

Semantic Flow Regularization: Teaching LLMs to Generate Diverse Yet Coherent Responses

Kerui Peng, Feifei Li, Xingyu Fan, Wenhui Que*

WeChat, Tencent Inc., Beijing, China

{keruipeng, niyali, fanxfan, victorque}@tencent.com

Abstract

When large language models are fine-tuned to generate persona- or tone-conditioned responses, their output diversity is severely limited—a failure we term **Cross-Style Collapse**. We trace this collapse to the cross-entropy objective, which under shared representations tends to suppress diverse continuations. We propose **Semantic Flow Regularization (SFR)**, a lightweight auxiliary objective that supervises the backbone with continuous sentence-encoder embeddings of future segments via conditional flow matching. The stochastic flow source preserves multi-modality by construction; the flow-matching head is discarded at inference, adding **zero deployment cost**. On a large-scale industrial dialogue dataset (Qwen3-32B, 9 personas), SFR improves output diversity, style fidelity, and response quality over SFT. We further validate on the public LiveCodeBench-v5 (Qwen2.5-Coder-7B-Instruct), where SFR consistently improves pass@k, confirming generality beyond stylized dialogue. A controlled comparison on MBPP reveals Multi-Token Prediction to be a degenerate special case of SFR.¹

1 Introduction

LLM-based conversational agents are increasingly used in applications requiring both factual correctness and stylistic diversity. Our production deployments include personalized question answering for AI avatars on WeChat Official Accounts and personalized generation for WeChat Moments’ AI-assisted writing, where responses must reflect each user’s communication style while remaining grounded in relevant published content. In these settings, the same query can admit many valid, persona-dependent responses.

Standard supervised fine-tuning (SFT) optimizes a token-level cross-entropy (CE) loss. When valid

continuations are multi-modal, whether across different style conditions or across multiple valid solutions under the same condition, CE under finite-capacity shared representations tends to encourage shared high-probability continuations, suppressing output diversity even when the underlying target distribution is rich—a link between the likelihood objective and generic, diversity-poor responses established at least since Li et al. (2016). In the stylized-dialogue setting this manifests as what we term Cross-Style Collapse, echoing the text degeneration (Holtzman et al., 2020) and persona homogenization (Xiao et al., 2026) observed in prior work: swapping the persona on a fixed query often flips only a greeting token while the response body stays nearly unchanged. A related symptom appears in code generation, where pass@k gains plateau because the model collapses to a single template despite multiple correct solutions existing.

Several families of existing work address aspects of this gap, but none fully resolves it. *Decoding-time* methods such as nucleus sampling (Holtzman et al., 2020), contrastive decoding (Li et al., 2023), classifier-guided decoders (Keskar et al., 2019; Yang and Klein, 2021; Krause et al., 2021) reweight logits without changing the backbone’s internal representations, so the underlying mode-averaging persists. *Per-style fine-tuning* (Fan et al., 2025) trains a separate parameter set for each persona, achieving strong stylistic control but incurring storage and maintenance costs that grow linearly with the number of styles. *Training-time discrete extensions* such as unlikelihood (Welleck et al., 2020), SimCTG (Su et al., 2022), and multi-token prediction (Gloeckle et al., 2024) extend the supervision horizon yet remain in one-hot geometry and do not explicitly model the multi-modal distribution over styles.

We propose **Semantic Flow Regularization (SFR)**, a single auxiliary objective added during fine-tuning that addresses the problem from

*Corresponding author.

¹Code: <https://github.com/WeAgentAI/SFR>.

a different angle: teaching the backbone to *see* the future stylistic configuration of each response. At every response position t a lightweight flow-matching (FM) head v_ϕ learns to transport a random unit-sphere source z_0 to the sentence-encoder embedding $z_1^{(t)} = E(y_{t+1:t+k})$ of the upcoming suffix, conditioned on the backbone hidden state h_t . Because the source is stochastic, the learned vector field preserves the full conditional distribution rather than collapsing to its mean (Section A); the backbone is thus encouraged to encode style-discriminative information at every position. At inference time v_ϕ and E are simply deleted—the deployed model is bit-identical to a standard AR LM with no additional latency or memory.

Our contributions:

- We introduce SFR, a training-time auxiliary that injects continuous, segment-level future supervision into AR language models at zero inference cost, specifically targeting the mode-averaging problem in multi-style fine-tuning.
- We show theoretically that stochastic flow-matching supervision preserves conditional multi-modality where deterministic regression alternatives provably cannot.
- We validate SFR on a large-scale industrial stylized-dialogue dataset, showing significant improvements in both output diversity and style fidelity; since this dataset is proprietary, we further confirm the generality of SFR on the public LiveCodeBench-v5 benchmark with consistent pass@k gains. A controlled comparison on MBPP additionally connects SFR to Multi-Token Prediction as a degenerate special case.

2 Related Work

2.1 Stylized and persona-grounded generation

Style-controllable dialogue spans both benchmarks and methods. PersonaChat (Zhang et al., 2018; Dinan et al., 2020), LIGHT (Urbanek et al., 2019), RoleLLM (Wang et al., 2024) and CharacterGLM (Zhou et al., 2023) supply persona profiles and evaluate consistency. On the method side, per-style fine-tuning (Fan et al., 2025) trains a dedicated parameter set (or LoRA adapter) per persona, yielding strong control but linear storage cost in the number of styles. Decoding-time approaches steer logits without touching the backbone, such as CTRL (Keskar et al., 2019), FUDGE (Yang and Klein, 2021), GeDi (Krause et al., 2021),

PPLM (Dathathri et al., 2020), so the backbone’s representations remain mode-averaged and insensitive to the persona prompt. Closest in spirit are latent-variable dialogue models that also inject a stochastic signal to escape mode-averaging: VHRED (Serban et al., 2017) introduces a hierarchical Gaussian latent over the whole utterance, and the CVAE dialog model of Zhao et al. (2017) conditions generation on a sampled discourse-level latent; StyleFusion (Gao et al., 2019) targets the diversity–relevance trade-off in stylized dialogue most directly, fusing a style-specific and a content-specific latent space. All three inject stochasticity into a *discrete-token* decoder and are trained end-to-end as the primary generative model, so the stochastic component cannot be removed at inference without retraining; SFR instead couples the stochastic source to a *continuous, sentence-level* embedding target via flow matching, and drops the auxiliary head entirely after training, leaving a plain AR decoder identical to one trained with SFT. SFR thus takes a middle path relative to per-style parameters and decoding-time classifiers: one unified training objective that makes the backbone itself style-sensitive, requiring neither per-persona parameters nor decoding-time classifiers.

