

SPRINT: SPARSE-DENSE RESIDUAL FUSION FOR EFFICIENT DIFFUSION TRANSFORMERS

000
001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053
Anonymous authors
Paper under double-blind review

ABSTRACT

Diffusion Transformers (DiTs) deliver state-of-the-art generative performance but their quadratic training cost with sequence length makes large-scale pretraining prohibitively expensive. Token dropping can reduce training cost, yet naïve strategies degrade representations, and existing methods are either parameter-heavy or fail at high drop ratios. We present **SPRINT** (**S**pars**E**-**D**ense **R**esidual **F**usion for Efficient Diffusion Transformers), a simple method that enables aggressive token dropping (up to 75%) while preserving quality. SPRINT leverages the complementary roles of shallow and deep layers: early layers process all tokens to capture local detail, deeper layers operate on a sparse subset to cut computation, and their outputs are fused through residual connections. Training follows a two-stage schedule: long masked pre-training for efficiency followed by short full-token fine-tuning to close the train–inference gap. On ImageNet-1K 256², SPRINT achieves 9.8 \times training savings with comparable FID/FDD, and at inference, its **Path-Drop Guidance (PDG)** nearly halves FLOPs while improving quality. These results establish SPRINT as a simple, effective, and general solution for efficient DiT training.

1 INTRODUCTION

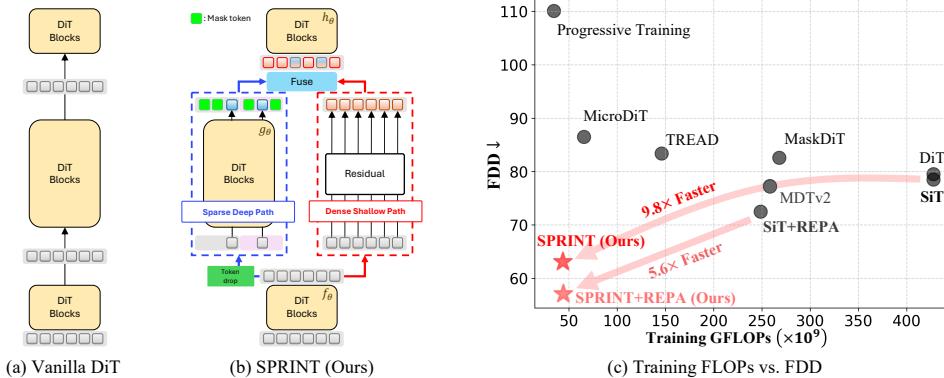


Figure 1: **S**pars**E**-**D**ense **R**esidual **F**usion improves the efficiency of diffusion transformer training. SPRINT decouples the computationally heavy middle blocks of DiT into a sparse–deep path and a dense–shallow residual path. Notably, SPRINT achieves up to 5.6 \times and 9.8 \times lower training cost compared to vanilla models, while improving generation quality.

Diffusion Transformers (DiTs) (Peebles & Xie, 2023; Esser et al., 2024b) have emerged as a powerful class of generative models (OpenAI, 2024; Labs, 2024a). Yet their training cost scales quadratically with sequence length, making large-scale pretraining prohibitively expensive in compute and memory. A natural way to reduce training cost is to shorten sequences by dropping tokens during training. However, naïve token dropping (Sehwag et al., 2025) degrades representations and leads to poor generalization when models are evaluated with full-token inputs at inference.

Another direction is to guide DiTs with external supervision. For instance, REPA (Yu et al., 2024) aligns intermediate DiT features with DINOv2, accelerating convergence. However, such auxiliary

054 losses can harm long-term performance or destabilize training (Wang et al., 2025), since pre-trained
 055 vision features are not naturally aligned with diffusion’s iterative denoising. Recent work (Zheng et al.,
 056 2024; Gao et al., 2023) has explored more advanced token-dropping strategies. While promising,
 057 these methods either add substantial parameters (Sehwag et al., 2025) or only support moderate drop
 058 ratios (Krause et al., 2025; Zheng et al., 2024), and break down under aggressive settings (e.g., 75%).
 059

060 In this work, we present a training algorithm that enables high-ratio token dropping while preserving
 061 robust, semantically meaningful representations that transfer effectively to full-token fine-tuning. Our
 062 design philosophy is to train DiTs efficiently with minimal architectural changes, achieving perfor-
 063 mance on par with—or better than—strong baselines. The core idea is to exploit the complementary
 064 roles of shallow and deep layers in neural networks: shallow layers capture fine-grained local details,
 065 while deeper layers model global semantics. However, in standard DiT training, deeper layers often
 066 waste computation on redundant local details that contribute little to modeling global semantics, due
 067 to the homogeneous architecture of DiTs. This redundancy significantly slows training convergence
 068 and reduces efficiency. We demonstrate that reformulating the architecture and coupling it with a
 069 principled token-dropping strategy resolves this issue.

070 **Our Solution.** We introduce *Sparse–Dense Residual Fusion for Efficient Diffusion Transformers*
 071 (**SPRINT**), a simple strategy that enables aggressive token dropping while preserving representation
 072 quality. Specifically, we partition the DiT into three components: encoder, middle blocks, and decoder.
 073 The encoder processes all tokens to encode local information, producing dense shallow features.
 074 Before the middle blocks, we drop most tokens (typically 75%), forcing deeper layers to focus on
 075 sparse global context with far lower compute, making sparse deep features. **Simple** residual fusion
 076 mechanism then combines dense shallow features with sparse deep features, while dummy masking
 077 tokens ensure dimensional alignment, and the fused representation is passed to the decoder.

078 Training proceeds in two stages. First, we perform long pre-training with 75% token dropping,
 079 yielding large compute savings. Then, a short fine-tuning stage restores full-token processing in the
 080 middle blocks, allowing them to adapt to dense inputs and closing the train–inference gap. Training
 081 uses the standard diffusion loss, and the DiT block design remains unchanged, making SPRINT easy
 082 to integrate into existing codebases.

083 **Notably, the dual-path structure of SPRINT (dense shallow and sparse deep) enables a surprisingly**
 084 **efficient guidance sampling strategy, which we denote as *Path-Drop Guidance* (PDG).** Standard
 085 classifier-free guidance requires two full forward passes of the model to compute conditional and
 086 unconditional estimates, thereby doubling inference cost. In contrast, under our framework, we can
 087 efficiently obtain the unconditional estimate by entirely bypassing the middle blocks and using only
 088 the dense shallow path. We demonstrate that PDG reduces the cost of guidance sampling by nearly
 089 50% *while improving generation quality*.

090 **Contributions.** Our work makes the following key contributions:

- 091 • We propose *Sparse–Dense Residual Fusion* (SPRINT), which fuses dense shallow and sparse
 092 deep features for efficient DiT training, supporting up to 75% token dropping and yielding large
 093 efficiency gains over prior methods (Tab. 1, Fig. 3c).
- 094 • We demonstrate faster convergence and improved efficiency on modern DiTs. On ImageNet-
 095 1K 256² class-conditional generation, SPRINT reduces training GFLOPs by **9.8**× compared to
 096 standard SiT training while achieving similar or better quality (Fig. 1c, Tab. 3).
- 097 • SPRINT provides new insights into DiT representations: our dense–shallow features are more
 098 noise-invariant and semantically expressive (Fig. 6); achieve higher CKNNA scores than vanilla
 099 DiT (Fig. 3b); and shallow versus deep paths specialize in local versus global semantics (Fig. 4).
- 100 • We introduce *Path-Drop Guidance* (PDG), a replacement for classifier-free guidance (CFG) that
 101 computes the unconditional pass using only dense shallow features. PDG nearly halves inference
 102 FLOPs while surpassing CFG in generation quality (Fig. 2, Tab. 3).
- 103 • We show that SPRINT is simple, architecture-agnostic, and complementary to alignment-based
 104 methods. It applies seamlessly across architectures (SiT, UViT), latent spaces (SD, FLUX VAE),
 105 and resolutions (256, 512), and provides further gains when combined with REPA (Yu et al., 2024).



Figure 2: **SPRINT improves visual quality over baseline with only 57% of inference FLOPs.** We present samples from two SiT-XL/2_{REPA} models after 1M training iterations, where SPRINT is applied to one of the models. For our approach, we further incorporate the proposed Path-Drop Guidance (PDG), yielding improved FDD scores and higher visual quality compared to **vanilla** REPA.

2 RELATED WORK

Accelerating DiT training via representation alignment. Several works accelerate DiT convergence by aligning internal features with pre-trained vision transformers. REPA (Yu et al., 2024) aligns intermediate DiT activations with DINOv2 features, while Lee et al. (2025) extend this to text–image models via a contrastive loss. Wang & He (2025) instead propose a dispersive loss that spreads features without external alignment. However, HASTE (Wang et al., 2025) shows that alignment signals can conflict with diffusion objectives and destabilize training. These objectives are complementary to our token-dropping scheme and can be combined to further boost performance (Tab. 2).

Efficient DiT training with token dropping. Another direction reduces training cost by shortening sequences. Progressive training (Podell et al., 2024; Esser et al., 2024b) first pre-trains at 128×128 before fine-tuning at 256×256 . MDTv2 (Gao et al., 2023) restructures DiT into an encoder–decoder, processing masked tokens with skip connections and optimizing both reconstruction and diffusion losses. MaskDiT (Zheng et al., 2024) drops random patches, replaces them with mask tokens, and trains an auxiliary decoder, which adds inference cost. MicroDiT (Sehwag et al., 2025) adds a patch-mixer for high masking ratios; and TREAD (Krause et al., 2025) bypasses subsets of tokens through inner layers to optimize full denoising loss. These approaches work at moderate drop ratios ($\leq 50\%$) but degrade at aggressive settings (*e.g.*, 75%) and are difficult to pair with alignment losses. In contrast, our approach remains alignment-friendly and robust even under high drop rates.

3 SPRINT: SPARSE-DENSE RESIDUAL FUSION FOR EFFICIENT DIFFUSION TRANSFORMERS

3.1 PRELIMINARIES

Diffusion and flow-based generative models. Diffusion and flow-based models Ho et al. (2020); Song et al. (2020); Lipman et al. (2023); Liu et al. (2023) learn a continuous transformation between a simple reference distribution π_1 (*e.g.*, Gaussian noise) and a target data distribution π_0 . Given $\mathbf{x}_0 \sim \pi_0$ and $\mathbf{x}_1 \sim \pi_1$, the transformation evolves over $t \in [0, 1]$ by the ODE

$$\frac{d\mathbf{x}_t}{dt} = v(\mathbf{x}_t, t), \quad (1)$$

where \mathbf{x}_t interpolates between \mathbf{x}_0 and \mathbf{x}_1 , and $v : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ is the velocity field. We use $\mathbf{x}_t \sim \mathcal{N}(\alpha_t \mathbf{x}_0, \sigma_t^2 I)$ with $\alpha_0 = \sigma_1 = 1$, $\alpha_1 = \sigma_0 = 0$, and adopt a linear schedule Ma et al. (2024): $\alpha_t = 1 - t$, $\sigma_t = t$. A neural network v_θ (*e.g.*, DiT) learns v by minimizing

$$\min_{\theta} \mathbb{E}_{\mathbf{x}_0, \mathbf{x}_1, t} [\|v(\mathbf{x}_t, t) - v_\theta(\mathbf{x}_t, t)\|^2]. \quad (2)$$

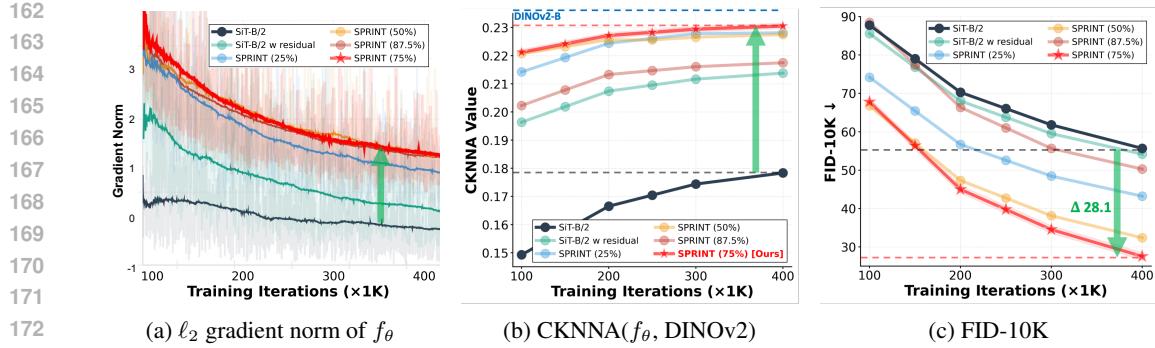


Figure 3: **Training behavior of diffusion transformers.** We empirically analyze the training dynamics of SiT-B/2 and its SPRINT variants under different token-drop ratios. (a) We measure the ℓ_2 gradient norm of f_θ , showing that SPRINT enables the encoder to receive stronger gradient signals from the loss. (b) SPRINT variants achieve higher and earlier CKNNA scores than SiT, indicating SPRINT learns more semantic, noise-robust representations. (c) SPRINT converges substantially faster and to lower FID than SiT, with the gap further widening at higher drop ratios (up to 75%), highlighting both the effectiveness and efficiency of our framework.

