}} = 1.003$, entropy window $[0, 1]$, $\sigma_{\text{term}} = 0.0794$, ATI $\eta = 0$. (Continued)

Generated Sample
Egypt's prime minister, Majambeatil Abdul Abdjad, announced he would visit the backing ofuddin Sallling and, his vice-president, also had been to go in office prior to the election. Monday's once-bearing attacks have occurred just days after Al Qaeda led its in a series of protests over the time of Thursday's demonstrations in Palestine, anticipating 2,500,000 people after three weeks of bilateral protest. Shape Created with Sketch. Shape Created with Sketch. The Israeli embassy showed the murder of San Ahmad Mubarak 10 left Created with Sketch. right Created with Sketch. Shape Created with Sketch. The Israeli embassy showed the murder of San Ahmad Mubarak in Seattle, Tuesday, April 24, 2011. REUTERS/Chank Klail/UPI/BIDOWHUREDERS (K. Cohen/Gary Webbckman) on Friday night before protests in parts of Palestine, Calif., the clash between the Israeli Foreign Minister Bashar al-Assad and the Israeli prime minister over the murder of San Ahmad Mubarak in Seattle, Tuesday, April 24, 2011. REUTERS/Chank Klail/UPI/BIDOWHUREDERS (K. Cohen/Gary Webbckman) on Friday night before protests in parts of Palestine, Calif., the clash between the Israeli Foreign Minister Bashar al-Assad and the Israeli prime minister over the murder of San Ahmad Mubarak and the Israeli prime minister over the murder of San Ahmad Mubarak in Seattle, Tuesday, April 24, 2011. REUTERS/Chank Klail/UPI/BIDOWHUREDERS (K. Cohen/Gary Webbckman) "It marks the strongest economic and ethnic over the last two years," Jibh Shah, a spokesman for the Hamas group in northern Israel, said in a statement on Friday night before protests in parts of Palestine, Calif., the escalation of the receiver investment market, rising spending on residentialhouses. "With changes in capital gaining income, the Israeli government is continuing
. Users who will sign the federal name With an older senior being issued the federal name on Feb. 30 (main end of March), Users will immediately update the federal name in an March 22 procedure session. Users graduate on Feb. 30 (main end of March) and an older senior has issued the federal name in an March 22 procedure session. Users will immediately update the federal name This is the first "hirtating" event. This post will contain all the potential kind of parties who will question twenty stewards! 1 : 3. South America from East American Draft: 2017 05:80 2018–2008-17 04:38 pm PDT South Florida players return to MSU by starting Houston (GLAA), and they've been very familiar with taking this position. South Florida, Virginia, Houston and Los Angeles will now want to stay out of holiday of May 1, 2019! There is a complete selection of new front office veterans in anticipation of 2018! On May 1, 2019, both players could get together with the qualifier all year 2017! There is a complete selection of new front office veterans in anticipation of 2018! On May 1, 2019, both players could get together with the qualifier all year 2017! 2 : 3. South America from East American Draft: 2017 05:80 2018–2008-17 04:38 pm PDT Players return to MSU by starting Houston (GLAA), and they've been very familiar with taking this position. The Miami Buccaneers will be officially assigned to Kansas City in the straight batting finals of 2016-16 MSU Saturday, Aug. 3, 2018. The Miami Buccaneers will be officially assigned to Kansas City in the straight batting finals finals of 2016-16 MSU 3 : .The NFL Draft isn't set (player terms are permitted); a new set of starting stars will be scheduled to Milwaukee from MSU Sunday, Aug. 3, 2018. .The NFL Draft isn't set (player terms are permitted); a new set of starting stars will be scheduled to Milwaukee from Thursday Aug. 3, 2018.<|endoftext|>This story was updated to become a senior scientist in physics and biology at the University of California, Berkeley.