2.2 Training-time future supervision

Several methods augment CE with longer-horizon targets. Unlikelihood (Welleck et al., 2020), SimCTG (Su et al., 2022) and MixCE (Zhang et al., 2023) penalise degeneration at horizon 1. MTP (Gloeckle et al., 2024; DeepSeek-AI et al., 2025) predicts K future tokens but remains in discrete one-hot space. Continuous-target alternatives move into embedding space: representation alignment via MSE or cosine (LCM team et al., 2024) regresses to a fixed embedding, while JEPA (Assran et al., 2023) and V-JEPA (Bardes et al., 2024) predict latent representations of future inputs in vision. JEPA is architecturally the closest predecessor—it also uses a predictor network conditioned on context to forecast future embeddings. However, JEPA operates in the vision domain without an autoregressive generator, its predictor is deterministic (collapsing to the conditional mean; Proposition 3.1), and it is not designed as a drop-in auxiliary for zero-overhead AR deployment. SFR differs on all three counts: it is integrated with an AR LLM that directly generates text, it uses a stochastic FM source to preserve multi-modality, and the FM head is deleted at deployment with zero infer-




Figure 1: **Left:** Training and inference architecture of SFR. During training, the backbone F_θ produces hidden states h_t consumed by both the LM head W_{AR} (standard next-token loss $\mathcal{L}_{AR}$) and the auxiliary FM head v_ϕ (flow-matching loss $\mathcal{L}_{SFR}$, supervised by the frozen encoder target $z_1 = E(y_{t+1:t+k})$). At inference, v_ϕ and E are discarded; the deployed model is a standard AR LM. **Right:** Illustration of style-discriminative generation. The input prompt consists of a style description concatenated with the user query. Without SFR, the next-token-only LM head tends toward style-collapsed outputs. The FM head regularizes the backbone toward distinct regions on the target embedding manifold for each style, enabling the model to produce genuinely differentiated responses across personas.

ence overhead.

Concurrent work (Zhang et al., 2026) observes that token positions divide into *Locks* (one dominant continuation) and *Forks* (multiple valid continuations), but does not explicitly focus on locating which positions fall into which category. In Section 5.2 we show that the FM head trained by SFR yields an explicit per-position Lock/Fork diagnostic.

2.3 Flow and diffusion language models

Diffusion-LM (Li et al., 2022), SEDD (Lou et al., 2024), Plaid (Gulrajani and Hashimoto, 2023) and RFFlow (Liu, 2022) model continuous multi-modal sequence distributions but replace the AR decoder with an iterative denoiser, breaking compatibility with standard serving stacks. The key distinction is that these methods use flow/diffusion as the *generation* mechanism, whereas SFR uses it purely as a *training-time regularizer*: the flow head is discarded before deployment, leaving a plain AR LM with zero architectural or inference-cost overhead.

3 Method

3.1 Overall framework

The backbone F_θ and the LM head W_{AR} are unchanged. We attach a lightweight conditional vector-field head $v_\phi(z_\tau, \tau; h_t)$ that participates only in the auxiliary loss and is dropped at inference. Training minimizes

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{AR} + \lambda \cdot \mathcal{L}_{SFR}, \quad (1)$$

where $\mathcal{L}_{AR}$ is the mean next-token cross-entropy over all response positions and $\mathcal{L}_{SFR}$ is the mean flow-matching loss defined in (3) below. Figure 1 (left) summarizes the training-time data flow and the deployment-time graph, which is identical to a vanilla AR LM, while Figure 1 (right) illustrates how the FM head encourages style-discriminative backbone representations.

In principle, v_ϕ can be attached after any intermediate layer of the backbone: an earlier tap gives the FM head access to lower-level features but also requires a larger predictor to bridge the representational gap between that layer and the final semantic target z_1 . We attach v_ϕ after the last Transformer layer, so that the conditioning h_t is the same representation that the LM head (i.e. the final un-embedding projection followed by softmax) operates on. This choice maximizes the information available to v_ϕ while keeping the predictor small, and ensures that the FM gradient reshapes precisely the representation that the AR head consumes. Concretely, v_ϕ fuses its three inputs as follows: the scalar τ is first mapped, via a sinusoidal timestep embedding (as in standard diffusion time-conditioning), to a vector of the same dimension as z_τ ; this embedding, z_τ , and the conditioning hidden state h_t are concatenated along the feature dimension, passed through an input projection, and then through several residual blocks with Layer-Norm, before a zero-initialized output projection produces $\hat{v}$. Hidden width and depth scale with

the backbone (Section D). A simple architecture is intentional: the FM head’s role is to supply gradient signal to the backbone rather than to solve the flow task itself, so limited capacity prevents it from short-circuiting the auxiliary objective (cf. Section 3.4).

We validate the framework in the SFT setting; extending it to pretraining, LoRA, or RLHF is left to future work.

3.2 Auxiliary objective

For each response position t , the auxiliary target is the L2-normalized sentence-encoder embedding of the next k tokens,

$$z_1^{(t)} = E(y_{t+1}, \dots, y_{t+k}) \in \mathbb{S}^{d_z-1}, \quad (2)$$

where E is a frozen sentence encoder (BGE-Large (Xiao et al., 2024) in our experiments) and k is the suffix horizon. We supervise the FM head with conditional FM (CFM; Lipman et al., 2023; Liu et al., 2023b): given a random source $z_0 \sim \text{Unif}(\mathbb{S}^{d_z-1})$ (a standard Gaussian sample subsequently renormalized to the unit sphere, matching the space in which z_1 lives), uniform $\tau \sim \mathcal{U}[0, 1]$, and straight-line interpolation $z_\tau = (1-\tau)z_0 + \tau z_1^{(t)}$, the loss at position t is

$$\mathcal{L}_{\text{SFR}}^{(t)} = \mathbb{E}_{\tau, z_0} \|v_\phi(z_\tau, \tau; h_t) - (z_1^{(t)} - z_0)\|_2^2. \quad (3)$$

The full training loss $\mathcal{L}_{\text{SFR}}$ is the mean of (3) over all response positions.

Each h_t therefore receives a *continuous*, d_z -dimensional target rather than a one-hot symbol, and a *segment-level* target spanning the next k tokens rather than only y_{t+1} . The randomness of z_0 is also essential rather than cosmetic, formalized by the following two results (proofs in Section A.2):

Proposition 3.1 (Mode averaging). *For a conditional distribution $P(z_1 | c)$ with finite second moment, the unique minimizer of $\min_f \mathbb{E} \|f(c) - z_1\|_2^2$ is the conditional mean $\mathbb{E}[z_1 | c]$. If $P(z_1 | c)$ is multi-modal, this minimizer coincides with no mode and may lie in a low-density region between them.*

Proposition 3.2 (Multi-modal capacity of CFM). *Under the straight-line path $z_\tau = (1-\tau)z_0 + \tau z_1$ with stochastic $z_0 \sim p_0$, the minimizer of (3) is the marginal velocity $v^*(z, \tau, c) = \mathbb{E}[z_1 - z_0 | z_\tau = z, c]$, and the ODE $\dot{z}_\tau = v^*(z_\tau, \tau, c)$ transports p_0 to the full conditional target distribution $p_1(\cdot | c)$ at $\tau=1$, preserving all of its modes.*

Proposition 3.1 shows that replacing z_0 by a fixed point collapses (3) into deterministic regression whose optimum is the conditional mean; Proposition 3.2 shows that a stochastic source instead recovers the full conditional multi-modal structure of $P(z_1 | h_t)$. The motivation for choosing flow matching over MSE / cosine alignment or diffusion denoising is discussed in Section 6.1.

The exact per-position encoding in (2) costs $O(L)$ calls to E per training example. We reduce this to $O(L/s)$ by *sparse anchor encoding*: E is invoked only every s tokens, and intermediate positions reuse the nearest anchor’s target. The stride ablation in Section G shows that quality is roughly flat from $s=1$ to $s=16$ while training time drops by $\sim 24\%$; in practice s can be chosen jointly with the suffix length k .

3.3 Inference

At deployment, v_ϕ and E are removed and only $(F_\theta, W_{\text{AR}})$ is loaded. The resulting checkpoint has the same architecture and inference graph as a vanilla AR LM and can be served by any standard tool (vLLM (Kwon et al., 2023), SGLang (Zheng et al., 2024), LoRA merging, GPTQ (Frantar et al., 2022)/AWQ (Lin et al., 2024) quantization) without modification; at deployment SFR adds no extra parameters, no extra latency, and no extra decoding logic.