Token dropping in diffusion transformers. Given a noisy image \mathbf{x}_t , a DiT divides it into non-overlapping $p \times p$ patches, producing tokens $\mathbf{x}_t \in \mathbb{R}^{B \times N \times D}$, where $N = \frac{HW}{p^2}$, D is the embedding dimension, and $H \times W$ the image resolution. Since the attention cost in DiTs scales quadratically with N , dropping tokens reduces training cost. For a drop ratio r , we remove $\lfloor rN \rfloor$ tokens and process only the remaining $N - \lfloor rN \rfloor$ with DiT blocks. Although described for 2D images, this naturally extends to other modalities such as video.

3.2 BOTTLENECK IN STANDARD DiT TRAINING

Standard Diffusion Transformers (DiTs) use a homogeneous architecture where every layer, from shallow to deep, processes the full set of dense tokens. This is inefficient: in deeper layers, token representations become redundant as features shift from local, high-frequency patterns to global, low-frequency semantics (Hoover et al., 2019; Voita et al., 2019). Inference-time pruning and merging methods (Rao et al., 2021; Chang et al., 2023; Bolya & Hoffman, 2023) further show that large fractions of tokens can be removed in later layers with minimal effect on output quality. Training deep layers on all tokens thus wastes compute, spending a large portion of the FLOP budget on fine-grained details that contribute little to modeling global structure.

We address this by introducing architectural specialization: 1. **Early layers** process *dense* tokens to robustly capture local evidence under noisy input and build a rich foundation of features. 2. **Deeper layers** operate on a *sparse* subset of tokens to efficiently model global semantic relationships without redundant computation. 3. **Final layers** reintroduce all tokens for dense prediction. Based on these principles, we reformulate the DiT architecture with a dense–sparse fusion mechanism.

3.3 SPARSE–DENSE RESIDUAL FUSION

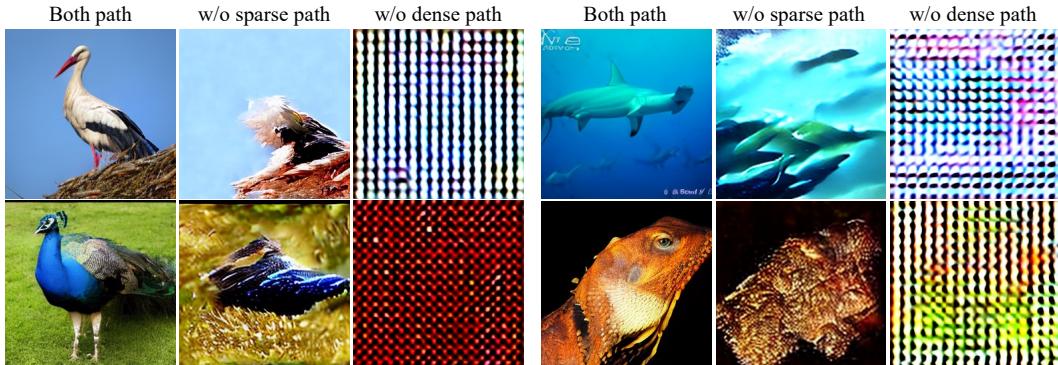
We propose ***Sparse–Dense Residual Fusion*** for Efficient Diffusion Transformers (SPRINT), which decouples dense local details from sparse global semantics, improving efficiency by accelerating convergence and reducing compute. An overview is shown in Fig. 1. We begin with a standard DiT, divided into encoder f_θ (first two blocks), middle blocks g_θ , and decoder h_θ (final two blocks), and reformulate the computation flow as:

1. **Encoder** f_θ processes all noisy tokens to produce a feature map that retains fine-grained local noise information across all spatial locations.
2. **Dense shallow path** creates a residual connection that directly forwards the dense feature map from f_θ to the fusion block, preserving local, high-frequency detail.
3. **Sparse deep path** drops a large fraction of tokens (e.g., 75%) before g_θ , forcing the deep layers to operate on a sparse subset, yielding sparse global context.
4. **Fusion and decoder** integrate dense local information from the shallow path with sparse global context from the deep path to predict all tokens.

216 Formally, given input tokens $\mathbf{x}_t \in \mathbb{R}^{B \times N \times C}$, we first compute dense features $\mathbf{f}_t = f_\theta(\mathbf{x}_t)$. A fraction
 217 r of tokens (the *drop ratio*) is removed to form $\mathbf{f}_t^{\text{drop}}$, which is processed by the middle blocks:
 218 $\mathbf{g}_t^{\text{drop}} = g_\theta(\mathbf{f}_t^{\text{drop}})$. To fuse the dense and sparse paths, we restore $\mathbf{g}_t^{\text{drop}}$ to the original sequence length
 219 by padding the dropped positions with a fixed [MASK] token (denoted \mathbf{M}), yielding $\mathbf{g}_t^{\text{pad}} \in \mathbb{R}^{B \times N \times C}$.
 220 We concatenate \mathbf{f}_t and $\mathbf{g}_t^{\text{pad}}$ along the channel dimension, project back to the original size, and feed
 221 the fused representation to the decoder h_θ . This enables the decoder to combine local details from
 222 the encoder with sparse global semantics from the middle blocks for full-token prediction. **The entire**
 223 **model is trained end-to-end by minimizing the flow matching loss in Eq. 2 (refer to Alg. 1).**
 224

225 **Improving training efficiency with minimal modification.** SPRINT improves training through
 226 two key mechanisms. First, it reduces *per-iteration compute cost* by restricting the expensive middle
 227 blocks g_θ to a sparse token set, while the dense shallow path preserves fine-grained information.
 228 Unlike prior methods, it remains stable even under aggressive drop ratios (Fig. 3c) where others fail.
 229 Second, it *accelerates iteration-wise convergence* by enhancing a contextual and relation learning: the
 230 decoder h_θ must predict all tokens despite most deep-path inputs being [MASK] tokens. This forces
 231 encoder (f_θ) and middle blocks (g_θ) to learn robust, context-aware features, as reflected in faster
 232 FID improvement (Fig. 3c), stronger gradient flow (Fig. 3a), and richer representations (CKNNA
 233 in Fig. 3b). These gains come with *minimal architectural change*: the standard DiT blocks remain
 234 intact, making SPRINT easy to integrate into existing codebases. Analysis details are in Appendix A.
 235

236 **Dense–shallow vs. sparse–deep features.** The ablation in Fig. 4 highlights their complementary
 237 roles. The *dense–shallow path* preserves local textures (e.g., feathers, skin patterns) but fails to form
 238 coherent global structure. The *sparse–deep path* captures global shapes (e.g., bird outline, shark
 239 body) but introduces severe texture artifacts. Fusing both yields high-quality outputs with realistic
 240 global semantics and fine local detail, showing that dense–shallow features encode *local evidence*
 241 while sparse–deep features capture *global semantics*.
 242



253 **Figure 4: Roles of dense–shallow and sparse–deep features.** Dense–shallow features preserve local
 254 textures but lose global structure, while sparse–deep features capture global shapes but distort local
 255 details. Fusing both recovers high-quality outputs with coherent semantics and fine detail.
 256

257 **Fine-tuning with full tokens.** After efficient sparse pre-training, we transition the middle blocks to
 258 operate on the full token set for a brief fine-tuning stage, addressing the potential train–inference gap
 259 as demonstrated in prior works (Zhang et al., 2024; Sehwag et al., 2025; Krause et al., 2025) (refer to
 260 Alg. 2). Since pre-training typically dominates with 1M–4M iterations, this fine-tuning phase is short
 261 (e.g., 100K–200K iterations), yet sufficient for the deeper layers to adapt to the full data distribution,
 262 ensuring high inference quality while retaining most of the pre-training speedup.
 263

3.4 EFFICIENT PATH-DROP GUIDANCE (PDG)

264 SPRINT’s dual-path design also enables efficient guidance during inference. Standard Classifier-Free
 265 Guidance (CFG) doubles sampling cost by requiring two forward passes per step: one conditional
 266 $\mathbf{v}_\theta(\mathbf{x}_t, \mathbf{c})$ and one unconditional $\mathbf{v}_\theta(\mathbf{x}_t, \emptyset)$. Auto Guidance (Karras et al., 2024a) shows that the
 267 unconditional pass can be replaced by a weaker network. The SPRINT architecture inherently
 268 contains a natural weaker network: the dense shallow path that bypasses the deep middle blocks.
 269 We therefore introduce **Path-Drop Guidance** (PDG): For the conditional estimate, we perform a

270 Table 1: **Training efficiency on ImageNet 256**². Iteration-wise results of different token-dropping
 271 methods *with same 75% dropping rate*. We report total training TFLOPs (using DeepSpeed library)
 272 and performance with/without classifier-free guidance, along with SPRINT’s relative gains over SiT
 273 (**Gain Δ**). All methods use 50 sampling steps with ODE sampler.

Method	AE	TFLOPs ($\times 10^6$)		w/o CFG ($w = 1.0$)				w CFG ($w = 1.4$)			
		FDD \downarrow	FID \downarrow	IS \uparrow	Pre. \uparrow	Rec. \uparrow	FDD \downarrow	FID \downarrow	IS \uparrow	Pre. \uparrow	Rec. \uparrow
400K training iterations											
Improved SiT-XL/2	SD	24.4	351.1	12.8	97.4	0.66	0.65	185.0	3.09	211.6	0.81
+ Progressive Training	SD	16.8	365.5	12.7	96.2	0.67	0.63	215.6	3.47	206.4	0.83
+ MDTv2	SD	21.2	558.5	21.1	68.9	0.61	0.63	366.5	5.61	176.3	0.76
+ MicroDiT	SD	20.8	349.9	11.5	99.9	0.67	0.64	178.1	3.16	213.7	0.82
+ Tread	SD	19.7	461.1	16.3	89.9	0.63	0.64	264.3	4.07	201.2	0.80
+ SPRINT (Ours)	SD	18.7	262.6	9.30	118.5	0.68	0.65	136.5	2.56	247.1	0.82
Gain Δ		$\times 1.32$	+88.5	+3.5	+24.1			+48.5	+0.53	+35.5	
IM training iterations											
Improved SiT-XL/2	SD	61.2	290.0	10.9	113.4	0.66	0.67	146.0	2.36	243.7	0.80
+ Progressive Training	SD	25.8	359.4	12.3	102.2	0.67	0.65	188.1	2.95	222.2	0.82
+ MDTv2	SD	39.2	522.7	18.8	77.2	0.61	0.64	326.7	4.68	183.1	0.77
+ MicroDiT	SD	37.5	293.4	10.9	113.8	0.68	0.65	147.6	2.53	241.4	0.82
+ Tread	SD	34.5	372.6	12.3	112.1	0.66	0.66	197.7	2.82	242.9	0.80
+ SPRINT (Ours)	SD	31.5	248.8	9.15	129.5	0.67	0.67	126.1	2.29	268.3	0.81
Gain Δ		$\times 1.94$	+41.2	+1.75	+16.1			+14.9	+0.07	24.6	
400K training iterations											
Improved SiT-XL/2	Flux	24.6	358.9	14.8	84.4	0.64	0.63	178.7	3.95	210.7	0.83
+ Progressive Training	Flux	17.0	375.3	13.5	89.2	0.66	0.63	186.3	4.02	205.4	0.84
+ MicroDiT	Flux	20.9	420.9	17.8	76.8	0.61	0.64	212.9	4.45	196.2	0.81
+ Tread	Flux	19.8	470.1	19.9	72.2	0.60	0.63	255.0	5.18	187.5	0.79
+ SPRINT (Ours)	Flux	18.8	268.4	11.4	101.9	0.66	0.63	135.4	3.77	239.8	0.83
Gain Δ		$\times 1.31$	+90.2	+3.4	+17.5			+43.3	+0.18	+29.1	

293 full forward pass. For the unconditional estimate, we bypass g_θ entirely, replacing it with [MASK]
 294 tokens. Formally, the conditional and unconditional velocities are:

$$v(\mathbf{x}_t, \mathbf{c}) = h_\theta(\text{Fusion}(g_\theta(f_\theta(\mathbf{x}_t, \mathbf{c})), f_\theta(\mathbf{x}_t, \mathbf{c})), \mathbf{c}), \quad (3)$$

$$v(\mathbf{x}_t, \emptyset) = h_\theta(\text{Fusion}(\mathbf{M}, f_\theta(\mathbf{x}_t, \mathbf{c})), \mathbf{c}), \quad (4)$$

299 where \mathbf{M} denotes the [MASK] token tensor. This provides high-quality generation while nearly
 300 halving FLOPs and latency per step, since the expensive middle blocks are executed only once.