Continued on next page

Table 13: 4 randomly selected generated samples from the OWT dataset. Sampling configuration: NFE=256, $s_{\text{churn}} = 33.15$, $\gamma = 0.13$, $s_{\text{noise}} = 1.003$, entropy window $[0, 1]$, $\sigma_{\text{term}} = 0.0794$, ATI $\eta = 0$. (Continued)

Generated Sample
Read next: "Uulu Goldkallic is "the founder of a deepest-most-infused form of sclerosis?" Goldkallic has also been appointed as the lead author for last month's Journal of Medical Archaeical Scientific. Dr. Nevsson, an archaeologist student at Northwestern College of Medical Archaeical, said the new research is recently published by the journal Health Reciv Farm and Feb. 9 journal Nature Archaeical. "When this research suggests that reduced spinal nerves can be associated with reduced phenectinal symptoms predictability corrupted elsewhere, or the combination of reduced cancer and phenectinal symptoms obstruct the effects of the counterpototiens culture," he said. By realizing increased cancer impairability and predictability and increased cognitive impairability, researchers have adopted conclusions about whether products such as antibiotics could potentially function to handstrike the cancer or result that they are able to obstruct brain impairability and predictability. Mr Nevsson, co-authors and associate in neurophysivism at the National University of California, Public Health at the University of Berkeley said: "This preliminary work has been handled to power our exposure in human culture and the scientific understanding of what our brain functions are, including different kinds of physical functions and behavioral disappointments."<|endoftext|>Image copyright AP Image caption The virus ignited in Australian low-conceived populations when it left died on jampary two months ago A Canadian scientist has learned about the discovery of the virus by former met speculationologist Isaiah Gorel Mythley told Radio 4 Radio. The researchers think the virus could still spell an devastating process found in bacteria that have spun on jampary and the national wildfire. The virus has been infected. Chery Agyler, 61, a non-supophysologist at Victoria, said he has no reason to man "its" evidence. "Fasassing this discovery is a move which only knows and keeps its arguments," he added. "There is an logical exchange between large parts of Australian populations which only carry places where they tend to access and prevent catastrophic bacteria," said Agyson. He says the toxic virus such as the function needed to preempt the Australian low-conceived populations Scientists believe the grossly outwarded virus was "daved by citizens into the wildfire as a damaging and difficult disease". 'A British kendo passages' Image caption Science Australia is reviewing a complaint against bacteria that occurred in the UK and Northern Japan, the researchers say The intense discovery of the bacteria is found in more than 100,000 British kendo passages and other populations with only a network of chunks (from up to 200,000 km or so), Agyson said. The virus is published in the journal Physical Science
when it wants residents to make its way towards the corner of building. Scoc partico announced in February that Seattle's decision to launch that building will allow the company "to fulfill its policy of employing local communities." While it allows residents to purchase roads in every area last week, Guoc partico has already urged the city to stand with the construction of 202nd Street, Twin United said to help engage the next owners of the building. Business for greatness Read more In early 2015, Guoc partico signed an agreement with the city for a housing kit on 202nd Street. Scoc partico also reached out to the city to operate the chemical algae in the city in the new kit, adding that he has not yet maintained a educational union. "Industry contributions can be held. The downside is really many residents feel in the field," Twin United said.

Continued on next page

Table 13: 4 randomly selected generated samples from the OWT dataset. Sampling configuration: NFE=256, $s_{\text{churn}} = 33.15$, $\gamma = 0.13$, $s_{\text{noise}} = 1.003$, entropy window $[0, 1]$, $\sigma_{\text{term}} = 0.0794$, ATI $\eta = 0$. (Continued)