The argument for why this is sound is that v_ϕ has already done its work by the end of training: through backpropagation it has shaped θ so that h_t encodes future-aware, condition-discriminative information, and this property travels with θ regardless of whether v_ϕ is loaded at inference. Section 5.1 verifies this directly by showing that the matched-position hidden states $h_t^{q, s_1}, \dots, h_t^{q, s_m}$ become well-separated clusters under SFR, even though v_ϕ is no longer in the forward graph.

3.4 Training stabilisation

Naïvely adding $\lambda \mathcal{L}_{\text{SFR}}$ to $\mathcal{L}_{\text{AR}}$ exhibits two failure modes. The first is *head short-circuiting*: v_ϕ learns the suffix mapping faster than the backbone reorganizes, $\mathcal{L}_{\text{SFR}}$ collapses to near zero, and the auxiliary gradient stops flowing back into θ . The second is *early-step shock*: a randomly initialized v_ϕ produces large initial gradients that destabilize the pretrained backbone before any useful auxiliary signal has formed. We address both with three mechanisms that must act jointly—removing any one of them causes training to diverge (either the

backbone loss oscillates violently or the FM head ceases to influence the backbone at all).

First, the output projection of v_ϕ is zero-initialized, so that $\hat{v} \equiv 0$ at step 0 and $\mathcal{L}_{\text{SFR}}$ starts at the benign constant $\mathbb{E}\|z_1 - z_0\|_2^2$, preventing early-step shock. Second, a two-phase warmup blocks FM gradients from propagating into the backbone for the first T_h steps (only v_ϕ updates during this period), after which λ ramps linearly from 0 to its target value over the next T_ℓ steps. Third, an exponential moving average with decay μ is maintained on v_ϕ 's parameters, making the predictor a slow-moving target that the backbone can track without the predictor racing ahead. Implementation details and schedule values are given in Sections C and D; Section F provides the complete training algorithm and Section E a taxonomy of the design space.

4 Experiments

We evaluate SFR in three settings: (i) a large-scale industrial stylized-dialogue task that stress-tests output diversity under persona conditioning (§4.1), (ii) the public LiveCodeBench-v5 benchmark that confirms generality on code generation (§4.2), and (iii) a controlled comparison with MTP on MBPP (§4.3). Single-factor ablations on λ , k , s , and the three stabilizers are deferred to Section G.

4.1 Stylized dialogue

To the best of our knowledge, there exists no public dataset that simultaneously provides articles, user queries, and large-scale stylized replies from millions of authors. We therefore evaluate directly on proprietary data from a widely-used, real-world industrial official-accounts platform. We fine-tune **Qwen3-32B** (Yang et al., 2025) on $\sim 170\text{k}$ multi-persona dialogue examples covering nine stylistic personas (Section B.1 details the full construction pipeline); the held-out evaluation set contains ~ 150 queries, and for each method we generate one response per (query, persona) pair.

We compare against a **Base** reference (the frozen Qwen3-32B backbone with persona prompt, without any fine-tuning) and **SFT** (standard supervised fine-tuning on the same multi-persona corpus). Our method, **SFR**, is trained on top of the same SFT recipe with the auxiliary flow-matching objective of (3). For the LLM-judge evaluation we additionally include **DeepSeek-R1-prompt**, a strong prompt-only baseline that calls a frozen DeepSeek-R1 (Guo et al., 2025) with the same

	SFT	SimCTG	CD	SFR
CS-SB ₁	0.798 $\pm$.007	0.813 $\pm$.006	0.812 $\pm$.009	0.783$\pm$.006
CS-SB ₂	0.655 $\pm$.011	0.680 $\pm$.011	0.652 $\pm$.015	0.628$\pm$.011
CS-SB ₃	0.542 $\pm$.014	0.574 $\pm$.014	0.526 $\pm$.020	0.509$\pm$.013
CS-SB ₄	0.452 $\pm$.015	0.487 $\pm$.016	0.432 $\pm$.022	0.418$\pm$.014

Table 1: Cross-style Self-BLEU (CS-SB_n, $n \in \{1, 2, 3, 4\}$, lower is better) at $T = 0.6$ on Qwen3-32B, **averaged** over the 148-query test set, with 95% paired-bootstrap CIs. SimCTG and CD are additional diversity-enhancing baselines (Section K.1).

	R1-prompt	SFT	SFR
style-0 Ctx	4.5	4.568	4.622
style-0 Rel	4.662	4.764	4.77
style-0 SS	4.654	4.588	4.905
style-0 Flu	4.838	4.831	4.905
style-1 Ctx	4.175	4.216	4.257
style-1 Rel	4.562	4.581	4.628
style-1 SS	3.117	2.858	3.135
style-1 Flu	4.781	4.838	4.791
style-2 Ctx	4.117	4.169	4.128
style-2 Rel	4.285	4.311	4.223
style-2 SS	4.766	4.73	4.824
style-2 Flu	4.964	4.98	4.993
style-3 Ctx	4.511	4.568	4.696
style-3 Rel	4.759	4.818	4.851
style-3 SS	3.891	3.932	4.73
style-3 Flu	4.796	4.831	4.98
<i>The remaining five styles are reported in Table 15.</i>			
Average Ctx	4.368	4.427	4.463
Average Rel	4.604	4.661	4.666
Average SS	4.329	4.243	4.401
Average Flu	4.811	4.86	4.881

Table 2: Stylized-dialogue LLM-judge scores (1–5) at $T = 0.8$ on Qwen3-32B. *Average*: mean across all nine styles. Full per-style breakdown in Section I.

persona prompt.

For cross-style multi-modality we report *Cross-Style Self-BLEU* (CS-SB; Zhu et al., 2018): for each query, we generate one reply per style and compute the pairwise Self-BLEU among the nine resulting replies; the final score averages over all queries in the test set. Lower CS-SB indicates that the model produces genuinely distinct replies under different persona prompts on the same query. We additionally report four LLM-judge metrics scored by Qwen3-235B on a 1–5 Likert scale: *context coherence* (Ctx), measuring whether the response stays on-topic with respect to the query; *query relevance* (Rel), measuring whether the response actually answers what the user asked; *style strength* (SS), measuring how strongly the response exhibits the target persona; and *fluency* (Flu). We aggregate these scores at the style level, reporting both the

per-style value and the across-style average.

Table 1 reports CS-SB averaged across all queries, with 95% paired-bootstrap confidence intervals (Section K.2). Note that absolute CS-SB values are relatively high because all replies are grounded in the same source article, so the semantic core of all replies is shared by design; the metric therefore isolates surface-level stylistic variation. SFR consistently outperforms SFT across all n -gram orders, with non-overlapping CIs at every order (e.g. $0.418 \pm .014$ vs. $0.452 \pm .015$ at $n=4$), indicating that the flow-matching regularizer further diversifies stylistic surface form beyond what SFT alone provides. **SimCTG** and **CD** are two additional diversity-enhancing baselines under the identical setup (Section K.1); neither matches SFR. SimCTG’s CS-SB is close to SFT’s, indicating that its contrastive representation loss does not translate into cross-style separation. CD’s CS-SB is competitive with SFR’s (partially overlapping CIs at $n \geq 2$), but the LLM-judge breakdown in Section K.1 shows this comes with an abnormally low style-strength score, indicating a surface-level decoding artifact rather than genuine style-semantic diversification. The gain is particularly pronounced on open-ended, opinion-rich queries where the stylistic space is large (over 33% relative CS-SB reduction over SFT), and smallest on factual or procedural queries with constrained answer forms; see Section H for a per-query breakdown and qualitative analysis.