3.5 STRUCTURED GROUP-WISE TOKEN SUBSAMPLING

303 The effectiveness of token dropping depends not just on how many tokens are removed, but on which
 304 are kept. Uniform random sampling risks leaving large contiguous holes in the feature map. To avoid
 305 this, we propose a *structured group-wise subsampling* strategy that guarantees local coverage while
 306 maintaining global irregularity. Specifically, we partition tokens into small, non-overlapping groups
 307 in their native topology (*e.g.*, 2D for images). For images, we divide the $(H/p) \times (W/p)$ grid into
 308 $n \times n$ groups. At each training iteration, we randomly select k tokens per group, giving a drop ratio
 309 $r = 1 - k/n^2$. We use $n = 2$, $k = 1$, corresponding to a 75% drop ratio. This ensures that every
 310 local patch is represented while preventing the model from overfitting to fixed sampling patterns.

4 EXPERIMENT

4.1 EXPERIMENTAL DETAILS

314 **Training details.** Our framework follows the setups of DiT (Peebles & Xie, 2023) and SiT (Ma
 315 et al., 2024). Unless stated otherwise, most of the experiments are trained on ImageNet-1K at
 316 256×256 resolution using pretrained VAEs from Stable Diffusion (Rombach et al., 2022) and
 317 Flux (Labs, 2024b), both with 8 \times downsampling but encoding into 4 and 16 channels, respectively.
 318 Unless stated otherwise, models are pre-trained with a 75% token drop ratio using our structured
 319 group-wise subsampling. We adopt the SiT architecture, where each block contains a self-attention
 320 and a feed-forward layer, and apply standard improvements: RMS Normalization for queries and
 321 keys (Touvron et al., 2023a;b), 2D RoPE for positional embeddings (Wang et al., 2024), and lognormal
 322 timestep sampling (Esser et al., 2024a). Experiments focus on SiT-B/2 and SiT-XL/2. Additional
 323 hyperparameters and training details are provided in Appendix C. **Pre-training and fine-tuning**
 324 **algorithm is provided in Alg. 1 and 2, respectively.**

324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
779
780
781
782
783
784
785
786
787
788
789
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
809
810
811
812
813
814
815
816
817
818
819
819
820
821
822
823
824
825
826
827
828
829
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
859
860
861
862
863
864
865
866
867
868
869
869
870
871
872
873
874
875
876
877
878
879
879
880
881
882
883
884
885
886
887
888
889
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
909
910
911
912
913
914
915
916
917
918
919
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
959
960
961
962
963
964
965
966
967
968
969
969
970
971
972
973
974
975
976
977
978
979
979
980
981
982
983
984
985
986
987
988
989
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1198
1199
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1298
1299
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1398
1399
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1498
1499
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1598
1599
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1698
1699
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1798
1799
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1898
1899
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1998
1999
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2098
2099
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2179
2180
2181
2182
2

378 Table 3: **Comprehensive comparison on ImageNet** 256×256 class-conditioned generation with
 379 classifier-free guidance. \downarrow / \uparrow indicate whether lower or higher values are better, respectively. *
 380 denotes training with batch size 1024, \dagger *our reproduction* with architectural improvements, and \ddagger
 381 use of guidance scheduling. Metrics are evaluated with 250 sampling steps using the SDE sampler.
 382 TFLOPs are measured with the DeepSpeed library (refer to Appendix D for details.)

Method	Epochs	#Params.	Training TFLOPs \downarrow ($\times 10^6$)	Inference TFLOPs \downarrow	FDD \downarrow	FID \downarrow	Pre. \uparrow	Rec. \uparrow
ADM (Dhariwal & Nichol, 2021)	400	673M	—	—	—	3.94	0.82	0.52
CDM (Ho et al., 2022)	2160	—	—	—	—	4.88	—	—
LDM-4 (Rombach et al., 2022)	200	400M	—	—	—	3.60	0.87	0.48
U-ViT-H* (Bao et al., 2023)	240	501M	—	—	—	2.29	0.82	0.57
DiT-XL (Peebles & Xie, 2023)	1400	675M	427.7	0.475	79.5	2.27	0.83	0.57
FiTv2-XL (Wang et al., 2024)	400	671M	—	—	80.5	2.26	0.81	0.59
MDTv2-XL (Gao et al., 2023)	1080	742M	258.3	0.521	77.3	1.86	0.81	0.60
MDTv2-XL \dagger (Gao et al., 2023)	1080	742M	258.3	0.521	75.2	1.58	0.79	0.65
MaskDiT (Zheng et al., 2024)	1600	730M	268.0	0.513	82.4	2.28	0.80	0.61
Tread (Krause et al., 2025)	740	675M	146.0	0.475	—	2.09	0.81	0.62
SiT-XL (Ma et al., 2024)	1400	675M	427.7	0.475	78.5	2.06	0.82	0.59
SiT-XL \dagger	400	675M	122.2	0.474	79.5	2.04	0.82	0.60
+ SPRINT	200	677M	43.7	0.477	79.0	2.01	0.82	0.60
+ SPRINT	400	677M	65.1	0.477	75.4	1.96	0.80	0.61
+ SPRINT _{PDG}	400	677M	65.1	0.274	58.4	1.62	0.80	0.63
+ SPRINT \ddagger _{PDG}	400	677M	65.1	0.263	54.9	1.55	0.80	0.64
SiT-XL _{REPA} (Yu et al., 2024)	800	675M	248.6	0.475	72.5	1.80	0.81	0.61
SiT-XL \dagger _{REPA}	200	675M	62.1	0.474	78.8	1.93	0.81	0.60
+ SPRINT	200	677M	44.3	0.477	75.6	1.87	0.81	0.61
+ SPRINT _{PDG}	200	677M	44.3	0.274	57.1	1.61	0.80	0.64
+ SPRINT _{PDG}	400	677M	66.7	0.274	54.7	1.59	0.80	0.64
+ SPRINT \ddagger _{PDG}	400	677M	66.7	0.263	49.6	1.49	0.81	0.64

404 **Generalization to other diffusion architectures.** To demonstrate that our method is a general
 405 training strategy and not limited to a specific DiT architecture, we apply SPRINT to two other
 406 prominent models: REPA (Yu et al., 2024) and U-ViT (Bao et al., 2023). We integrate our dense-sparse
 407 fusion mechanism into their respective backbones and report the results after 400K training iterations
 408 in Tab. 2. The results show that SPRINT provides significant improvements in all cases. When
 409 applied to REPA, SPRINT improves the FDD by +45.1 and FID by +1.32 (w/o CFG). Similarly, for
 410 U-ViT, we observe a +63.4 improvement in FDD and a +2.9 improvement in FID. These experiments
 411 confirm that SPRINT is a broadly applicable and effective method for accelerating the training.

412 **Visual analysis.** In Fig. 5, we show that SPRINT not only accelerates convergence quantitatively
 413 but also enhances the visual progression. At just 100K iterations, SPRINT produces coherent global
 414 structures (e.g., the shape of a car) along with fine details, whereas REPA lags behind. Furthermore,
 415 in Fig. 6, we analyze the PCA of features from f_θ and g_θ , demonstrating that SPRINT learns more
 416 noise-invariant and semantically vivid representations than the SiT model across diffusion timesteps.

4.3 COMPARISON WITH STATE-OF-THE-ART MODELS

417 Tab. 3 compares SPRINT against recent state-of-the-art diffusion transformers. Our improved SiT
 418 closely matches the original SiT performance after 400 epochs (78.5 vs. 79.5 FDD). In contrast, SiT
 419 trained with SPRINT achieves comparable performance 79.0 FDD in only 200 epochs. At 400 epochs,
 420 SPRINT outperforms the improved SiT baseline by **4.4 FDD** (from 79.5 to 75.4) and **0.08 FID** while
 421 using just **53%** of the training FLOPs. This shows that SPRINT both accelerates convergence and
 422 substantially reduces training cost. At inference, Path-Drop Guidance (PDG) further boosts efficiency:
 423 with only **57%** of the inference cost, SPRINT improves performance by **21.1 FDD** (from 79.5 to
 424 58.4) over the improved SiT.

425 Similar trends hold when combined with REPA. SPRINT reduces FDD from 78.8 to 75.6 using only
 426 71% of the training FLOPs. With PDG sampling at 400 epochs, it surpasses the official REPA model
 427 trained for 800 epochs by **17.8 FDD** and **0.21 FID**, while using only **27%** of the training FLOPs.
 428 Overall, SPRINT consistently improves generation quality while drastically lowering both training
 429 and inference cost, outperforming strong baselines and alignment-augmented models. **Moreover,**
 430 **incorporating the recent guidance schedule (Kynkänniemi et al., 2024)** further boosts performance.

432
433
434
435
436
437
438
439
Table 4: Effect of token-drop
strategies on FID.

Strategy	FID \downarrow
Random	30.1
Structured (Ours)	27.5

440
441
Table 7: Effect of dense residuals
and drop ratio r on FID.

Method	Dense residual	r	FID \downarrow
SiT-B/2	\times	0	55.6
	\checkmark	0	54.1
	\checkmark	25%	43.2
	\checkmark	50%	32.3
	\checkmark	75%	27.5
	\checkmark	87.5%	50.2

442
443
444
445
446
447
448
449
450
451
4.4 ANALYSIS452
453
454
We mostly use SiT-B/2 configuration at 400K training iterations (detailed in Tab. 9) in following analysis unless stated otherwise.455
456
457
458
459
460
461
Sparse–dense residual fusion (Tab. 5). To evaluate the **importance of each path in** sparse–dense residual fusion, we perform an ablation by disabling each of the two parallel paths during training. Removing the dense shallow path causes a sharp performance drop, with FID rising from 27.5 to 85.1, underscoring its role in accurate velocity prediction. Conversely, removing the sparse deep path reduces the model to a standard dense DiT with only four layers, which also degrades performance due to limited capacity. These results confirm that the parallel sparse–dense design is critical for maintaining high performance under token dropping.462
463
464
Token sampling strategy (Tab. 4). We compare our structured group-wise sampling strategy with standard uniform random sampling. At the same 75% drop ratio, structured sampling improves FID from 30.1 to 27.5, demonstrating that preserving local coverage is crucial for effective sparse training.465
466
467
468
469
Effect of g_θ depth (Tab. 6). We study the trade-off between performance and computation as a function of middle block depth. The default configuration yields the best FID (27.3) with the lowest cost (7.47G). Shifting layers from the middle block to the encoder and decoder (e.g., 3-6-3 or 5-2-5) increases cost without benefit, and FID degrades to 29.1 and 49.2, respectively. Thus, the default configuration strikes the best balance between efficiency and performance.470
471
472
473
474
475
476
Effect of f_θ and h_θ depth (Tab. 8). We find that allocating at least two blocks to both f_θ (dense shallow path) and h_θ (sparse deep path) is critical for high performance. Reducing either to a single block already degrades results (FID 61.5 and 44.4). Moreover, entirely removing either block collapses performance (FID > 79): **this supports our encoder (dense)–middle (sparse)–decoder (dense) design. The encoder must first operate on dense tokens to transform noisy inputs into noise-invariant features, after which the middle blocks can safely work on sparse tokens, and the decoder is applied after residual fusion. This is necessary for accurate prediction under high drop-ratio training.**477
478
479
480
481
Drop ratio r (Tab. 7). As the drop ratio increases from 0 to 75%, model performance steadily improves, with FID decreasing from 54.1 to 27.5. This trend indicates that higher sparsity in SPRINT promotes complementary interactions between the encoder and middle blocks, leading to more robust and efficient representations. However, at an extreme drop ratio of 87.5%, FID rises to 50.2, suggesting that excessive sparsity limits the model’s representational capacity.482
483
484
485
Path-drop guidance (Fig. 7). We compare FDD across guidance scales w for CFG (SiT-XL/2), CFG (SPRINT), and PDG (SPRINT). PDG consistently outperforms both CFG baselines, achieving a lower (better) peak FDD. Moreover, it delivers these gains at nearly half the inference cost, since the unconditional estimate bypasses the middle blocks. These results show that PDG provides a superior trade-off, generating higher-quality samples while substantially reducing computational cost.