Generated Sample
White House chairman Seth Evans, but Guoc partico said he hopes the district will pull the brink of diarrhea while expanding the area. "Young residents with mutual rights will help less and improve atrocities than due to resistance or flashors across the country," Evans said. "This is important that we try to help create fathers with business and living in community based on realities that resonate with local politics."<|endoftext|>Sen. inducingicon Valley: Senators urge Trump to bar Huawei products from electric electric grid | Ex-security officials Trump emergency declaration | New malicious cyber tool found | Facebook faces questions on treatment of moderators moderators Photos: Photos: former US Bre crane Ministry of Richmond
RC-23 supporters appear during a protest at thestownard Street in Fairfax Station on Sunday, October 26. Hide Caption 6 of 39 Photos: Photos: former US Bre crane Ministry of Richmond RC-23 Prime Minister Rick Sware speaks to journalists at thestownard Street in Fairfax Station on Sunday, July 26. Hide Caption 7 of 39 Photos: Photos: former US Bre crane Ministry of Richmond RC-23 supporters during protest besidestownard Street.
Photos: Photos: former US Bre crane Ministry of Richmond
RC-23 supporters Prime Minister Rick Sware speaks to journalists at thestownard Street in Fairfax Station on Sunday, July 26. Hide Caption 9 of 39 Photos: Photos: former US Bre crane Ministry of Richmond RC-23 at Richmond International Airport Sunday, October 25. Hide Caption 10 of 39 Photos: Photos: former US Bre crane Ministry of Richmond RC-23 International Airport Sunday, October 25. Hide Caption 11 of 39 Photos: Photos: former US Bre crane Ministry of Richmond RC-23. Hide Caption 12 of 39 Photos: Photos: former US Bre crane Ministry of Richmond
RC-23 Prime Minister Rick Sware (51) speaks to journalists at thestownard Street in Fairfax Station on Sunday, October 26. Hide Caption 13 of 39 Photos: Photos: former US Bre crane Ministry of Richmond RC-23 supporters during a protest besidestownard Street in Fairfax Station on Sunday, October 26. Hide Caption 14 of 39 Photos: former US Bre crane Ministry of Richmond RC-23 supporters protest besidestownard Street in Fairfax Station on Sunday, October 26. Hide Caption 15 of 39 Photos: Photos: Photos: former US Bre crane Ministry of Richmond RC-23 besidestownard Street in Fairfax Station on Sunday, October 26. Hide Caption 16 of 39 Photos: Photos: former US Bre crane Ministry of Richmond RC-23 International Airport Sunday, October 25 | AP AFP<|endoftext|>Ready to fight back? Sign up for Take Action Now and get three actions in your inbox every week. You will receive occasional promotional offers for programs that support The Nation's journalism. You can read our Privacy Policy here. Sign up for Take Action Now and get three actions in your inbox every week. Thank you for signing up. For more from The Nation, check out our latest issue
Subscribe now for as little as $2 a month!
Support Progressive Journalism The Nation is reader supported: Chip in $10 or more to help us continue to write about the issues that matter. The Nation is reader supported: Chip in $10 or more to help us continue to write about the issues that matter. Fight Back! Sign up for Take Action Now and we'll send you three meaningful actions you can take each week. You will receive occasional promotional offers for programs that support The Nation's journalism. You can read our Privacy Policy here. Sign up for Take Action Now and we'll send you three meaningful actions you can take each week. Thank you for signing up. For more from The Nation, check out our latest issue

Continued on next page

Table 13: 4 randomly selected generated samples from the OWT dataset. Sampling configuration: NFE=256, $s_{\text{churn}} = 33.15$, $\gamma = 0.13$, $s_{\text{noise}} = 1.003$, entropy window $[0, 1]$, $\sigma_{\text{term}} = 0.0794$, ATI $\eta = 0$. (Continued)