Table 2 further shows the four LLM-judge metrics broken down by representative styles and averaged at $T = 0.8$. SFR improves over both DeepSeek-R1-prompt and SFT on most styles in aggregate across the four metrics, with the largest gain on style strength and context coherence; the simultaneous lift on the average scores indicates that the gain is genuinely representational rather than achieved by sacrificing relevance. A random subsample was additionally scored by human annotators; rankings agreed with the LLM-judge on all four dimensions.

Reading the two tables together with the temperature sweep, the picture is consistent: SFR simultaneously raises the four quality scores of Table 2 and lowers the cross-style sameness of Table 1, and this remains true across the temperature sweep of Figure 2. This is the regime in which SFR is designed to help: at low temperature SFT collapses to a near-shared template across personas, while SFR keeps the persona signal alive throughout the




Figure 2: Style strength (SS) versus other quality metrics (Ctx, Rel, Flu) at varying decoding temperatures $T \in \{0.2, 0.4, 0.6, 0.8, 1.0, 1.2, 1.4, 1.6\}$ on Qwen3-32B. SFR yields a better SS–overall-quality trade-off across temperatures, with small metric-specific exceptions at high T where both methods degrade.

response.

Judge robustness and public-data validation.

A cross-family validation judge, **GLM-5.2**, rescored all dialogue outputs alongside our primary judge (Qwen3-235B), rates SFR best on every dimension under both judges, with no same-family bias detected in any SFR-vs-SFT comparison (Section K.3). To rule out that the diversity gain is an artifact of the proprietary dataset, we additionally fine-tune Qwen3-4B on the public PersonaChat (Zhang et al., 2018) dataset (Section K.4 for setup); Table 4 shows the same CS-SB reduction of SFR over SFT reproduces across all four n -gram orders.

4.2 Generalization to open-ended code generation

The diversity-preserving mechanism of SFR is not specific to stylistic variation. It applies whenever the model must maintain multiple valid generation modes under different conditions. Code generation, where a single specification admits diverse correct implementations, provides a natural stress test: here “conditions” are problem descriptions rather than style prompts, and “modes” are functionally equivalent but structurally distinct solutions.

Since the stylized-dialogue dataset is proprietary, we further validate on a public benchmark—LiveCodeBench-v5 (Jain et al., 2025). We fine-tune

Budget	Method	pass@1				pass@5			
		all	easy	medium	hard	all	easy	medium	hard
3B	SFT	44.8	92.3	51.0	13.9	58.2	97.2	75.1	24.4
	SFR	46.7	92.0	55.6	15.1	60.1	96.7	78.5	26.6
6B	SFT	46.5	92.9	53.8	15.4	60.7	97.8	77.3	28.1
	SFR	50.2	94.1	59.8	19.0	64.2	97.5	80.6	33.9
12B	SFT	52.2	94.9	62.6	21.0	65.7	97.5	84.5	34.7
	SFR	54.4	95.9	64.0	24.4	68.7	97.5	86.7	39.1

Table 3: LiveCodeBench-v5 pass@ k at $T = 0.8$ on Qwen2.5-Coder-7B-Instruct (10 samples per problem), across three matched training-token budgets (3B/6B/12B).

Method	CS-SB ₁	CS-SB ₂	CS-SB ₃	CS-SB ₄
SFT	0.653	0.507	0.391	0.299
SFR	0.598	0.441	0.322	0.233

Table 4: CS-SB on public PersonaChat data (Qwen3-4B). Lower is better (= larger inter-persona diversity).

Qwen2.5-Coder-7B-Instruct (Hui et al., 2024) on $\sim 380k$ chain-of-thought samples cleaned from rStar-Code (Liu et al., 2025) (Section B.2) and stop at three matched training-token budgets: **3B**, **6B**, and **12B**, comparing SFT against SFR under matched data, optimizer, and step counts. Decoding uses temperature $T = 0.8$; we report pass@1 and pass@5 computed over 10 independent samples per problem, broken down by LiveCodeBench’s official difficulty splits (Easy, Medium, Hard) and overall (All).

The result is reported in Table 3. On the overall (All) split, SFR improves pass@1 over SFT at every budget (+1.9, +3.7, +2.2 at 3B/6B/12B), and improves pass@5 consistently (+1.9 at 3B, +3.5 at 6B, +3.0 at 12B). The gain is most pronounced on the Medium and Hard splits, where SFR dominates SFT at every budget on pass@1 (e.g. Medium: +4.6, +6.0, +1.4; Hard: +1.2, +3.6, +3.4); on the Easy split the two methods are comparable, which is unsurprising because Easy problems leave little room for multi-modal coverage to matter. The advantage holds at the largest budget (12B) on both pass@1 and pass@5 across all splits (e.g. pass@5 Medium: 86.7 vs 84.5; Hard: 39.1 vs 34.7). The consistent pass@1 improvement is in line with the reading of SFR as enhancing conditional discrimination at the *single-sample* level: even when coverage is already high, per-sample quality continues to benefit.

4.3 Controlled comparison on MBPP

This set of experiments compares SFR against both SFT and multi-token prediction under strictly matched conditions. On **Qwen3-4B-Base** (Yang et al., 2025) we train three objectives on $\sim 500k$ quality-filtered samples from OpenCode-Instruct (Ahmad et al., 2025) (Section B.3), under matched optimizer and step counts: SFT, MTP with $K=4$ prediction heads, and SFR encoding the next $k=32$ tokens as its target (Equation (3)), stride 1. All three checkpoints are evaluated on MBPP and on the harder MBPP+ split (Austin et al., 2021; Liu et al., 2023a) at $T = 0.6$; we report pass@1, pass@5, and pass@10, each computed over 16 independent samples per problem.

Table 5 gives the result. SFR achieves the highest scores on both splits across all three metrics, with a clear margin over both SFT and MTP on pass@5 and pass@10 (e.g. MBPP+ pass@10: 79.0 vs 77.3 for MTP and 76.7 for SFT; MBPP-Base pass@5: 87.4 vs 85.8 for MTP and 85.6 for SFT). On pass@1, SFR also leads on both splits (MBPP-Base: 78.3 vs 77.2/77.1; MBPP+: 67.7 vs 67.2/67.0), though the margin is narrower; the advantage widens at higher k , consistent with the interpretation that continuous future supervision enriches the backbone’s coverage of alternative solution paths in addition to single-shot accuracy. Paired-bootstrap 95% CIs on (method – SFR), computed over problems on the Base split, confirm this pattern is not noise: neither comparison is significant at pass@1 (all three methods sit near a shared ceiling), but SFR’s advantage becomes significant against SFT already at pass@5, and against both SFT and MTP at pass@10.

Combined with the dialogue and the open-ended code-generation results, the three experiments tell a consistent story: the auxiliary flow-matching loss reshapes the backbone’s hidden state so that h_t en-

Method	Split	pass@1	pass@5	pass@10
SFT	Base	77.2	85.6	88.1
MTP	Base	77.1	85.8	88.1
SFR	Base	78.3	87.4	89.2
SFT	Plus	67	74.4	76.7
MTP	Plus	67.2	75.4	77.3
SFR	Plus	67.7	76.8	79.0

Table 5: MBPP and MBPP+ pass@ k at $T = 0.6$ on Qwen3-4B-Base (16 samples per problem). SFR leads on all metrics across both splits.

codes future-aware, condition-discriminative information, which manifests as cross-style distinctiveness in dialogue and as higher pass@ k coverage in code generation. Inference cost is identical to vanilla SFT in all three cases.