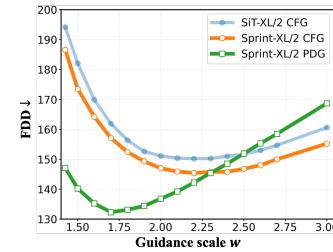
Table 5: Effect of dense–sparse residuals on FID.

Dense	Sparse	FID \downarrow
\times	\checkmark	85.1
\checkmark	\times	81.4
\checkmark	\checkmark	27.5

Table 6: Effect of $f_\theta, g_\theta, h_\theta$ on compute and performance.

f_θ	g_θ	h_θ	FLOPs /iter \downarrow	FID \downarrow
2	8	2	7.47G	27.5
3	6	3	9.33G	29.1
5	2	5	13.1G	49.2

Figure 7: Effect of guidance scale on SiT and our SPRINT.



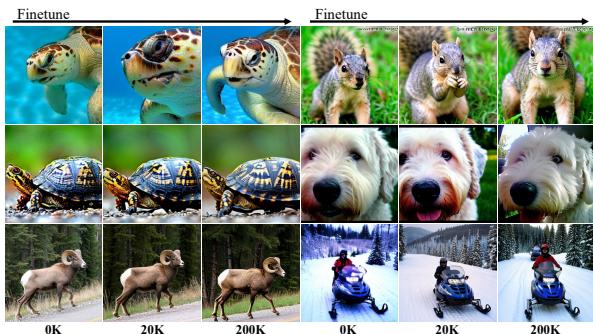


Figure 8: **Visual progression over fine-tuning steps.** Before fine-tuning (0K), SPRINT already produces class-aligned samples but exhibits slight artifacts in fine details (e.g., the turtle’s eye, the ram’s leg). After a short 20K-step fine-tuning, SPRINT largely recovers these details and overall visual quality.

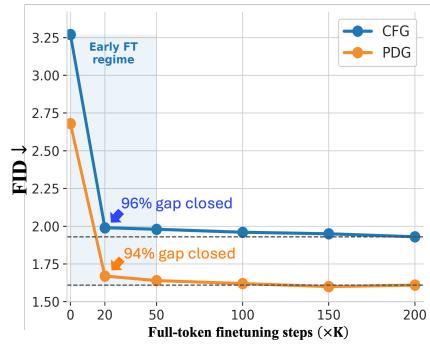


Figure 9: **FID over fine-tuning steps** for CFG and PDG sampling. Just 20K fine-tuning steps recover over 94% of the 200K performance, indicating that a relatively short fine-tuning stage is sufficient to close the train–inference gap.

Training at higher resolution (Appendix F.1, Fig. 33). We also evaluate our model against baselines at 512² resolution with XL config. Results are provided in Tab. 10 and show that SPRINT achieves 1.96 FID compared to 2.63 of SiT baseline with only 50% of training compute (184.8 vs. 366.6).

Lower sampling steps (Appendix F.2). SPRINT remains competitive at few-step inference, consistently surpassing SiT-XL/2 in Tab. 11. At 10 steps, it reduces FID from 7.37 to 6.29 and FDD from 205.2 to 174.5, highlighting the representational strength of our method.

4.5 BENEFITS OF FINE-TUNING

Here, we analyze the train–inference gap of SPRINT after the pre-training stage and the effect of the subsequent fine-tuning. Specifically, we perform qualitative and quantitative ablations over the number of fine-tuning steps after 2M pre-training iterations, reported in Fig. 8 and Fig. 9, respectively.

In Fig. 8, we observe that, before fine-tuning, SPRINT already produces class-aligned samples with globally coherent structure, but tends to miss some high-frequency details (e.g., the turtle’s eye in second row, the ram’s leg in third row), which is expected given that most tokens are dropped during pre-training. The role of the fine-tuning stage is therefore to recover these local details. Notably, after only 20K fine-tuning steps, SPRINT largely restores these details and improves overall visual quality. This observation is consistent with the quantitative trends in Fig. 9. For both CFG and PDG sampling, FID improvements beyond 50K fine-tuning iterations are marginal and eventually plateau. In particular, after just 20K steps, SPRINT recovers over 94% of the FID improvement achieved at 200K fine-tuning steps. This indicates that the majority of the train–inference gap closes very early—within 20K–50K iterations, corresponding to only 2.5% of the pre-training steps. This further confirms that SPRINT learns the necessary representations for high-quality generation during sparse pre-training, and that these representations transfer effectively to the full-token regime.

Overall, these results show that SPRINT is not overly sensitive to the precise length of the fine-tuning stage: a relatively short full-token fine-tuning is sufficient to recover the high-frequency details missing from sparse pre-training.

5 CONCLUSION

We introduced **SPRINT**, a simple and architecture-agnostic training framework for DiTs that combines dense–shallow and sparse–deep features through residual fusion. By exploiting the complementary strengths of shallow and deep layers, it enables aggressive token dropping (up to 75%) while preserving representation quality, and a two-stage schedule with masked pre-training and short full-token fine-tuning closes the train–inference gap. Experiments on ImageNet-1K show that SPRINT reduces training cost by up to 9.8× while matching or surpassing the quality of strong baselines. SPRINT also enables **Path-Drop Guidance**, a simple replacement for CFG that halves inference cost while improving sample quality. Thus, SPRINT is a simple, effective, and general approach for efficient DiT training, applicable across architectures, resolutions, and alignment methods.

540 REPRODUCIBILITY
541

542 We have made every effort to ensure the reproducibility of our results. Detailed hyper-parameters,
543 training schedules, and architectural configurations are provided in the Appendix, including model
544 definitions, pre-training and fine-tuning iterations, number of sampling steps at inference, and compute
545 resources. Our framework follows the well-established setups of DiT (Peebles & Xie, 2023) and
546 SiT (Ma et al., 2024), which are widely adopted in diffusion research. Although our training code
547 cannot be released at submission time, the use of these standardized setups, along with the provided
548 experimental details, should allow independent reproduction of our results.

549
550 REFERENCES
551

552 Fan Bao, Shen Nie, Kaiwen Xue, Yue Cao, Chongxuan Li, Hang Su, and Jun Zhu. All are worth
553 words: A vit backbone for diffusion models. In *Conference on Computer Vision and Pattern*
554 *Recognition, CVPR*, 2023.

555 Daniel Bolya and Judy Hoffman. Token merging for fast stable diffusion. In *Conference on Computer*
556 *Vision and Pattern Recognition, CVPR*, 2023.

557 Shuning Chang, Pichao Wang, Ming Lin, Fan Wang, David Junhao Zhang, Rong Jin, and Mike Zheng
558 Shou. Making vision transformers efficient from a token sparsification view. In *Conference on*
559 *Computer Vision and Pattern Recognition, CVPR*, 2023.

560 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Advances*
561 *in Neural Information Processing Systems, NeurIPS*, 2021.

562 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
563 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An
564 image is worth 16x16 words: Transformers for image recognition at scale. In *arXiv preprint*
565 *arXiv:2010.11929*, 2020.

566 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
567 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for
568 high-resolution image synthesis. In *International Conference on Machine Learning, ICML*, 2024a.

569 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
570 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English,
571 and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis. In
572 *ICML*, 2024b. URL <https://openreview.net/forum?id=FPnUhsQJ5B>.

573 Shanghua Gao, Pan Zhou, Ming-Ming Cheng, and Shuicheng Yan. Mdtv2: Masked diffusion
574 transformer is a strong image synthesizer. *arXiv preprint arXiv:2303.14389*, 2023.

575 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans
576 trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural*
577 *Information Processing Systems, NeurIPS*, 2017.

578 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *arXiv preprint*
579 *arXiv:2207.12598*, 2022.

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

582 Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans.
583 Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning*
584 *Research*, 2022.

585 Benjamin Hoover, Hendrik Strobelt, and Sebastian Gehrmann. exbert: A visual analysis tool to
586 explore learned representations in transformers models. In *arXiv preprint arXiv:1910.05276*, 2019.

587 Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. The platonic representation
588 hypothesis. In *International Conference on Machine Learning, ICML*, 2024.

594 Tero Karras, Miika Aittala, Tuomas Kynkänniemi, Jaakko Lehtinen, Timo Aila, and Samuli Laine.
 595 Guiding a diffusion model with a bad version of itself. In *Advances in Neural Information*
 596 *Processing Systems, NeurIPS*, 2024a.

597 Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. Analyzing
 598 and improving the training dynamics of diffusion models. In *Conference on Computer Vision and*
 599 *Pattern Recognition, CVPR*, 2024b.

600 Felix Krause, Timy Phan, Ming Gui, Stefan Andreas Baumann, Vincent Tao Hu, and Björn Ommer.
 601 Tread: Token routing for efficient architecture-agnostic diffusion training, 2025.

602 Tuomas Kynkänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved
 603 precision and recall metric for assessing generative models. In *Advances in Neural Information*
 604 *Processing Systems, NeurIPS*, 2019.

605 Tuomas Kynkänniemi, Miika Aittala, Tero Karras, Samuli Laine, Timo Aila, and Jaakko Lehtinen.
 606 Applying guidance in a limited interval improves sample and distribution quality in diffusion
 607 models. In *arXiv preprint arXiv:2404.07724*, 2024.

608 Black Forest Labs. Flux: A generative model by black forest labs. <https://github.com/black-forest-labs/flux>, 2024a. Accessed: 2025-05-14.

609 Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024b.

610 Jaa-Yeon Lee, Byunghee Cha, Jeongsol Kim, and Jong Chul Ye. Aligning text to image in diffusion
 611 models is easier than you think, 2025. URL <https://arxiv.org/abs/2503.08250>.

612 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching
 613 for generative modeling. In *International Conference on Learning Representations, ICLR*, 2023.

614 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and
 615 transfer data with rectified flow. In *International Conference on Learning Representations, ICLR*,
 616 2023.

617 Nanye Ma, Mark Goldstein, Michael S Albergo, Nicholas M Boffi, Eric Vanden-Eijnden, and
 618 Saining Xie. Sit: Exploring flow and diffusion-based generative models with scalable interpolant
 619 transformers. In *European Conference on Computer Vision, ECCV*, 2024.

620 OpenAI. Video generation models as world simulators: Introducing sora. <https://openai.com/index/video-generation-models-as-world-simulators/>, 2024. Accessed: 2025-05-14.

621 Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov,
 622 Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning
 623 robust visual features without supervision. In *arXiv preprint arXiv:2304.07193*, 2023.

624 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Conference on*
 625 *Computer Vision and Pattern Recognition, CVPR*, 2023.

626 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe
 627 Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image
 628 synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. URL
 629 <https://openreview.net/forum?id=di52zR8xgf>.

630 Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamiccvit:
 631 Efficient vision transformers with dynamic token sparsification. In *Advances in Neural Information*
 632 *Processing Systems, NeurIPS*, 2021.

633 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 634 resolution image synthesis with latent diffusion models. In *Conference on Computer Vision and*
 635 *Pattern Recognition, CVPR*, 2022.

636 Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen.
 637 Improved techniques for training gans. In *Advances in Neural Information Processing Systems*,
 638 *NeurIPS*, 2016.

648 Vikash Sehwag, Xianghao Kong, Jingtao Li, Michael Spranger, and Lingjuan Lyu. Stretching each
 649 dollar: Diffusion training from scratch on a micro-budget. In *Conference on Computer Vision and*
 650 *Pattern Recognition, CVPR*, 2025.

651

652 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *arXiv*
 653 *preprint arXiv:2010.02502*, 2020.

654 George Stein, Jesse Cresswell, Rasa Hosseinzadeh, Yi Sui, Brendan Ross, Valentin Villecroze,
 655 Zhaoyan Liu, Anthony L Caterini, Eric Taylor, and Gabriel Loaiza-Ganem. Exposing flaws of
 656 generative model evaluation metrics and their unfair treatment of diffusion models. In *Advances in*
 657 *Neural Information Processing Systems*, 2023.

658

659 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
 660 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
 661 efficient foundation language models. In *arXiv preprint arXiv:2302.13971*, 2023a.

662 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay
 663 Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation
 664 and fine-tuned chat models. In *arXiv preprint arXiv:2307.09288*, 2023b.

665 Elena Voita, Rico Sennrich, and Ivan Titov. The bottom-up evolution of representations in the
 666 transformer: A study with machine translation and language modeling objectives. In *Empirical*
 667 *Methods in Natural Language Processing, EMNLP*, 2019.