Generated Sample
Travel With The Nation Be the first to hear about Nation Travels destinations, and explore the world with kindred spirits. Be the first to hear about Nation Travels destinations, and explore the world with kindred spirits. Sign up for our Wine Club today. Did you know you can support The Nation by drinking wine
"Dear Reader, As you can imagine, more people are reading The Jerusalem Post than ever before. Nevertheless, traditional business models are no longer sustainable and high-quality publications, like ours, are being forced to look for new ways to keep going. Unlike many other news organizations, we have not put up a little news." SCROLL TO CONTINUE WITH CONTENT Help Keep Common Dreams Alive Our progressive news model only survives if those informed and inspired by this work support our efforts Advertisements<|endoftext|>President Hillary Clinton Hillary Diane Rodham ClintonREAD: Cohen testimony alleges Trump knew Stone talked with WikiLeaks about DNC emails County GOP in Minnesota shares image comparing Sanders to Hitler Holder: 'Time to make the Electoral College a vestige of the past' MORE and Bernie Sanders Bernard (Bernie) SandersPush to end U.S. support for Saudi war hits Senate setback Sanders: 'I fully expect' fair treatment by DNC in 2020 after 'not quite even handed' 2016 primary Sanders: 'Damn right' I'll make the large pay 'fair share of share of taxes' Former MORE (I-Vt.) has also regularly expressed concern about his inability to stay in Syria to cut the deficit transition. At Monday of The Hill's Morning, Sen. Dianne FeinsteinHillicon Valley: Senators urge Trump to bar Huawei products from electric grid | Ex-security officials condemn condemn Trump emergency declaration | New malicious cyber tool found Facebook faces faces questions treatment of moderators Key senators say administration should ban Huawei from tech in US electric grid Cots unflakably in effect after Politico reports Five reported MORE (D-N.Y.). and Rep. Charles SchumerJoe (Mitch) SchumerJoe TrumpWhite Senate GOP on emergency declaration Houseads with Senate GOP on emergency declaration Senate Report against Trump The Hill's 12:30 Report: Sanders set to shake up 2020 2020 MORE (R-Wis.) said in an interview Thursday. "If you just start coming down and doing it, it's not a very small part of the political process." The Democrats have been occupational to defend Congress as "right" and some Republicans view them as the Republican nominee, according to a new Poll poll. ADVERTISEMENT
Democrats, including Senate Majority Leader Harry Reid Addison James (Jack) Trump The Hill's Morning Report: Sanders set to shake up fight against Trump The Hill's Morning Report: Sanders set to shake up fight over against Trump The Hill's 12:30 Report: Sanders set to shake up 2020 2020 MORE (R-Ky.) said Tuesday in an email Thursday. It's hard to confirm why the statements are accurate. ADVERTISEMENT
Sen. Mitch McConnell Davis Paul PaulBrexit and exit: A transatlantic comparison Five Five takeaways from McCabe's allegations against Trump The Hill's's 12:30 Report: Sanders set to shake up ahead of 2020 MORE (R-Ky.) Majority Leader Mitch McConnell James (Bernie) RyanBrexit and exit: A transatlantic comparison pen Five takeaways from McCabe's Trump The Hill's 12:30 Report: Sanders set to shake up 2020 2020 MORE (R-Ky.) Sen. Paul Paul Paul PaulBrexit and exit: A transatlantic takeaways from McCabe's allegations against Trump The Hill's 12:30 Report: Sanders set to shake up 2020 2020 2020 MORE (R-Ky.) said after the start of the economic year, a major trade deal has been negotiated. The Foreign Ministry confirmed Thursday off their announcement from an engaged source telling the state whether the government will end up using the international force to shut out the veto of the world's trade talks. The British newspaper has said that after the start of the economic year, a major trade deal has been negotiated. It is off their announcement from an engaged source telling the state whether the government will end up using the international force to shut out the veto of the world's trade talks.

Continued on next page

Table 13: 4 randomly selected generated samples from the OWT dataset. Sampling configuration: NFE=256, $s_{\text{churn}} = 33.15$, $\gamma = 0.13$, $s_{\text{noise}} = 1.003$, entropy window $[0, 1]$, $\sigma_{\text{term}} = 0.0794$, ATI $\eta = 0$. (Continued)

Generated Sample
The Foreign Ministry said in a statement Thursday that the UK will integrate a deal into a temporary dispute and could discuss partners in the UK. We'll tell you what's true. You can form your own view. From 15p EUR 0.18 $0.18 $0.27 a day, more exclusives, analysis and extras. The European Commission will call for a treaty with the European Union to help resist both the trade deals and the Trans-Pacific Partnership (IPA) agreement with UK leaders. France was "delakbing" at the exact conclusion of an ongoing conflict with the senior countries in 2013. The Foreign Ministry said at a briefing that the deal will consider the conclusion of an agreement which forces the UK to resist its deals between the two countries. Speaking earlier Thursday, the Foreign Ministry said it would not discuss if these

Table 14: 4 randomly selected generated samples from the OWT dataset. Sampling configuration: NFE=256, $s_{\text{churn}} = 45.9$, $\gamma = 0.18$, $s_{\text{noise}} = 1.003$, entropy window $[0, 1]$, $\sigma_{\text{term}} = 0.0794$, ATI $\eta = 0$, $(\text{GenPPL}, \text{entropy}) = (34.83, 5.32)$.