5 Analysis

We probe two properties of the trained model: how cleanly the backbone hidden state separates by conditioning prompt (Section 5.1), and how the discarded flow-matching head, when re-attached *offline*, can be used to read out a token-level uncertainty signal that the autoregressive head alone does not expose (Section 5.2).

5.1 Hidden states separate by persona

For each method we collect $\{h_t^{q,s} : q \in Q_{\text{eval}}, s \in S, t = t_0\}$, where Q_{eval} is the evaluation query set and S the set of nine personas, at a fixed response position t_0 (the third response token in the headline result), visualize the resulting vectors with t-SNE (van der Maaten and Hinton, 2008) in Figure 3, and quantify style separation with a *Style Separability Index*, $\text{SSI} = \mathbb{E}_s \text{var}_q(h_t^{q,s}) / \mathbb{E}_q \text{var}_s(h_t^{q,s})$, i.e. within-style scatter divided by between-style scatter (lower is better). At $t_0=3$ the personas overlap heavily under SFT and SSI is large (0.8052); under SFR the same styles separate cleanly and SSI drops to 0.5992, yielding a SFT/SFR ratio of 1.34. The persona prompt has therefore been written into the backbone’s representation under SFR and only weakly so under SFT.

Table 6 sweeps the response position t_0 and shows the same trend: SFR’s SSI is at or below SFT’s at every t_0 we examined, with the largest gap concentrated in the $t_0 \in [2, 5]$ range, where the persona signal is most actively shaping the early body of the reply; later positions converge as both methods commit to a particular continuation.

t_0	SFT	SFR	ratio
1	0.9705	0.9589	1.0121
2	0.6793	0.5434	1.2501
3	0.8052	0.5992	1.3439
5	0.7915	0.6943	1.1401
10	1.0467	1.0029	1.0437
20	1.0431	1.0234	1.0193
40	1.1030	1.0962	1.0061

Table 6: Style Separability Index (SSI, lower is better) of the backbone hidden state at increasing response positions t_0 . Ratio is SFT/SFR; values above 1 indicate that SFR produces better-separated persona groups.




Figure 3: Hidden-state geometry under cross-persona conditioning. t-SNE of $h_t^{q,s}$ for ten queries $\times$ nine personas. Under SFT the per-persona styles mix across queries; under SFR the styles show clearer persona-wise separation across the sampled queries.

5.2 Semantic endpoint dispersion as a Lock/Fork diagnostic

Recent work (Zhang et al., 2026) distinguishes *Lock* positions (one dominant continuation; precision-critical) from *Fork* positions (multiple valid continuations; exploration-critical), but does not provide a mechanism to *locate* which positions are Locks or Forks—the distinction is implicit in the self-distillation process. SFR’s discarded FM head offers an explicit diagnostic: at position t we fix h_t , draw $N=64$ independent source samples $z_0^{(i)} \sim \text{Unif}(\mathbb{S}^{d_z-1})$, and integrate the learned ODE $\dot{z}_\tau = v_\phi(z_\tau, \tau; h_t)$ for 50 Euler steps from $\tau=0$ to $\tau=1$, yielding endpoints $\hat{z}_1^{(i)}$. We then compute the *semantic endpoint dispersion*

$$\text{SED}_t = \text{tr}[\text{Cov}_i(\hat{z}_1^{(i)})], \quad (4)$$

which measures the spread of predicted future semantic endpoints. Unlike softmax entropy, SED is intrinsic to the learned flow field and independent of any decoding temperature. Low SED $\Rightarrow$ Lock; high SED $\Rightarrow$ Fork at segment level rather than single-token level. Section 5.1 shows a concrete example on a quicksort completion. Tokens whose continuations are nearly deterministic

Prompt: Write Python code for quicksort.

Completion: ```python\ndef quicksort(arr):\n """\n Sorts an array using the quicksort algorithm.\n\n Parameters:\n arr\n (list): The list of elements to be sorted.\n\n Returns:\n list: A new list containing all elements from the input list in sorted\n order.\n """\n if len(arr) <= 1:\n return arr\n else:\n pivot = arr[len(arr) // 2]\n left = [x for x in arr if x < pivot]\n middle = [x for x in arr if x == pivot]\n right = [x for x in arr if x > pivot]\n return quicksort(left) + middle\n + quicksort(right)\n```

Token	Parameters	Returns	arr	left	right	return
SED _t	0.669	0.8239	0.7148	0.5951	0.4049	0.3873
Label	Fork	Fork	Fork	Lock	Lock	Lock

Table 7: Lock/Fork diagnosis via SED_t on a sampled completion (excerpt). SED_t values are normalized to [0, 1] by dividing by the per-sequence maximum; the Lock/Fork label is assigned by thresholding at the per-sequence median.

given the algorithm, such as the idiomatic variable names left/right and the keyword return, all receive low SED and are classified as Locks. In contrast, Parameters and Returns introduce free-form natural-language descriptions where multiple valid phrasings exist, and arr appears at the pivot selection line where the partitioning strategy is genuinely underdetermined; all three exhibit high SED and are classified as Forks. We note that the per-sequence median threshold is used here for illustration; calibrating a task-specific threshold is left for future work.

6 Discussion

6.1 Why flow matching

Deterministic regression (MSE / cosine to the encoder target) collapses to the conditional mean $\mathbb{E}[z_1 | c]$ (Proposition 3.1), which mode-averages across conditioning prompts—precisely the pathology we set out to cure. The randomness of z_0 in (3) is what makes the optimum equal to the full conditional vector field (Proposition 3.2); swapping FM for MSE or cosine, while keeping the same auxiliary head and target, erases most of the gain (Table 13, Section G.1). FM’s constant-magnitude target and single-step regression are also far simpler to tune than diffusion DSM as a small auxiliary loss, and the straight-line formulation makes the following equivalence, which underlies the MTP inclusion result Theorem A.3, almost trivial to prove:

Lemma 6.1 (Velocity-matching $\equiv$ endpoint-matching). *Restrict SFR to the discrete, constant-velocity MTP-instance of Definition A.2 (Section A.3): the head predicts a constant velocity $v_{\phi,k} = q_{\phi,k}(\cdot | h_t) - \pi$ toward a fixed reference π , with target velocity $u_{*,k}^{(t)} = e(y_{t+k}) - \pi$. Then, for any divergence D on the tangent space induced by a divergence D_Δ on the simplex via translation*

by π , matching $v_{\phi,k}$ to $u_{*,k}^{(t)}$ under D is identical to matching the endpoint prediction $q_{\phi,k}(\cdot | h_t)$ to the one-hot target $e(y_{t+k})$ under D_Δ .

Taking D_Δ to be the KL divergence recovers exactly the MTP cross-entropy loss (Theorem A.3); the full discussion, including why CFM rather than diffusion, is deferred to Sections A.3 and A.6.

6.2 Scope

SFR does not change the inference graph: the deployed checkpoint has the same architecture as a vanilla AR LM and runs on any existing serving stack. It also does not verify factual consistency of responses against the persona specification—that is an orthogonal fact-checking problem. As Figure 2 shows, SFR’s practical edge is largest in the low-temperature regime that production systems prefer; at $T \geq 1.0$ sampling randomness dominates and the two methods converge. The formal relationship between SFR and multi-token prediction is detailed in Section A.