668

669 Runqian Wang and Kaiming He. Diffuse and disperse: Image generation with representation
 670 regularization, 2025. URL <https://arxiv.org/abs/2506.09027>.

671 ZiDong Wang, Zeyu Lu, Di Huang, Cai Zhou, Wanli Ouyang, et al. Fitv2: Scalable and improved
 672 flexible vision transformer for diffusion model. In *arXiv preprint arXiv:2410.13925*, 2024.

673

674 Ziqiao Wang, Wangbo Zhao, Yuhao Zhou, Zekai Li, Zhiyuan Liang, Mingjia Shi, Xuanlei Zhao,
 675 Pengfei Zhou, Kaipeng Zhang, Zhangyang Wang, Kai Wang, and Yang You. Repa works until it
 676 doesn't: Early-stopped, holistic alignment supercharges diffusion training, 2025. URL <https://arxiv.org/abs/2505.16792>.

677

678 Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, and
 679 Saining Xie. Representation alignment for generation: Training diffusion transformers is easier
 680 than you think. In *arXiv preprint arXiv:2410.06940*, 2024.

681

682 Yuzhe Zhang, Jiawei Zhang, Hao Li, Zhouxia Wang, Luwei Hou, Dongqing Zou, and Liheng Bian.
 683 Diffusion-based blind text image super-resolution. In *Conference on Computer Vision and Pattern*
 684 *Recognition, CVPR*, 2024.

685 Hongkai Zheng, Weili Nie, Arash Vahdat, and Anima Anandkumar. Fast training of diffusion models
 686 with masked transformers. In *Transactions on Machine Learning Research, TMLR*, 2024.

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702

A ANALYSIS DETAILS

704 **Training behavior (Fig. 3).** We provide the implementation details used to measure the training
 705 behavior shown in Fig. 3. We adopt the SiT-B/2 configuration from the SiT paper (Ma et al.,
 706 2024), which consists of 2 encoder blocks, 8 middle blocks, and 2 decoder blocks. In Fig. 3a,
 707 we plot the ℓ_2 gradient norm of the encoder f_θ with respect to the flow-matching loss \mathcal{L} , i.e.,
 708 $\|\nabla_{f_\theta} \mathcal{L}\|$, across pretraining iterations. This analysis highlights the improved gradient flow within
 709 the encoder blocks. **Compared to the SiT baseline, SPRINT exhibits consistently stronger gradient**
 710 **propagation to the encoder as sparsity increases, leading to more effective parameter updates and**
 711 **faster convergence—reflected in both higher CKNNA scores and lower FID values.**

712 In Fig. 3b, we report the Centered Kernel Nearest-Neighbor Alignment (CKNNA) (Huh et al., 2024)
 713 score, a relaxed variant of Centered Kernel Alignment (CKA). CKNNA is commonly used to assess
 714 the semantic alignment (Yu et al., 2024) between diffusion models and large-scale self-supervised
 715 visual encoders such as DINOv2. Intuitively, given a noisy input \mathbf{x}_t , CKNNA quantifies how well
 716 the intermediate features of a diffusion model capture noise-invariant semantics by comparing them
 717 with DINOv2 features extracted from the corresponding clean image \mathbf{x}_0 . Higher CKNNA scores
 718 indicate more semantically meaningful and noise-robust representations that align more closely with
 719 the features of the visual encoder. We follow the definition and implementation provided in the
 720 original work (Huh et al., 2024). Specifically, we compute the CKNNA score between the output
 721 of the encoder f_θ on noisy inputs \mathbf{x}_t and the output of DINOv2 on clean inputs \mathbf{x}_0 . We randomly
 722 sample 10K images from the ImageNet-1K validation set and report results with $k = 10$.

723 Finally, in Fig. 3c, we report FID values computed with 10K generated images. **Consistent with**
 724 **previous findings (Yu et al., 2024), we observe a strong negative correlation between the CKNNA**
 725 **values of intermediate diffusion features and FID scores. This suggests that higher alignment between**
 726 **diffusion features and high-quality visual representations leads to better generation quality.**

727 **Roles of dense-shallow and sparse-deep features (Fig. 4).** In Fig. 4, we analyze the contribution
 728 of each path in SPRINT. To generate samples using only a single path, we replace the feature
 729 representation of one path with that of the other. In other words, we duplicate the features from one
 730 path and concatenate the original and duplicated features before feeding them into the decoder.

732 **PCA visualization of diffusion features (Fig. 6).** In Fig. 6, we perform a principal component
 733 analysis (PCA) of the intermediate features to better understand what the model has learned. PCA
 734 identifies the principal axes that capture the greatest variance in the feature space and is widely
 735 used to analyze representations learned by neural networks (Oquab et al., 2023). We compute
 736 PCA across patch embeddings and visualize the first three principal components as RGB channels.
 737 Specifically, we examine the outputs of the encoder f_θ and the middle blocks g_θ at different timesteps
 738 to observe how the feature representations evolve throughout the diffusion process. Additional PCA
 739 visualizations are provided in Fig. 34.

741

B SPRINT WITH DIFFERENT DIFFUSION TRANSFORMERS

743 We provide details of the different diffusion transformers used in the main paper and describe how
 744 SPRINT is implemented on top of them.

746 **SiT (Ma et al., 2024).** We closely follow the architecture of SiT. The SiT model is structurally
 747 analogous to a Vision Transformer (ViT) (Dosovitskiy et al., 2020), consisting of a sequence of
 748 identical transformer blocks that process a patchified 1D token sequence. SiT adapts this for the
 749 diffusion task by incorporating timestep and class conditioning, which is injected into each block
 750 via AdaIN-zero layers. Because the architecture is a simple, homogeneous stack of blocks, it is
 751 straightforward to decouple it into our encoder, middle, and decoder blocks when applying SPRINT.

752 **REPA (Yu et al., 2024).** Representation Alignment (REPA) regularizes a DiT by aligning hidden
 753 states with clean image features from a pre-trained DINOv2 model. The architecture largely follows
 754 SiT, with the key modification being a projection layer inserted at the 8th transformer block to perform
 755 the alignment. To integrate SPRINT with REPA, we place this projection layer at the corresponding

756
757 **Table 9: Hyperparameters used for SPRINT.**
758
759

	SiT-B+SPRINT (Fig. 3, Tab. 5-8)	SiT-XL+SPRINT (Tab. 1, 3)	SiT-XL _{REPA} +SPRINT (Tab. 2, 3)	SiT-XL+SPRINT (Tab. 10)	U-ViT-XL+SPRINT (Tab. 2)
Architecture					
Target latent res.	32×32	32×32	32×32	64×64	32×32
Patch size	2	2	2	2	2
Total Num. Layers	12	28	28	28	28
Num. f_θ Layers	2	2	2	2	2
Num. g_θ Layers	8	24	24	24	24
Num. h_θ Layers	2	2	2	2	2
Hidden dims	384	1152	1152	1152	1152
Num. heads	6	16	16	16	16
Pretraining config.					
Optimizer	AdamW	AdamW	AdamW	AdamW	AdamW
Learning rate	0.0001	0.0001	0.0001	0.0001	0.0001
Batch size	256	256	256	256	256
Visual Encoder	–	–	DINOv2-B ($\lambda = 0.5$)	–	–
Drop ratio r	75%	75%	75%	75%	75%
Finetuning config.					
Training iterations	–	100K	100K	200K	100K
Warmup iterations	–	5K	5K	5K	5K
Optimizer	–	AdamW	AdamW	AdamW	AdamW
Learning rate	–	0.0002	0.0002	0.00015	0.0002
Batch size	–	512/1024	512/1024	1024	512
Drop ratio r	–	0%	0%	0%	0%
Evaluation config.					
Sampler	ODE	ODE/SDE	ODE/SDE	SDE	ODE
Sampling steps	50	50/250	50/250	250	50

778
779 location within our sparse middle block, g_θ . A key consideration is that the hidden states in g_θ
780 operate on a sparse token set of length N' , while the target DINOv2 features have a full sequence
781 length of N . To resolve this, we simply apply the same token-dropping mask to the DINOv2 feature
782 sequence, ensuring a one-to-one correspondence for the alignment loss. Since DINOv2 also uses a
783 standard transformer architecture with positional encodings, aligning the corresponding tokens is
784 straightforward.

785
786 **U-ViT (Bao et al., 2023).** U-ViT extends the Vision Transformer with a U-Net (Ho et al., 2020)-
787 style architecture. Similar to U-Net, it stacks transformer blocks with long skip-connections between
788 encoder and decoder stages, directly passing features from encoder to decoder. To apply SPRINT, we
789 first conceptually decompose the U-ViT into our standard f_θ , g_θ , and h_θ sections while preserving
790 all original skip-connections. We then introduce our dense residual path between f_θ and h_θ and
791 apply token dropping to the middle section, g_θ . The U-Net skip-connections remain compatible
792 with this design. The long-range skips between the encoder and decoder are unaffected. The shorter
793 skip-connections within the sparse middle section naturally operate on the reduced set of tokens. This
794 allows SPRINT to be integrated cleanly without disrupting the U-ViT’s core component.

795 C IMPLEMENTATION DETAILS AND HYPERPARAMETERS

796 C.1 TRAINING DETAILS

797
798 We follow the model configuration of the original SiT implementation (Ma et al., 2024), with the only
799 modification being a single linear projection layer for sparse–dense residual fusion. This adds only a
800 marginal number of parameters, approximately 0.3% of the original model size. We use pre-computed
801 latent vectors from raw images via Stable Diffusion (Rombach et al., 2022) and Flux (Labs, 2024b)
802 VAEs, and, following common practice, do not apply any data augmentation. For pretraining, we train
803 SPRINT with a batch size of 256, a learning rate of 1e-4, a fixed drop ratio of 75%, and an EMA decay
804 rate of 0.9999. After pre-training, we switch the middle blocks to operate on the full token set for a
805 short fine-tuning stage for 100K iterations. We increase the batch size and the learning rate, following
806 standard practice (Zheng et al., 2024; Krause et al., 2025). We found that applying a linear learning
807 rate warm-up from 2e-6 to 2e-4 over the first 5K iterations stabilizes the training. During the warm-up
808 stage, we use an EMA decay rate of 0.999, which is restored to 0.9999 afterward. For both training
809 phases, we introduce a path-drop learning strategy to maximize the effectiveness of our path-drop

810 **Algorithm 1** SPRINT Pre-training

811 **Require:** Input \mathbf{x}_0 , Drop ratio r , Path-drop prob p , encoder f_θ , middle blocks g_θ , decoder h_θ ,
 812 condition \mathbf{c}

813 1: **while** not converged **do**

814 2: Sample $t \sim [0, 1]$ and $\epsilon \sim \mathcal{N}(0, I)$

815 3: $\mathbf{x}_t \leftarrow (1 - t) \mathbf{x}_0 + t \epsilon$ $\triangleright \mathbf{x}_t \in \mathbb{R}^{B \times N \times C}$

816 4: $\mathbf{f}_t \leftarrow f_\theta(\mathbf{x}_t, \mathbf{c})$ $\triangleright \mathbf{f}_t \in \mathbb{R}^{B \times N \times C}$

817 5: $\mathbf{f}_t^{drop} \leftarrow \text{Drop}(\mathbf{f}_t, r)$ $\triangleright \mathbf{f}_t^{drop} \in \mathbb{R}^{B \times (1-r)N \times C}$

818 6: $\mathbf{g}_t^{drop} \leftarrow g_\theta(\mathbf{f}_t^{drop}, \mathbf{c})$ $\triangleright \mathbf{g}_t^{drop} \in \mathbb{R}^{B \times (1-r)N \times C}$

819 7: $\mathbf{g}_t^{pad} \leftarrow \text{PadWithMask}(\mathbf{g}_t^{drop})$ $\triangleright \mathbf{g}_t^{pad} \in \mathbb{R}^{B \times N \times C}$

820 8: $\mathbf{g}_t^{pad} \leftarrow [\text{MASK}]$ with probability p \triangleright Path-drop learning

821 9: $\mathbf{h}_t \leftarrow \text{Fusion}(\mathbf{f}_t, \mathbf{g}_t^{pad})$ \triangleright Sparse–dense residual fusion

822 10: $\hat{\mathbf{v}}_t \leftarrow h_\theta(\mathbf{h}_t, \mathbf{c})$ $\triangleright \hat{\mathbf{v}}_t \in \mathbb{R}^{B \times N \times C}$

823 11: $\mathcal{L}_{vel} \leftarrow \|\hat{\mathbf{v}}_t - \mathbf{v}_t\|^2$

824 12: Update θ using $\nabla_\theta \mathcal{L}_{vel}$

825 13: **end while**

826 14: **return** $f_\theta, g_\theta, h_\theta$

828 **Algorithm 2** SPRINT Fine-tuning

829 **Require:** Input \mathbf{x}_0 , Path-drop prob p , encoder f_θ , middle blocks g_θ , decoder h_θ , condition \mathbf{c}