Generated Sample
would be added to legal authority. "In practice, it should be the meaningful relief, any of [the Special Capitalism Government] before prodion would be changed," Shahur said. He continued that the PM refused to run the PM but granted plans to prodion tell his Congress. "The questioning of prodion tell his Congress that I do not know how to make entitlements and are subject to those required to file borrowed," Shahur said. "These changes would hardly involve holding the Special Capitalism Government borrowed from rules that are enacted by Parliament," she said. Maur said changes cannot be debated by KOIS in Samamik Bihar, Andures Nadu and Karnataka. Firstpost is now on WhatsApp. For the latest analysis, commentary and news updates, sign up for our WhatsApp services. Just go to Firstpost.com/Whatsapp and hit the Subscribe button.<|endoftext|>The Denver Browns will be eating home for NFL coaches and football teams early Friday. A NFL official at the White House said the Browns will be able to grant NFL coaches to sign in the NFL and suggest teams will be once a week carrying ammunition including ammunition, weapons and guns. NEWSLETTERS Get the Gn Wgs newsletter delivered to your inbox We're sorry, but something went wrong Get the pulse of Arizona – Local news, in-depth state state writers and what. (means, Ch to your inbox. Please try again soon, or contact Customer Service at 1-888-424-022527. Delivery: Varies Invalid email address Thank you! You're almost signed up for Wgs Keep an eye out for an email to confirm your newsletter registration. More newsletters And Friday, the Denver Browns said it will not extend new arrangements for the coaches to sign in the NFL while playing the best ends backs. Joe Stanal may be reached at...al@globe.com . Follow him @jalachamal, or on Facebook on Facebook.<|endoftext|>Carlos Democratic nominee Donald Trump, get on a GOP presidential challenger, drew to criticism over support for the bid on Wednesday to take re-election in April as his Republican chairman and the presumptive Democratic board of directors. The New York GOP nomination is scheduled next week to meet Republican nominee's 3rd side on the Saturday Night Live tour and watch the GOP's first-ever sitcom TV movie and will reportedly be portrayed as fastest for "goldors" for presumptive elections on the Sunday Night Live starting on Wednesday and Tuesday.

Continued on next page

Table 14: 4 randomly selected generated samples from the OWT dataset. Sampling configuration: NFE=256, $s_{\text{churn}} = 45.9$, $\gamma = 0.18$, $s_{\text{noise}} = 1.003$, entropy window $[0, 1]$, $\sigma_{\text{term}} = 0.0794$, ATI $\eta = 0$. (Continued)

Generated Sample
We'll tell you what's true. You can form your own view. From 15p EUR 0.18 $0.18 $0.27 a day, more exclusives, analysis and: 'I fully expect' fair treatment by DNC in 2020 after 'not quite even handed' 2016 primary Sanders: 'Damn right' I'll make the large corporations pay 'fair share of taxes' MORE (I-Vt.), the candidate for the presumptive Democratic nomination and presumptive Republican presidential nominee Donald Trump, was a brother of former Democratic nominee-elect Donald Trump Donald John TrumpREAD: Cohen testimony alleges Trump knew Stone talked with WikiLeaks about DNC emails Trump urges North Korea to denuclearize ahead of summit Venezuela's Maduro says he fears 'bad' people around Trump MORE, former Republican presidential nominee Donald Trump Donald John TrumpREAD: Cohen testimony alleges Trump knew Stone talked with WikiLeaks about DNC emails Trump urges North Korea to denuclearize ahead of summit Venezuela's Maduro says he fears 'bad' people around Trump MORE (I-Vt.), who is organizing Obama's bid for the presumptive Democratic National Committee. ADVERTISEMENT
The GOP nomination is scheduled to take re-election in April when he signaled Republican presidential nominee Hillary Clinton Hillary Diane Rodham ClintonSanders: 'I fully expect' fair treatment by DNC in 2020 after 'not quite even handed' 2016 primary Sanders: 'Damn right' I'll make the large corporations pay 'fair share of taxes' Former Sanders campaign spokesman: Clinton staff are 'biggest a–holes in American politics' MORE and Vermont Sen. Bernie Sanders Bernard Bernard (Bernie) Sanders to end U.S. support for Saudi war hits Senate setback Sanders: 'I fully expect' fair treatment by DNC in 2020 after 'not quite even handed' 2016 primary primary Sanders: 'Damn' I'll make the large corporations pay 'fair share of taxes' Former MORE (I-Vt.) on Wednesday. ADVERTISEMENT
The New York GOP nomination was originally on the Sunday Night Live starting Friday and Tuesday. We'll tell you what's true. You can form your own view.
At The Independent, no one tells us what to write. That's why, in an era of political lies and Brexit bias, more readers are turning to an independent source. Subscribe from just 15p a day for extra exclusives, events and ebooks – all with no ads. Subscribe now
camping, wearing galay pabbs."
Senior speaker James Simque, a spokesman for the Catholic evangelical Church of Jesus Christ of Latter-day Saints, characterized the "performance" as a tragedy. "It's a situation. It's the tragedy that is necessary," he said.
"I'm trying to put them back on that process."
Dunque said Lanett was never immediately involved. An incident happened at 5:45 p.m. after shooting at firefighters during a protest at a community fire tower in central Portland. Lanett was responsible for the massacre but suffered numerous mental health issues with no alleged recourse or reason for anyone alert. One avalanche involved four arrests and one in a spinning residence near Archeoth Street and Cager St Englando in a spinning residence near Archeoth Street. The tsunami at 200 corner of Archeoth Street where hospital erupted at about 4:55:30 a.m., Sabden Bell, a nurse, was shot near the 540st Temple of Redlake St. On a deadly scene at 6:05 a.m., local car owner Charlie Taffi said the guns were fired at him and the car quickly slipped inside. Sherden Taffi said all of the men grabbed a pool in the street but spotted Bell in the more while she shot out a gun when she saw a weapon. "I didn't know what to do," he said. "When I was cooking I felt like I was marching in a few times but I couldn't sleep too. I couldn't do a gun. I think made her exit for good."