7 Conclusion

We identified Cross-Style Collapse, the failure of fine-tuned LLMs to produce diverse outputs under varying conditioning, and proposed Semantic Flow Regularization (SFR), a conditional flow-matching auxiliary loss whose head is discarded at inference for zero deployment cost. On industrial dialogue (Qwen3-32B, 9 personas), SFR improves cross-style diversity while raising quality. Public validation on LiveCodeBench-v5 and MBPP confirms generality: SFR consistently lifts pass@ k and theoretically subsumes Multi-Token Prediction as a special case. As a byproduct, the offline FM head yields a segment-level Lock/Fork uncertainty probe beyond the AR softmax. Extending SFR to pre-training and RLHF is future work.

Limitations

Hyperparameter sensitivity. SFR introduces several hyperparameters beyond vanilla SFT: the loss weight λ , suffix length k , stride s , warmup schedule, and EMA decay, as well as the choice of sentence encoder E itself. The ablations in Section G demonstrate broad robustness across reasonable ranges, but finding an optimal configuration still requires non-trivial grid search; no single default set works universally across all backbones and tasks.

Training dynamics of the FM head. The flow-matching loss interacts with the main CE loss in subtle ways. If λ is set too high or the FM head warms up too quickly, it can “short-circuit” the backbone by dominating the gradient signal before the AR head has converged, degrading next-token accuracy. Conversely, insufficient warmup leads to large initial FM errors that propagate noisy gradients into the backbone. In practice we find it necessary to monitor the FM loss curve throughout training and adjust the warmup length accordingly—an extra operational burden absent in standard SFT.

Style data coverage and distribution. Our experiments use a fixed set of nine hand-curated personas with roughly balanced data. When the number of styles grows substantially (e.g. hundreds of personas) or the per-style data distribution is highly imbalanced, minority styles may receive insufficient gradient signal from the flow-matching objective, causing the backbone to under-represent them. Whether curriculum sampling or per-style reweighting can mitigate this is left to future work. More broadly, our stylized benchmark does not cover continuous-spectrum styles or cross-language stylization.

Scaling to larger backbones. Our largest backbone is 32B parameters. Online encoding of suffix targets adds roughly 10% to training wall-clock time at stride 4 (and proportionally less at larger strides), which is the explicit cost paid for zero inference overhead. Whether the regularization effect strengthens or weakens at 70B+ scales, and whether the additional memory footprint of the FM head and encoder becomes a practical bottleneck at such scales, remain open questions.

Scope of the factuality evidence. Our stylized-dialogue task is a *grounded* setting: the model’s input always includes the relevant source article, so a high context-coherence score is itself a partial indicator that diversity gains are not com-

ing at the expense of drifting away from the given context (Section K). We have not evaluated SFR in a fully *open-ended* setting where the model must generate factual content from its own parametric knowledge without any supporting context (e.g. a TruthfulQA-style benchmark). Our current evidence on the diversity–factuality trade-off is therefore limited to grounded generation, and verifying it in open-ended settings is left to future work.

Potential risks. By explicitly encouraging output diversity, SFR may amplify the generation of undesirable content if the training data contains harmful, biased, or toxic material—the method does not distinguish between “diverse but safe” and “diverse but harmful” continuations. Standard content-filtering and alignment safeguards (e.g. RLHF, safety classifiers) remain necessary downstream. Additionally, improved style-mimicking ability could facilitate misuse in impersonation or disinformation scenarios; deployers should apply appropriate access controls and monitoring.

References

- Wasi Uddin Ahmad, Aleksander Ficek, Mehrzad Samadi, Jocelyn Huang, Vahid Noroozi, Somshubra Majumdar, and Boris Ginsburg. 2025. [Opencodeinstruct: A large-scale instruction tuning dataset for code llms](#). *CoRR*, abs/2504.04030.
- Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael G. Rabbat, Yann LeCun, and Nicolas Ballas. 2023. [Self-supervised learning from images with a joint-embedding predictive architecture](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 15619–15629. IEEE.
- Jacob Austin, Augustus Odena, Maxwell I. Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie J. Cai, Michael Terry, Quoc V. Le, and Charles Sutton. 2021. [Program synthesis with large language models](#). *CoRR*, abs/2108.07732.
- Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mido Assran, and Nicolas Ballas. 2024. [Revisiting feature prediction for learning visual representations from video](#). *Trans. Mach. Learn. Res.*, 2024.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2020. [Plug and play language models: A simple approach to controlled text generation](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

- DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, and 181 others. 2025. [Deepseek-v3 technical report](#). *Preprint*, arXiv:2412.19437.
- Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, Shrimai Prabhumoye, Alan W. Black, Alexander Rudnicky, Jason Williams, Joelle Pineau, Mikhail Burtsev, and Jason Weston. 2020. [The second conversational intelligence challenge \(convai2\)](#). In *The NeurIPS '18 Competition*, The Springer Series on Challenges in Machine Learning, pages 187–208, Cham. Springer.
- Xingyu Fan, Feifei Li, Wenhui Que, and Hailong Li. 2025. [One agent to serve all: a lite-adaptive stylized AI assistant for millions of multi-style official accounts](#). *CoRR*, abs/2509.17788.
- Elias Frantar, Saleh Ashkboos, Torsten Hoefer, and Dan Alistarh. 2022. [GPTQ: accurate post-training quantization for generative pre-trained transformers](#). *CoRR*, abs/2210.17323.
- Xiang Gao, Yizhe Zhang, Sungjin Lee, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2019. [Structuring latent spaces for stylized response generation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 1814–1823. Association for Computational Linguistics.
- Fabian Gloeckle, Badr Youbi Idrissi, Baptiste Rozière, David Lopez-Paz, and Gabriel Synnaeve. 2024. [Better & faster large language models via multi-token prediction](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, Proceedings of Machine Learning Research, pages 15706–15734. PMLR / OpenReview.net.
- Tilmann Gneiting and Adrian E. Raftery. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378.
- Ishaan Gulrajani and Tatsunori B. Hashimoto. 2023. [Likelihood-based diffusion language models](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shiron Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025. [Deepseek-r1 incentivizes reasoning in LLMs through reinforcement learning](#). *Nature*, 645(8081):633–638.
- Tiankai Hang, Shuyang Gu, Chen Li, Jianmin Bao, Dong Chen, Han Hu, Xin Geng, and Baining Guo. 2023. [Efficient diffusion training via min-snr weighting strategy](#). In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 7407–7417. IEEE.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Kai Dang, An Yang, Rui Men, Fei Huang, Xingzhang Ren, Xuancheng Ren, Jingren Zhou, and Junyang Lin. 2024. [Qwen2.5-coder technical report](#). *CoRR*, abs/2409.12186.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2025. [LiveCodeBench: Holistic and contamination free evaluation of large language models for code](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. 2022. [Elucidating the design space of diffusion-based generative models](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. 2019. [CTRL: A conditional transformer language model for controllable generation](#). *CoRR*, abs/1909.05858.
- Ben Krause, Akhilesh Deepak Gotmare, Bryan McCann, Nitish Shirish Keskar, Shafiq R. Joty, Richard Socher, and Nazneen Fatema Rajani. 2021. [Gedi: Generative discriminator guided sequence generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, Findings of ACL, pages 4929–4952. Association for Computational Linguistics.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with PagedAttention](#). pages 611–626.
- LCM team, Loïc Barrault, Paul-Ambroise Duquenne, Maha Elbayad, Artyom Kozhevnikov, Belen Alastruey, Pierre Andrews, Mariano Coria, Guillaume Couairon, Marta R. Costa-jussà, David Dale, Hady Elsahar, Kevin Heffernan, João Maria Janeiro,