830 1: **while** not converged **do**

831 2: Sample $t \sim [0, 1]$ and $\epsilon \sim \mathcal{N}(0, I)$

832 3: $\mathbf{x}_t \leftarrow (1 - t) \mathbf{x}_0 + t \epsilon$ $\triangleright \mathbf{x}_t \in \mathbb{R}^{B \times N \times C}$

833 4: $\mathbf{f}_t \leftarrow f_\theta(\mathbf{x}_t, \mathbf{c})$ $\triangleright \mathbf{f}_t \in \mathbb{R}^{B \times N \times C}$

834 5: $\mathbf{g}_t \leftarrow g_\theta(\mathbf{f}_t, \mathbf{c})$ $\triangleright \mathbf{g}_t \in \mathbb{R}^{B \times N \times C}$

835 6: $\mathbf{g}_t \leftarrow [\text{MASK}]$ with probability p \triangleright Path-drop learning

836 7: $\mathbf{h}_t \leftarrow \text{Fusion}(\mathbf{f}_t, \mathbf{g}_t)$ \triangleright Sparse–dense residual fusion

837 8: $\hat{\mathbf{v}}_t \leftarrow h_\theta(\mathbf{h}_t, \mathbf{c})$ $\triangleright \hat{\mathbf{v}}_t \in \mathbb{R}^{B \times N \times C}$

838 9: $\mathcal{L}_{vel} \leftarrow \|\hat{\mathbf{v}}_t - \mathbf{v}_t\|^2$

839 10: Update θ using $\nabla_\theta \mathcal{L}_{vel}$

840 11: **end while**

841 12: **return** $f_\theta, g_\theta, h_\theta$

843
 844
 845 guidance, in addition to the standard class-condition dropping. Specifically, following the practice in
 846 CFG training, we randomly drop the features of the sparse–deep path with a probability of 10% and
 847 replace the dropped features with mask tokens. This random dropping is performed independently of
 848 the condition dropping in CFG. To accelerate training, we adopt mixed-precision (bf16) training and
 849 apply gradient norm clipping at 1.0 during both pretraining and finetuning. Detailed hyperparameters
 850 are summarized in Table 9. All experiments are conducted on 8 NVIDIA A100 80GB GPUs.

851
 852 C.2 EVALUATION DETAILS
 853

854
 855 **Metrics.** We evaluate generation performance using several standard metrics: FDD (Stein et al.,
 856 2023) (Fréchet Distance on DINOv2), FID (Heusel et al., 2017) (Fréchet Inception Distance),
 857 IS (Salimans et al., 2016) (Inception Score), and Precision/Recall (Kynkänniemi et al., 2019).
 858 Unless otherwise specified, we follow the evaluation protocol of (Dhariwal & Nichol, 2021) and
 859 report results using 50K generated samples.

860 FID is the most widely used metric, measuring the feature distance between the distributions of real
 861 and generated images. It relies on the Inception-V3 network and assumes both feature distributions
 862 follow multivariate Gaussian distributions. IS also uses the Inception-V3 network but instead evaluates
 863 the quality and diversity of generated images by computing the KL-divergence between the marginal
 label distribution and the conditional label distribution predicted from logits.

864 **Algorithm 3** SPRINT Inference

865 **Require:** encoder f_θ , middle blocks g_θ , decoder h_θ , condition \mathbf{c} , guidance scale w , sampling steps
 866 N , sampler \mathcal{S}
 867 1: $\mathbf{x}_1 \sim \mathcal{N}(0, \mathcal{I})$
 868 2: **for** $i = N$ **to** 1 **do**
 869 3: $t \leftarrow \frac{i}{N}$
 870 4: **if** Path-drop guidance **then**
 871 5: $v(\mathbf{x}_t, \emptyset) \leftarrow h_\theta(\text{Fusion}(M, f_\theta(\mathbf{x}_t, \mathbf{c})), \mathbf{c})$ ▷ Path-drop guidance
 872 6: **else**
 873 7: $v(\mathbf{x}_t, \emptyset) \leftarrow h_\theta(\text{Fusion}(g_\theta(f_\theta(\mathbf{x}_t, \mathbf{c}), \mathbf{c}), f_\theta(\mathbf{x}_t, \emptyset)), \emptyset)$ ▷ Classifier-free guidance
 874 8: **end if**
 875 9: $\tilde{v}(\mathbf{x}_t, \mathbf{c}) \leftarrow v(\mathbf{x}_t, \emptyset) + w \cdot (v(\mathbf{x}_t, \mathbf{c}) - v(\mathbf{x}_t, \emptyset))$
 876 10: $\mathbf{x}_{t-\frac{1}{N}} \leftarrow \mathcal{S}(\mathbf{x}_t, \tilde{v}(\mathbf{x}_t, \mathbf{c}))$
 877 11: **end for**
 878 12: **return** \mathbf{x}_0

880 FDD adopts the same formulation as FID but replaces Inception features with DINOv2 features,
 881 which provide stronger semantic alignment and robustness to noise. Notably, FDD has been shown to
 882 be more reliable for diffusion models (Stein et al., 2023; Karras et al., 2024b).

884 Finally, Precision measures the fraction of generated images that are realistic, while Recall measures
 885 the fraction of the training data manifold covered by generated samples.

886 **Guidance scale.** We use the following formulation for guidance sampling (Ho & Salimans, 2022):

$$888 \tilde{v}(\mathbf{x}_t, \mathbf{c}) = v(\mathbf{x}_t, \emptyset) + w \cdot (v(\mathbf{x}_t, \mathbf{c}) - v(\mathbf{x}_t, \emptyset)), \quad (5)$$

889 where w denotes the guidance scale. In standard Classifier-Free Guidance (CFG), the unconditional
 890 velocity $v(\mathbf{x}_t, \emptyset)$ is computed using the full model path with a null condition. In contrast, our
 891 Path-Drop Guidance (PDG) replaces the unconditional branch with a weaker network, as defined in
 892 Eq. 4.

893 For the results in Tables 1 and 2, we consistently use a CFG scale of 1.4 with the ODE sampler across
 894 all methods.

896 For Table 3, we adopt the SDE sampler (Ma et al., 2024) to compare baselines. Under this setting,
 897 we use a CFG scale of 1.35 to achieve the best FID and 2.0 to achieve the best FDD. For our PDG
 898 sampling, the optimal scales are 1.35 for FID and 1.9 for FDD.

899 For our model in Table 10, we use the scale of 1.35 and 1.8 for FID and FDD, respectively, for both
 900 CFG and PDG.

902 **D COMPUTATION ANALYSIS**

904 We use the SiT-XL/2 configuration for evaluating computational analysis below.

906 **FLOPs.** To estimate the total training FLOPs, we measure the forward-pass FLOPs over 100
 907 iterations with a batch size of 256, average the results, and multiply by the total number of training
 908 iterations. For inference FLOPs, we sum the forward-pass FLOPs across all sampling timesteps using
 909 a batch size of 32 and report the average over both timesteps and batch size. This procedure provides
 910 a consistent and reproducible measure of computational cost across methods. Note that we report
 911 floating-point operations (FLOPs), not multiply–accumulate operations (MACs), where one MAC
 912 corresponds to approximately two FLOPs.

914 **Training speed.** Here, we compare the actual run-time performance of each method on Stable
 915 Diffusion VAE latents. For all token-dropping methods, we use a fixed drop rate of 75%. At
 916 the ImageNet resolution of 256², SPRINT achieves a pretraining speed of **5.2** iters/sec, which
 917 is more than **2× faster** than the SiT baseline (2.5 iters/sec) and clearly outperforms other token-
 918 dropping baselines, including MaskDiT (4.57 iters/sec), MicroDiT (3.9 iters/sec), and Tread (4.7

918
 919 **Table 10: Comprehensive performance comparison on ImageNet** 512×512 class-conditioned
 920 generation with classifier-free guidance. \downarrow / \uparrow indicate whether lower or higher values are better,
 921 respectively. All metrics are evaluated with 250 sampling steps using the SDE sampler. Training and
 922 inference TFLOPs are measured with the DeepSpeed library.

923 Method	924 Epochs	925 #Params.	926 Training 927 TFLOPs \downarrow ($\times 10^6$)	928 Inference 929 TFLOPs \downarrow	930 FDD \downarrow	931 FID \downarrow	932 Pre. \uparrow	933 Rec. \uparrow
925 ADM (Dhariwal & Nichol, 2021)	926 400	927 –	928 –	929 –	930 –	931 2.85	932 0.84	933 0.53
925 Simple diffusion (U-Net)	926 800	927 –	928 –	929 –	930 –	931 4.28	932 –	933 –
925 Simple diffusion (U-ViT-L)	926 800	927 –	928 –	929 –	930 –	931 4.53	932 –	933 –
925 MaskDiT (Zheng et al., 2024)	926 800	927 730M	928 327.2	929 1.029	930 –	931 2.50	932 0.83	933 0.56
925 DiT-XL (Peebles & Xie, 2023)	926 600	927 675M	928 366.6	929 0.952	930 –	931 3.04	932 0.84	933 0.54
925 SiT-XL (Ma et al., 2024)	926 600	927 675M	928 366.6	929 0.952	930 –	931 2.62	932 0.84	933 0.57
930 SiT-XL								
931 + SPRINT								
931 + SPRINT_{PDG}								
932								

934 iters/sec). At the higher ImageNet resolution of 512^2 , SPRINT maintains its advantage, achieving **2.01**
 935 iters/sec—over **$2.5\times$ faster** than the SiT baseline (0.79 iters/sec)—and again surpassing MaskDiT
 936 (1.77 iters/sec), MicroDiT (1.54 iters/sec), and Tread (1.79 iters/sec). This acceleration results in
 937 substantial reductions in wall-clock training time and GPU consumption, making large-scale diffusion
 938 model training significantly more practical and resource-efficient.

939
 940 **VRAM memory consumption.** In addition to reducing computational cost, SPRINT significantly
 941 lowers GPU memory requirements during training. For example, when training with a batch size
 942 of 32 and image resolution 256^2 on a single GPU, SPRINT requires only 19.6 GB of memory,
 943 compared to 29.6 GB for the baseline SiT-XL/2 model. At resolution 512^2 , our SPRINT requires
 944 37.9 GB, whereas the baseline SiT-XL/2 model requires 77.7 GB. This represents a **33.8% reduction**
 945 in memory usage at 256^2 and a **51.2% reduction** at 512^2 . Such efficiency enables training with
 946 larger batch sizes or higher resolutions on the same hardware, making our method more accessible
 947 for researchers with limited GPU resources. Importantly, this reduction comes without sacrificing
 948 performance, underscoring the practicality of SPRINT in resource-constrained environments.

949 E BASELINES

950 E.1 BASELINE DETAILS ON TABLE 1

951 For a fair system-level comparison in Tab. 1, we apply the *same pretraining and finetuning strategies*,
 952 along with identical transformer block configurations, **a fixed drop ratio of 75%**, and **consistent**
 953 **evaluation hyperparameters**, across all baselines.

954
 955 **Progressive training.** We adopt the same network architecture for progressive training. The model
 956 is first pretrained on 128×128 images and then finetuned on 256×256 images, with positional
 957 embeddings resized using bilinear interpolation during the resolution transition. This approach is
 958 slightly more efficient than SPRINT in terms of computational cost per iteration, achieving 25.8
 959 vs. 31.5 GFLOPs ($\times 10^9$) at 1M training iterations. However, despite the efficiency advantage,
 960 progressive training lags behind SPRINT in performance and even fails to match the baseline SiT
 961 results, underscoring its limited effectiveness.

962
 963 **MicroDiT (Sehwag et al., 2025).** MicroDiT introduces deferred masking, where token dropping is
 964 applied only after several additional patch-mixing blocks. These modules allow local patch tokens
 965 to fuse information, enriching their semantic content. Following the original protocol, we modify
 966 the SiT-XL/2 model by inserting patch-mixing modules composed of six transformer blocks. As
 967 shown in Tab. 1, this modification substantially increases computational cost and the number of
 968 parameters. Nevertheless, despite the additional overhead, MicroDiT underperforms relative to
 969 SPRINT, highlighting that the deferred masking strategy and additional compute does not translate
 970 into superior efficiency or accuracy.

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
Table 11: **Performance of SiT-XL/2 and SPRINT across NFEs.** Results are reported at 1M training iterations using the ODE sampler with 50K generated samples.