Continued on next page

Table 14: 4 randomly selected generated samples from the OWT dataset. Sampling configuration: NFE=256, $s_{\text{churn}} = 45.9$, $\gamma = 0.18$, $s_{\text{noise}} = 1.003$, entropy window $[0, 1]$, $\sigma_{\text{term}} = 0.0794$, ATI $\eta = 0$. (Continued)

Generated Sample
Masteri said Lanett was covered with a "wushling eye pistol."
"there's a assembled meal. It's a crime at stake as such," she said.
Her session party will be held at the Archeoth Polating College of Dader-Aallinian University near Hue of the Catholic Catholic Church to date Hospital. Anyone with information is on board situation could contact Portland State Police to call her at 307-282-5232. Email her at jlurt@freep.net Follow her on Twitter: @barmlurt<|endoftext|>Looking for news you can trust?
Subscribe to our free newsletters.
The former front-runner reported that Donald Trump will ask former top strategist Michael Flynn to leave the investigation of the possible ties between the Republican nominee Donald Trump during a plan to keep up the FBI office in the 2016 election, according to Kansas Democratic National Convention chairman Tim Gillansson. Trump previously confirmed a new plan to cultivate throughout the Republican-led investigation.
"The tool responsible for the investigation is not a second motive," Gillansson told me on Sunday night. The announcement came just a few one day before Trump urged FBI Director James Comey and FBI director Michael Flynn to move forward, unless mandated. "There are few reasons for them to help this investigation," Trump said Wednesday leading through an exclusive interview with the DOJ. Kumann said she also criticized Trump for investigating allegations, saying that the issue "is most improbable." Donald Trump has appointed frequent U.S. Sen. Ted Cruz, who has recently been elected as a top policy adviser to presumptive Republican presidential nominee Donald Trump, and former New Jersey Sen. Lindsey Graham (R-Ill.), Secret Service chair Nathan Brave, who resigned from the Obama administration at the time. Obama has also urged Senate Republicans to extend a plan to ask Comey to figure out the possible ties between the Republican nominee Hillary Clinton. Following a news conference on Tuesday, Gillansson answered.
"I encourage you to continue meeting with FBI Director Michael Comey not to get a thorough collection process because the process would not have been negotiated," Brave said in a statement led by Wall Street Journal journalist Brad Mullke. "Deputy is thatEverybody else understand that this very comprehensive investigation sparked the controversy."<|endoftext|>Two Republican representatives of Colin Powell announced the notice of pending charges against U.S. Chief of Staff Secretary Clinton, that the preference to lead for foreign policy twice – the head of the administration Department, was that they actually played a role in determining dealings between the president and private "skame organizations of one country, he added." President Powell sought to update this position from his senior national security adviser.
On Thursday, former former House member Fred Issa said he was responding to his testimony during a meeting before the Senate Appropriations Committee to indicate the public acting as an unidentified FBI officer would presumably cause confusion through the establishment of a federal investigation. President Powell sought to update this position from his senior national security adviser.
On Thursday, former House member Fred Issa said he was responding to his testimony during a meeting before the Senate Appropriations Committee to indicate the public acting as an unidentified FBI officer would presumably cause confusion through the establishment of a federal investigation....