- Tuan Tran, Christophe Ropers, Eduardo Sánchez, Robin San Roman, Alexandre Mourachko, and 2 others. 2024. [Large concept models: Language modeling in a sentence representation space](#). *CoRR*, abs/2412.08821.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 110–119. The Association for Computational Linguistics.
- Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. 2023. [Contrastive decoding: Open-ended text generation as optimization](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 12286–12312. Association for Computational Linguistics.
- Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B. Hashimoto. 2022. [Diffusion-lm improves controllable text generation](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. 2024. [AWQ: activation-aware weight quantization for on-device LLM compression and acceleration](#). In *Proceedings of the Seventh Annual Conference on Machine Learning and Systems, MLSys 2024, Santa Clara, CA, USA, May 13-16, 2024*. mlsys.org.
- Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. 2023. [Flow matching for generative modeling](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2023a. [Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Qiang Liu. 2022. [Rectified flow: A marginal preserving approach to optimal transport](#). *CoRR*, abs/2209.14577.
- Xingchao Liu, Chengyue Gong, and Qiang Liu. 2023b. [Flow straight and fast: Learning to generate and transfer data with rectified flow](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Yifei Liu, Li Lina Zhang, Yi Zhu, Bingcheng Dong, Xudong Zhou, Ning Shang, Fan Yang, and Mao Yang. 2025. [rstar-coder: Scaling competitive code reasoning with a large-scale verified dataset](#). *CoRR*, abs/2505.21297.
- Aaron Lou, Chenlin Meng, and Stefano Ermon. 2024. [Discrete diffusion modeling by estimating the ratios of the data distribution](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, Proceedings of Machine Learning Research, pages 32819–32848. PMLR / OpenReview.net.
- Phil Ostheimer, Mayank Nagda, Marius Kloft, and Sophie Fellenz. 2024. [Text style transfer evaluation using large language models](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 15802–15822, Torino, Italia. ELRA and ICCL.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. [Zero: memory optimizations toward training trillion parameter models](#). In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2020, Virtual Event / Atlanta, Georgia, USA, November 9-19, 2020*, page 20. IEEE/ACM.
- Tim Salimans and Jonathan Ho. 2022. [Progressive distillation for fast sampling of diffusion models](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron C. Courville, and Yoshua Bengio. 2017. [A hierarchical latent variable encoder-decoder model for generating dialogues](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 3295–3301. AAAI Press.
- Yixuan Su, Tian Lan, Yan Wang, Dani Yogatama, Lingpeng Kong, and Nigel Collier. 2022. [A contrastive framework for neural text generation](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Jack Urbanek, Angela Fan, Siddharth Karamcheti, Saachi Jain, Samuel Humeau, Emily Dinan, Tim Rocktäschel, Douwe Kiela, Arthur Szlam, and Jason Weston. 2019. [Learning to speak and act in a fantasy text adventure game](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International*

- Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 673–683. Association for Computational Linguistics.
- Laurens van der Maaten and Geoffrey Hinton. 2008. [Visualizing data using t-SNE](#). *Journal of Machine Learning Research*, 9(86):2579–2605.
- Noah Wang, Zhongyuan Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhao Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, Man Zhang, Zhaoxiang Zhang, Wanli Ouyang, Ke Xu, Wenhao Huang, Jie Fu, and Junran Peng. 2024. [RoleLLM: Benchmarking, eliciting, and enhancing role-playing abilities of large language models](#). In *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, Findings of ACL, pages 14743–14777. Association for Computational Linguistics.
- Sean Welleck, Ilia Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2020. [Neural text generation with unlikelihood training](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Shitao Xiao, Zheng Liu, Peitian Zhang, Niklas Muenighoff, Defu Lian, and Jian-Yun Nie. 2024. [C-pack: Packed resources for general chinese embeddings](#). pages 641–649.
- Yunze Xiao, Vivienne J. Zhang, Chenghao Yang, Ningshan Ma, Weihao Xuan, and Jen-Tse Huang. 2026. [The chameleon’s limit: Investigating persona collapse and homogenization in large language models](#). *CoRR*, arXiv:2604.24698.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Kevin Yang and Dan Klein. 2021. [FUDGE: controlled text generation with future discriminators](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 3511–3535. Association for Computational Linguistics.
- Ruixiang Zhang, Richard He Bai, Huangjie Zheng, Navdeep Jaitly, Ronan Collobert, and Yizhe Zhang. 2026. [Embarrassingly simple self-distillation improves code generation](#). *CoRR*, abs/2604.01193.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. [Personalizing dialogue agents: I have a dog, do you have pets too?](#)
- Shiyue Zhang, Shijie Wu, Ozan Irsoy, Steven Lu, Mohit Bansal, Mark Dredze, and David S. Rosenberg. 2023. [Mixce: Training autoregressive language models by mixing forward and reverse cross-entropies](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 9027–9050. Association for Computational Linguistics.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskénazi. 2017. [Learning discourse-level diversity for neural dialog models using conditional variational autoencoders](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 654–664. Association for Computational Linguistics.
- Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Sun, Jeff Huang, Cody Hao Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E. Gonzalez, Clark W. Barrett, and Ying Sheng. 2024. [Sglang: Efficient execution of structured language model programs](#). In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Jinfeng Zhou, Zhuang Chen, Dazhen Wan, Bosi Wen, Yi Song, Jifan Yu, Yongkang Huang, Libiao Peng, Jiaming Yang, Xiyao Xiao, Sahand Sabour, Xiaohan Zhang, Wenjing Hou, Yijia Zhang, Yuxiao Dong, Jie Tang, and Minlie Huang. 2023. [CharacterGLM: Customizing chinese conversational AI characters with large language models](#). *CoRR*, abs/2311.16832.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. [Texygen: A benchmarking platform for text generation models](#). In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, Ann Arbor, MI, USA, July 08-12, 2018*, pages 1097–1100. ACM.

A Theoretical Foundations

This appendix collects the formal statements and full proofs that support the claims made in Sections 3.2, 6.1 and 6.2. We first set up notation (Section A.1), then prove the multi-modality argument for choosing CFM over deterministic alignment (Section A.2), then prove the strict-equivalence results that place MTP as a special case of SFR (Section A.3), and finally summarize the resulting inclusion chain (Section A.4). Section A.6 addresses the related question of why CFM rather than diffusion.

A.1 Notation and the variational view

Let V be the vocabulary, $|V|$ its size, and $\Delta := \Delta^{|V|-1}$ the probability simplex. For a token $v \in V$

write $e(v) \in \Delta$ for its one-hot embedding. We use $h_t \in \mathbb{R}^d$ for a backbone hidden state and $z_1 \in \mathbb{R}^{d_z}$ for an auxiliary target (a continuous embedding in the body of the paper, a point on Δ^K in Definition A.2).

The token-level CE loss and the FM auxiliary loss used in the body are

$$\mathcal{L}_{\text{AR}} = \sum_t \text{CE}(\text{softmax}(W_{\text{AR}} h_t), y_{t+1}), \quad (5)$$

$$\mathcal{L}_{\text{SFR}} = \mathbb{E}_{\tau, z_0} \|v_\phi(z_\tau, \tau; h_t) - (z_1 - z_0)\|_2^2, \quad (6)$$

with $\tau \sim \mathcal{U}[0, 1]$, $z_0 \sim p_0$ ($p_0 = \text{Unif}(\mathbb{S}^{d_z-1})$ in our implementation), and the path $z_\tau = (1-\tau)z_0 + \tau z_1$.