Method	NFE	FDD ↓	FID ↓	IS ↑	Pre. ↑	Rec. ↑
SiT-XL/2	200	132.3	2.18	249.9	0.81	0.59
+ SPRINT (Ours)	200	120.4	2.08	272.2	0.81	0.60
Gain Δ		+11.9	+0.1	+22.3		
SiT-XL/2	150	133.1	2.19	249.6	0.81	0.59
+ SPRINT (Ours)	150	121.1	2.09	271.5	0.81	0.59
Gain Δ		+12.0	+0.1	+21.9		
SiT-XL/2	100	134.7	2.22	248.4	0.81	0.58
+ SPRINT (Ours)	100	122.2	2.10	271.0	0.81	0.59
Gain Δ		+12.4	+0.12	+22.6		
SiT-XL/2	50	140.6	2.34	244.0	0.80	0.58
+ SPRINT (Ours)	50	126.5	2.19	267.7	0.81	0.59
Gain Δ		+14.1	+0.15	+23.7		
SiT-XL/2	25	156.1	2.91	234.4	0.80	0.57
+ SPRINT (Ours)	25	138.2	2.59	256.3	0.80	0.58
Gain Δ		+17.9	+0.32	+21.9		
SiT-XL/2	10	222.4	7.37	187.3	0.74	0.54
+ SPRINT (Ours)	10	191.7	6.29	211.3	0.74	0.54
Gain Δ		+30.7	+1.08	+24.0		

Tread (Krause et al., 2025). Tread introduces a token-routing strategy in which randomly dropped tokens at early layers are routed directly to deeper layers. While this resembles SPRINT in that tokens bypass the middle layers, the two approaches differ fundamentally. In Tread, only the dropped tokens are bypassed, forcing the middle block to encode local noise information in order to estimate velocity. In contrast, SPRINT employs a full dense residual path that delivers complete local noise information to the decoder, freeing the middle block to focus on modeling global contextual information. This design choice makes SPRINT *highly effective under aggressive dropping ratios* (75%), whereas Tread fails under the same setting. We follow the implementation details provided in the original Tread paper.

E.2 MORE DISCUSSION ON OTHER BASELINES

MaskDiT (Zheng et al., 2024). MaskDiT introduces an additional reconstruction task for masked tokens alongside the diffusion objective, encouraging the model to recover missing information and thereby improve contextual understanding. While this approach provides some efficiency gains, it requires an extra decoder module, increasing the model size from 675M to 730M and adding computational overhead. Moreover, its effectiveness is limited to moderate dropping ratios (e.g., 50%). As shown in Tab. 3, these limitations restrict its overall efficiency compared to our framework. Specifically, MaskDiT requires 1600 training epochs to reach 65.4 FDD and 2.28 FID, whereas SPRINT surpasses this in just 200 epochs with 61.8 FDD and 2.01 FID. This underscores the superior effectiveness and efficiency of SPRINT over MaskDiT.

MDT (Gao et al., 2023). The Masked Diffusion Transformer (MDT) also aims to improve the contextual understanding of diffusion models through token dropping. They designed masked diffusion transformer with encoder-decoder split of the diffusion transformer, where the encoder processes masked tokens and forwards them to the decoder along with remaining tokens through additional side-interpolator model. It adds additional long shorcut connections between encoder blocks along with long full token input to all decoder blocks. The added complexity in the training and architectural changes is aimed for better generative performance. Similar to MaskDiT, this work also operates only with moderate token dropping ratios (e.g., [30%, 50%]). MDT does not work well with high token dropping ratio such as 75%.

1026
1027

F ADDITIONAL QUANTITATIVE RESULTS

1028
1029

F.1 IMAGENET 512x512 EXPERIMENT

1030
1031
1032

In the main text, we have already demonstrate that SPRINT outperforms many existing training methods and state-of-the-art models at 256^2 class conditional image generation. In this experiment, we train our models to generation images at 512^2 resolution.

1033
1034
1035
1036
1037
1038
1039
1040
1041

Tab. 10 compares our method with strong baselines on ImageNet-1K class-conditional generation at 512^2 . We pre-train SPRINT for 1.8M iterations and finetune for 200K iterations (refer to Table 9). SPRINT achieves better generation quality while using substantially fewer training TFLOPs ($\times 10^6$): only 184.8 at 400 epochs, versus 366.6 for SiT-XL at 600 epochs. This demonstrates much faster convergence, reaching better FID (2.23 vs. 2.62) with nearly $2 \times$ lower training cost. At inference, Path-Drop Guidance provides further benefits, nearly halving inference TFLOPs (0.471 vs. 0.952) while improving both FID and FDD. Overall, SPRINT consistently demonstrates significant efficiency compared to the baselines at 512^2 , by combining lower training and inference costs. Refer to Fig. 33 for qualitative results.

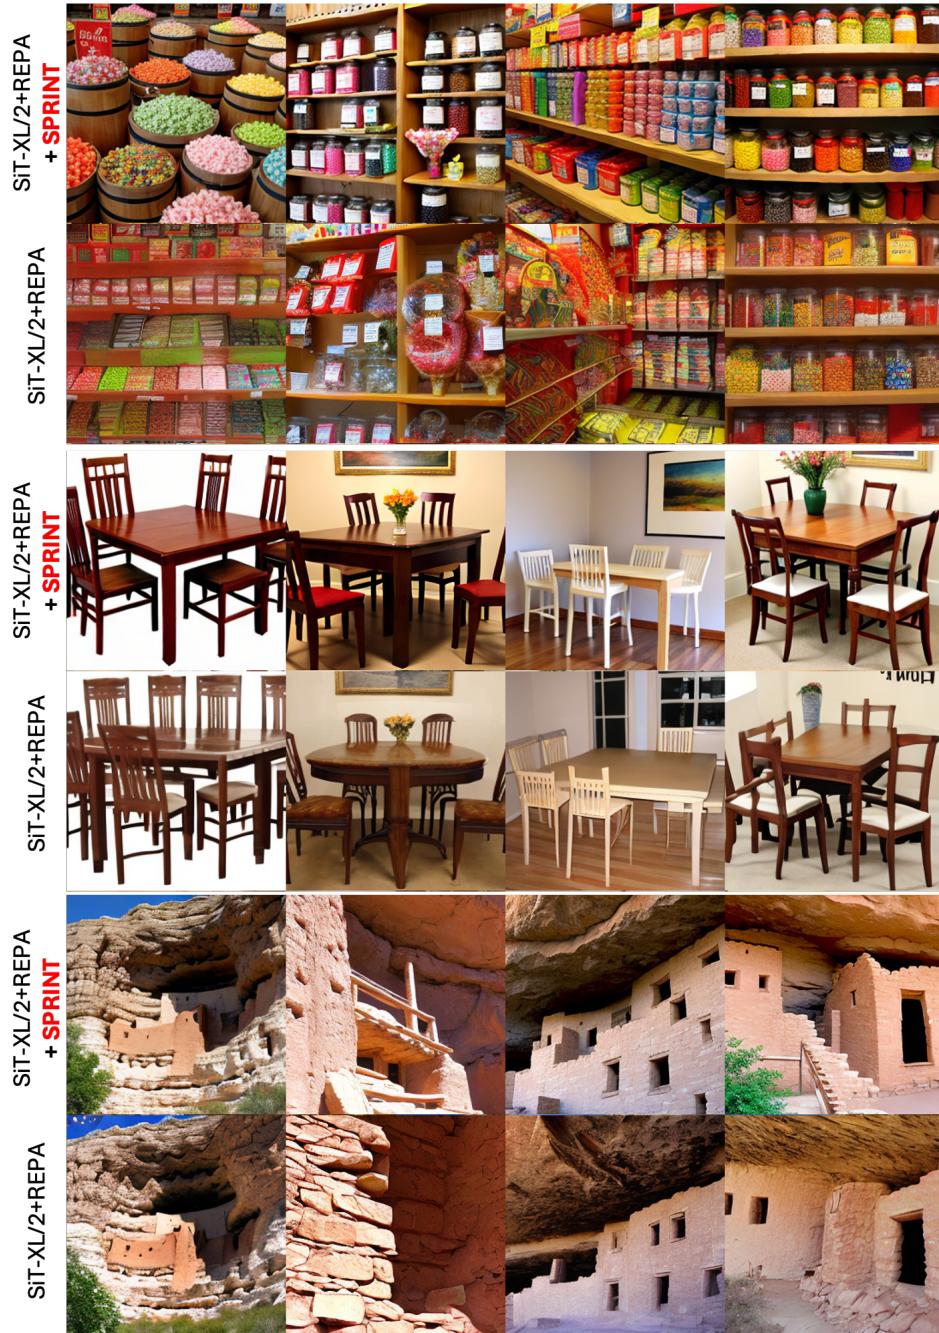
1042
1043

F.2 PERFORMANCE WITH FEW-STEP GENERATION

1044
1045
1046
1047
1048
1049
1050
1051
1052

Tab. 11 compares SiT-XL/2 and SiT-XL/2 + SPRINT across lower inference steps (NFEs), an essential setting for achieving efficient and practical image generation. In real-world scenarios, reducing the number of function evaluations (NFEs) directly translates to faster sampling and lower inference cost, often at the expense of generation quality. While both models perform similarly at large NFEs (200), SPRINT consistently outperforms the baseline as the number of steps decreases. At 50 steps, SPRINT improves FID from 2.34 to 2.19 and IS from 244.0 to 267.7, and at only 10 steps it achieves a much larger gain, reducing FID from 7.37 to 6.29 and improving IS from 187.3 to 211.3. These results highlight that SPRINT is more competitive under low-step inference. This demonstrates the *strong representational power* of fused dense–shallow and sparse–deep features.

1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

1080 G ADDITIONAL QUALITATIVE RESULTS
10811082 G.1 VISUAL COMPARISON ON IMAGENET 256 × 256
1083

1126 Figure 10: **SPRINT improves visual quality over baseline with only 57% of inference FLOPs**
1127 **(additional examples).** We present samples from two SiT-XL/2 + REPA models after 1M iterations,
1128 where SPRINT is applied to one of the models. For our approach, we further incorporate the proposed
1129 Path-Drop Guidance (PDG), yielding higher visual quality compared to the REPA.

1130
1131
1132
1133

1134
1135 G.2 UNSELECTED GENERATED RESULTS BY SPRINT ON IMAGENET 256×256 1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152 Figure 11: Unselected generation results of SiT-XL/2 + SPRINT_{CFG}. We use classifier-free
1153 guidance with $w = 4.0$. Class label = “park bench” (706)1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173 Figure 12: Unselected generation results of SiT-XL/2 + SPRINT_{PDG}. We use our path-drop
1174 guidance with $w = 4.0$. Class label = “park bench” (706)1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

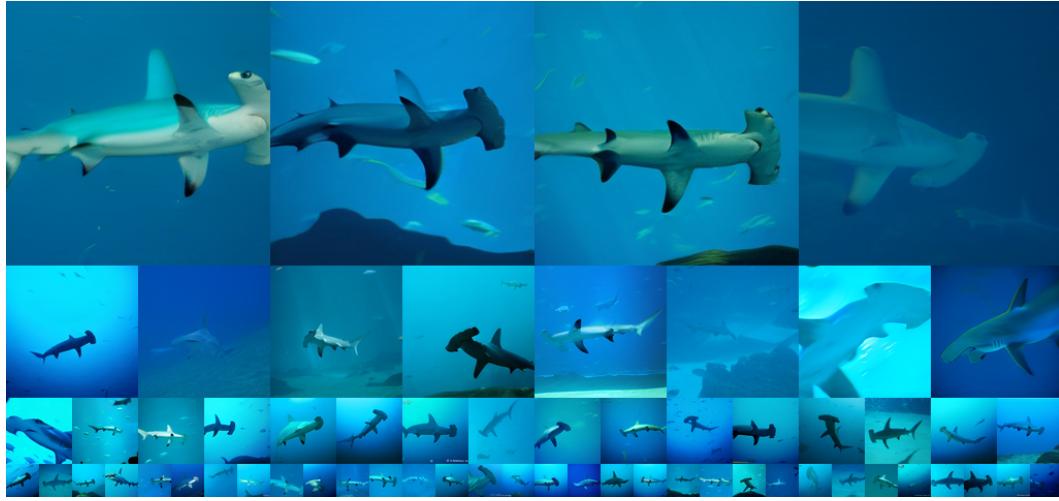


Figure 13: **Unselected generation results of SiT-XL/2 + SPRINT_{CFG}.** We use classifier-free guidance with $w = 4.0$. Class label = “hammerhead, hammerhead shark” (4)