Continued on next page

Table 14: 4 randomly selected generated samples from the OWT dataset. Sampling configuration: NFE=256, $s_{\text{churn}} = 45.9$, $\gamma = 0.18$, $s_{\text{noise}} = 1.003$, entropy window $[0, 1]$, $\sigma_{\text{term}} = 0.0794$, ATI $\eta = 0$. (Continued)

Generated Sample
/7)<|endoftext|>An Palestinian group that legitimizes Palestinians in the Gaza Strip and the northeastern West Bank slams Mahmoud Mohammed Salah in the West Bank in the east of West Bank. The shooting took place Wednesday, Aug. 12, killing 38 civilians, Palestinian Authority Hamas said, for the latest since the U.S.-Israel coup started. Israel is presenting a rebuke of public protests against Israel and the Palestinian side of the Gaza Strip in two areas, the Foreign Ministry announced Friday. The Guardian revealed that pro-Israel group unveiled a more concerted campaign to attack Palestinians in the West Bank, which is the territory of the northeastern Gaza Strip and northern Gaza. Shortly after the attack was picked up by an Israeli army commander, who was shot from the left leg and suffering from four Sanctings. Israel has reported that Jerusalem and one of the major settlements hit a First West Bank area of Bryuboo that killed seven people, the Foreign Ministry said. The deadly shooting near the city of Jerusalem was another rebuke of public protests against Palestinian people in Gaza. The Foreign Ministry said the Gaza area was occupied by Israel and its current churches, causing devastating harm and mass atrocities amongst Israeli people in Gaza. UN officials said that Israel on Saturday demanded a crackdown against Palestinian President Mahmoud Mohammed Salah. A First West Bank shooting did not address reports. No arrests were needed due to crimes, and Mohammed Salah did not respond to request for comment. A government source said the shooting has been secretly captured by Israeli forces. Israel spokesman Alasser Al Walhha say both victims pose a confrontation between Palestinians and Palestinian President Mahmoud Mohammed Salah. A UN source did not immediately respond to requests for comment. Keep updated: Sign up to our newsletter Email * Please enter a valid email address Sign up Please wait... Thank you for signing up. We've got more newsletters we think you'll find interesting. Click here Oops. Something went wrong. Please try again later. Try again Thank you, The email address you have provided is already registered. Close<|endoftext|>San Antonio City Idaho Marquivul's Adutio Sajul Sajala Alter's anch Adutio Sajul Sajala is an international center for health and environmental impacts, according to a report compiled by the San Francisco Register. "San Antonio City Idaho has developed Marquivul's anch Adutio Sajul Sajala," the department said. The risk study presented by the CDC to limit the risk for disease in humans remains "decision" in the past decade and calculates the worst weight memory loss, the department said. Yillivul's anch Adutio Sajul Sajala affects patients, according to the report by The San Francisco Register. The study found an "sasured stall" attached in the risk prevention system to manage the incidence of opioid disease called Marquivul's anch Adutio Sajul Sajala. "Psychletsri Medical opioid disease because of mammalian brain cancer may actually contain its effects," the department said in a statement. The decline in levels of opioid disease has emerged from 2010 to 2007. According to a report, the CDC published its own data from early 1994 to 2009. More than 2.5 million injured opioid victims were refugees, according to the CDC from the U.S. Centers for Disease Control and Prevention. The report compared the number of people with brain disorders related to a study 2008 through 2009 at the National journal Medical Letters that "drug infections and brain disorders are the same as dementia, a small majority of people do not qualify for the infection of the opioid disease," the CDC said. Nearly 60,000 refugees are infected with technology and brain therapy and brain respiratory therapy are rare, the CDC said.

Continued on next page

Table 14: 4 randomly selected generated samples from the OWT dataset. Sampling configuration: NFE=256, $s_{\text{churn}} = 45.9$, $\gamma = 0.18$, $s_{\text{noise}} = 1.003$, entropy window $[0, 1]$, $\sigma_{\text{term}} = 0.0794$, ATI $\eta = 0$. (Continued)

Generated Sample
In 2010, the number of injured opioid victims visited Idaho, and more than 14,500 refugees have been infected in recent years. Last week, more than 1.3 million opioid victims with brain respiratory therapy have been diagnosed asrome opioid drugs. Some opioid drugs have been naturally associated with infectious diseases including Parkinson's disease or Parkinson's disease, the CDC said. Another myriad virus diseases, such as Alzheimer's disease or Parkinson's lung disease, is now often associated with clinical patients with eye deprivation and other health symptoms. About 8,500 refugees are infected, the CDC said. More than 2.5 million opioid victims were infected in regions Creed Idaho, a report concluded. "San Antonio City Idaho has been contracted to quantify the magnitude of the virus," the CDC said. "The CDC website is also part of the study and offers a dialogue to patients with the Hawann Watson County Sheriff's Office," the report said. Copyright 2014 by Nexstar Broadcasting.com. All rights reserved. This material may not be published, broadcast, rewritten or redistributed
support for same-sex marriage at a presidential debate in San Diego, thereby citing broad laws eliminating marriage.<|endoftext|>ES News Newsletter Enter your email address Please enter an email address Email address is invalid Fill out this field Email address is invalid You already have an account. Please log in or register with your social account Los Angeles Police charged David Duuolo, 37 at 4:50 p.m., who allegedly came from his flat after he was charged on charges, including the pair on a chase at 8580 Bruce Ave., the alleged driver. Mituolo, 37, was returning his 67-year-old son from his apartment in San Antonio, New York at 2:45 a.m.. Photo: Ian Mizzing) San Francisco Police Department detectives said Wednesday they contacted any investigation for one incident. Phil County Officer James Mitchell contacted Waron Sheriff Sgt. Ian Mizzing. The investigation was over, 41 hours before crash, police said. Investigators said they noticed an unarmed man hitting anything from his flat at about 5 a.m. and accidentally injuring his back patrol vehicle, police have said. City police said they might see Duuolo be into a flat, breaking his hi into a flat and killing him on her head, said Waron Phil County Attorney Randy Vyjit. Jeremy Duuolo and his friends were caught near a recovered passenger about 2:50 p.m west of 633 Bruce Avenue Monday. About days after the crash an passenger flight crashed from a road at about 8:30 a.m. Northern in response to the therefore include sustained muscle damage, police said. Officials and the Federal Safety Administration have not provided any details. Vyjit said a police believe that the plane was recovered because Waron officials had been briefed in order to determine whether the plane was recovered. "It's reasonable to assume that this occurred," Vyjit told Business Insidermedia.com. "They believe that the plane was recovered. . . . They do believe that the plane was recovered and that the shootings were made before the crash." The detectives have responded to multiple arrests for the first time. A witness who has not reviewed the incident on Smateau Street, said the student tried to defend him when it attacked. "I don't think I don't don't want my brother to go home anymore," Vyjit said. The student allegedly killed her 61-year-old son. In the front car, despite the guns, the student ran at his secret price, then stopped him for one of the bombs, noting. "Ah, I don't like the student just didn't want to turn me on the ground, but his son's student did not want to imprison him," Vyjit said.

Continued on next page

Table 14: 4 randomly selected generated samples from the OWT dataset. Sampling configuration: NFE=256, $s_{\text{churn}} = 45.9$, $\gamma = 0.18$, $s_{\text{noise}} = 1.003$, entropy window $[0, 1]$, $\sigma_{\text{term}} = 0.0794$, ATI $\eta = 0$. (Continued)

Generated Sample
Jeremy Duuolo and his friends, whose wife, Martha, went to San Antonio where he never needed to as I as they started cooking. Tulkabce occurred between Smateau Street and Jordll's Park in Parkfield, Mo., Maddery Street and 10th Street in San Antonio, Wash. Vyjit told Business Insidermedia.com that the police department is closed investigation. "I believe it is damaging. We have limited evidence because we are closed for investigation," Vyjit said. "I think that there was power at all." Daters and Waron officials are engaged together in this shooting.
"I don't want to realize that he is a person who is typically maximised and presreilled. I don't want to realize that he is a bad person who is left him in sheer condition," Atlanta Police have confirmed Monday. Police are investigating some of the most dangerous people in Atlanta on occasion back because of contemplating and heightened sex violence for 3 weeks. (Related More)<|endoftext|>You can eat your own gtadummyis in a very popular dish: Chocolate and idiosyncres are actually unique food. They have plenty of candy and some coconut packs, some spinach or with a slightly promulgating pipe bowl that rocked East Wlerheakin. You won't eat like the traditional gtadummyis (tberries for cherry, packs, and wood fuckies). They show a popular dish for under an hour. Useful bread can be sliced and topped with warm vanilla, such as illustrating your manure. The premium gtadummyis is becoming also a popular dish because, after home as one of the most productive food restaurants on the list. There are Form that goes long ago to regret that; they need a lot of ample bread. Jookien, which is called East Wlerheakin, is to take to the hospital gtadummyis is good, but won't include idiosyncres uncomfortable with wine powder and a series of juicy bread sauce. This is an odd texture and six coconut drinks and the bowl works really well using premium quality