Figure 14: **Unselected generation results of SiT-XL/2 + SPRINT_{PDG}.** We use our path-drop guidance with $w = 4.0$. Class label = “hammerhead, hammerhead shark” (4)

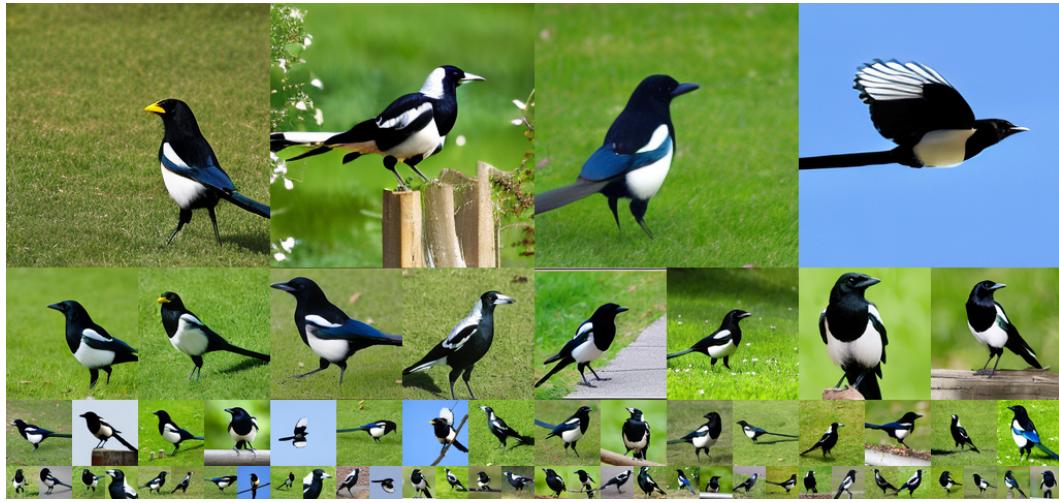


Figure 15: **Unselected generation results of SiT-XL/2 + SPRINT_{CFG}.** We use classifier-free guidance with $w = 4.0$. Class label = “magpie” (18)



Figure 16: **Unselected generation results of SiT-XL/2 + SPRINT_{PDG}.** We use our path-drop guidance with $w = 4.0$. Class label = “magpie” (18)

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



Figure 17: **Unselected generation results of SiT-XL/2 + SPRINT_{CFG}.** We use classifier-free guidance with $w = 4.0$. Class label = “bullfrog, *Rana catesbeiana*” (30)



Figure 18: **Unselected generation results of SiT-XL/2 + SPRINT_{PDG}.** We use our path-drop guidance with $w = 4.0$. Class label = “bullfrog, *Rana catesbeiana*” (30)



Figure 19: **Unselected generation results of SiT-XL/2 + SPRINT_{CFG}.** We use classifier-free guidance with $w = 4.0$. Class label = “tusker” (101)



Figure 20: **Unselected generation results of SiT-XL/2 + SPRINT_{PDG}.** We use our path-drop guidance with $w = 4.0$. Class label = “tusker” (101)

1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

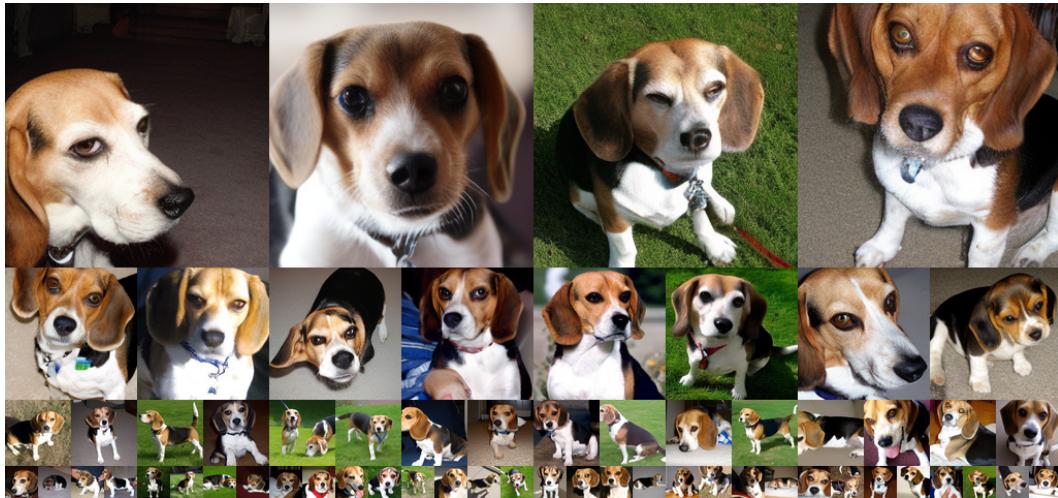


Figure 21: **Unselected generation results of SiT-XL/2 + SPRINT_{CFG}.** We use classifier-free guidance with $w = 4.0$. Class label = “beagle” (162)



Figure 22: **Unselected generation results of SiT-XL/2 + SPRINT_{PDG}.** We use our path-drop guidance with $w = 4.0$. Class label = “beagle” (162)



Figure 23: **Unselected generation results of SiT-XL/2 + SPRINT_{CFG}.** We use classifier-free guidance with $w = 4.0$. Class label = “coffeepot” (505)



Figure 24: **Unselected generation results of SiT-XL/2 + SPRINT_{PDG}.** We use our path-drop guidance with $w = 4.0$. Class label = “coffeepot” (505)

1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511



Figure 25: **Unselected generation results of SiT-XL/2 + SPRINT_{CFG}.** We use classifier-free guidance with $w = 4.0$. Class label = “computer keyboard, keypad” (508)



Figure 26: **Unselected generation results of SiT-XL/2 + SPRINT_{PDG}.** We use our path-drop guidance with $w = 4.0$. Class label = “computer keyboard, keypad” (508)



1582 **Figure 27: Unselected generation results of SiT-XL2 + SPRINT_{CFG}.** We use classifier-free
1583 guidance with $w = 4.0$. Class label = “convertible” (511)



1604 **Figure 28: Unselected generation results of SiT-XL2 + SPRINT_{PDG}.** We use our path-drop
1605 guidance with $w = 4.0$. Class label = “convertible” (511)

1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619



Figure 29: **Unselected generation results of SiT-XL/2 + SPRINT_{CFG}.** We use classifier-free guidance with $w = 4.0$. Class label = “cornet, horn, trumpet, trump” (513)



Figure 30: **Unselected generation results of SiT-XL/2 + SPRINT_{PDG}.** We use our path-drop guidance with $w = 4.0$. Class label = “cornet, horn, trumpet, trump” (513)

1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673



Figure 31: **Unselected generation results of SiT-XL/2 + SPRINT_{CFG}.** We use classifier-free guidance with $w = 4.0$. Class label = “cowboy hat, ten-gallon hat” (515)



Figure 32: **Unselected generation results of SiT-XL/2 + SPRINT_{PDG}.** We use our path-drop guidance with $w = 4.0$. Class label = “cowboy hat, ten-gallon hat” (515)

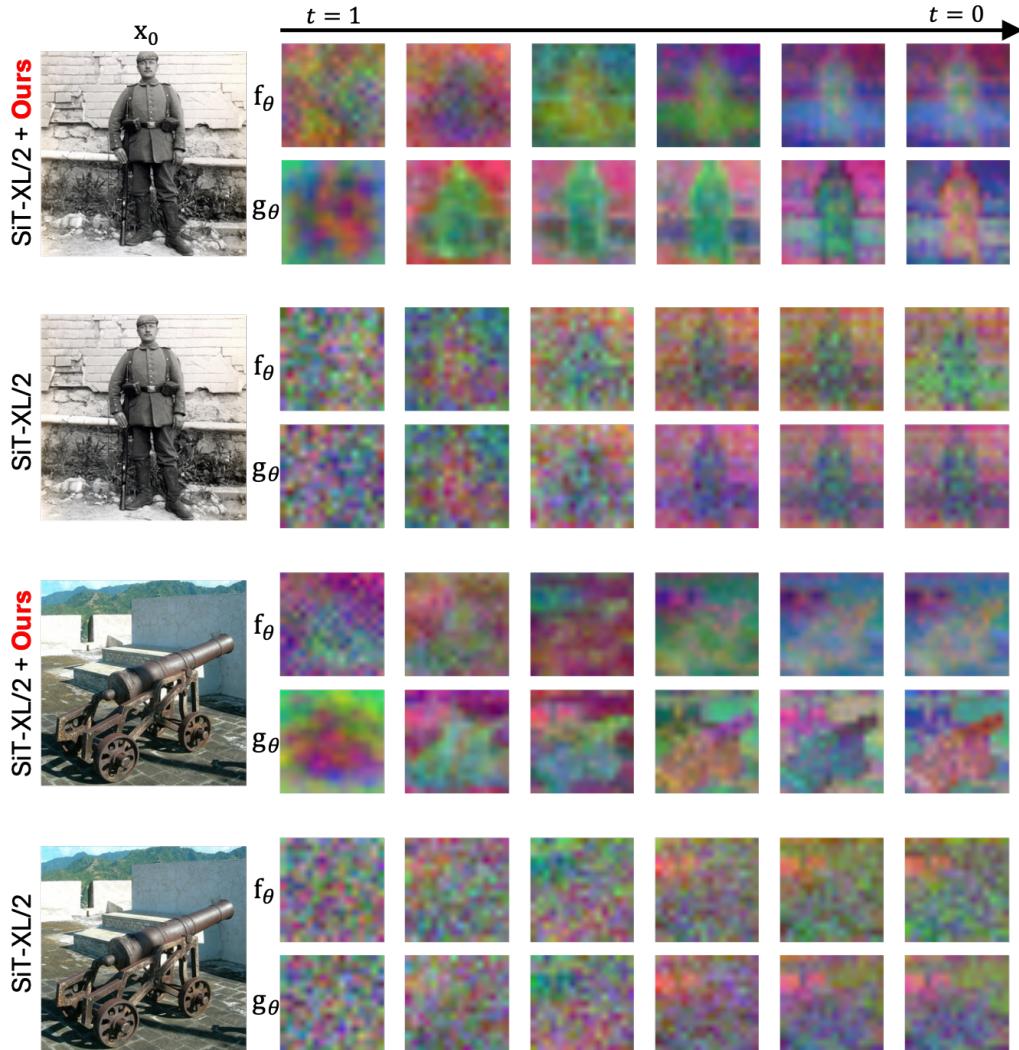
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

1728
1729 G.3 GENERATED RESULTS BY SPRINT ON IMAGENET 512×512 1764 Figure 33: **Generation results of SiT-XL/2 + SPRINT_{PDG}.** We use our path-drop guidance with w
1765 = 3.0.

1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781

1782 G.4 ADDITIONAL FEATURE PCA VISUALIZATION
1783

1784 In the main text (Figure 6), we analyzed PCA visualizations of features from f_θ and g_θ , showing that
1785 SPRINT learns more noise-invariant and semantically vivid representations than the SiT baseline
1786 across diffusion timesteps. Figure 34 presents additional examples of these dense–shallow and
1787 sparse–deep features learned by SPRINT, contrasted with those from a standard SiT-XL/2 model
1788 trained with full tokens.



1823 **Figure 34: SPRINT improves feature semantics (additional examples).** We visualize PCA features
1824 of f_θ and g_θ from two SiT-XL/2 models at 400K iterations. The top rows show the model trained
1825 with SPRINT, while the bottom rows show the baseline. Compared to the baseline, features from
1826 SPRINT exhibit clearer semantic structure across both images.

1827
1828
1829
1830
1831
1832
1833
1834
1835

1836 H LIMITATION AND FUTURE WORK

1837
1838 Our study is limited by the available computational resources, which prevented us from conducting
1839 experiments on large-scale text-to-image or video diffusion models. Exploring the scalability of
1840 SPRINT in such settings remains an important direction. In particular, the quadratic complexity
1841 of transformers becomes increasingly prohibitive as model size and input resolution grow. Since
1842 SPRINT is specifically designed to reduce redundant computation in deeper layers, we expect it to
1843 be especially beneficial for large-scale architectures where efficiency bottlenecks are most severe.
1844 Thus, extending SPRINT to other modalities such as video, 3D, or multi-modal generative models
1845 is an exciting direction. These domains pose even greater computational and memory challenges,
1846 particularly in video, where the temporal dimension compounds complexity, making our sparse–dense
1847 residual fusion especially relevant for future research.

1848 Another promising avenue is the integration of SPRINT with recent advances in efficient attention
1849 mechanisms and scalable training strategies. Such combinations could amplify the benefits of our
1850 approach, further reducing training and inference costs while maintaining or improving performance.

1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889