Variational view (post-hoc interpretation). The following is a *retrospective* probabilistic reading of the training objective, not the generative model from which the algorithm is derived. Introduce a continuous latent $z \in \mathcal{Z}$ representing the global semantic intent of a response. Marginalizing,

$$P(y | x, s) = \int_{\mathcal{Z}} P_{\text{AR}}(y | z, x, s) P_{\text{SFR}}(z | x, s) dz. \quad (7)$$

Pure SFT corresponds to the degenerate choice $P_{\text{SFR}}(z | x, s) = \delta_{z^*(x,s)}$, i.e. a single prescribed semantic intent for each (x, s) . SFR replaces this Dirac by an explicit continuous distribution that the auxiliary head v_ϕ approximates via flow matching. An informal ELBO-style motivation that rationalizes adding $\lambda \mathcal{L}_{\text{SFR}}$ to $\mathcal{L}_{\text{AR}}$ is given in Section A.5 below.

A.2 Why FM and not deterministic regression

The following two results, stated in Section 3.2 (Propositions 3.1 and 3.2), formalise the choice made there: any deterministic predictor of z_1 is a conditional-mean estimator and therefore mode-averages, whereas CFM with a random source preserves all modes of $P(z_1 | c)$. We give the proofs below.

Proof of Proposition 3.1. Let $\mu(c) = \mathbb{E}[z_1 | c]$ and decompose $z_1 = \mu(c) + \varepsilon$ with $\mathbb{E}[\varepsilon | c] = 0$. For any measurable f ,

$$\begin{aligned} \mathbb{E}\|f(c) - z_1\|_2^2 &= \mathbb{E}\|f(c) - \mu(c) - \varepsilon\|_2^2 \\ &= \mathbb{E}\|f(c) - \mu(c)\|_2^2 + \mathbb{E}\|\varepsilon\|_2^2, \end{aligned} \quad (8)$$

because the cross term $\mathbb{E}\langle f(c) - \mu(c), \varepsilon \rangle$ vanishes by the tower property. Both terms are non-negative; the second is independent of f . Equality holds iff $f = \mu$ a.e. For two well-separated modes $\{z_1^{(a)}, z_1^{(b)}\}$ with masses $p, 1-p$, $\mu(c) = p z_1^{(a)} + (1-p) z_1^{(b)}$ lies on the line segment between them, and hence in a low-density region of $P(z_1 | c)$. $\square$

Proof of Proposition 3.2. Step 1. By the tower property, $\mathbb{E}_{\tau, z_0, z_1} \|v_\phi(z_\tau, \tau, c) - (z_1 - z_0)\|^2 = \mathbb{E}_{\tau, z_\tau} \mathbb{E}_{z_0, z_1 | z_\tau, c} \|v_\phi(z_\tau, \tau, c) - (z_1 - z_0)\|^2$. The inner expectation is the L^2 distance from a deterministic point to the random vector $z_1 - z_0$, whose L^2 projection is its conditional mean $v^*(z, \tau, c)$ by Proposition 3.1.

Step 2. Using $z_\tau = (1-\tau)z_0 + \tau z_1$, $\dot{z}_\tau = z_1 - z_0$. By the tower property $\mathbb{E}[\dot{z}_\tau | z_\tau = z, c] = v^*(z, \tau, c)$. The marginal density satisfies the continuity equation

$$\partial_\tau p_\tau(z | c) + \nabla_{z \cdot} (p_\tau(z | c) v^*(z, \tau, c)) = 0, \quad (9)$$

so the deterministic ODE generated by v^* transports p_0 to $p_1(\cdot | c)$ at $\tau = 1$. Multi-modality of p_1 is preserved at the distributional level because the terminal marginal of the probability path equals $p_1(\cdot | c)$ under the exact marginal velocity field. See (Lipman et al., 2023, Thm. 3) for the full statement. $\square$

Remark A.1 (Engineering consequence). Replacing $z_0 \sim \text{Unif}(\mathbb{S}^{d_z-1})$ in (6) by a constant $z_0 \equiv c_0$ collapses $\mathcal{L}_{\text{SFR}}$ to $\mathbb{E}\|f(h_t) - z_1\|^2$ with $f(h_t) := v_\phi(\cdot; h_t) + c_0$, hence to deterministic regression. By Proposition 3.1 this objective then mode-averages. The randomness of z_0 is therefore essential and not a hyperparameter.

A.3 MTP as a strict special case of SFR

Definition A.2 (MTP-instance of the SFR family). For a fixed reference $\pi \in \text{int}(\Delta)$ and horizon $K \geq 0$, the *MTP-instance* of SFR is obtained by:

1. taking the state space to be Δ^K ;
2. taking the per-position target to be the one-hot stack $z_1^{(t)} = [e(y_{t+1}), \dots, e(y_{t+K})] \in \Delta^K$;
3. taking the source to be the constant $p_0 = \delta_{[\pi, \dots, \pi]}$;
4. using the straight-line path $z_{\tau, k}^{(t)} = (1-\tau)\pi + \tau e(y_{t+k})$;
5. restricting the head to a constant-velocity field $v_{\phi, k}(z_\tau, \tau; h_t) := q_{\phi, k}(\cdot | h_t) - \pi$ that is in-

dependent of (z_τ, τ) , where $q_{\phi,k}(\cdot | h_t) = \text{softmax}(W_k h_t + b_k) \in \Delta$.

Lemma 6.1 (stated in Section 6.1) makes this precise; we give the proof below.

Proof of Lemma 6.1. The tangent space of the simplex Δ is $T\Delta = \{u \in \mathbb{R}^{|V|} \mid \sum_i u_i = 0\}$. By Definition A.2(5), $v_{\phi,k} = q_{\phi,k} - \pi$ and $u_{\star,k}^{(t)} = e(y_{t+k}) - \pi$. Since $q_{\phi,k}, e(y_{t+k}) \in \Delta$, both $v_{\phi,k}$ and $u_{\star,k}^{(t)}$ lie in the tangent space $T\Delta$. Under the translation bijection $u \mapsto u + \pi$, i.e., $D(u||v) = D_\Delta(u + \pi || v + \pi)$ for $u, v \in T\Delta$, we have:

$$\begin{aligned} D(v_{\phi,k}(\cdot; h_t) || u_{\star,k}^{(t)}) \\ &= D_\Delta(v_{\phi,k}(\cdot; h_t) + \pi || u_{\star,k}^{(t)} + \pi) \\ &= D_\Delta(q_{\phi,k}(\cdot | h_t) || e(y_{t+k})). \end{aligned}$$

Euclidean case: For $D_\Delta(p||q) = \|p - q\|_2^2$, this translation yields:

$$\begin{aligned} \|v_{\phi,k} - u_{\star,k}^{(t)}\|_2^2 &= \|(q_{\phi,k} - \pi) - (e(y_{t+k}) - \pi)\|_2^2 \\ &= \|q_{\phi,k} - e(y_{t+k})\|_2^2. \end{aligned}$$

□

Bregman / KL case: Since the Bregman divergence $D_\Psi(p||q) = \sum_i p_i \log(p_i/q_i)$ is defined on Δ , and the constant-velocity parameterisation provides a bijection between velocity space $\Delta - \pi$ and endpoint space Δ , we define the velocity-matching loss in this geometry as $D_\Psi(e(y_{t+k})||q_{\phi,k})$ —the unique divergence on the endpoint space that the bijection $v_{\phi,k} \leftrightarrow q_{\phi,k}$ induces.

We define

$$\mathcal{L}_{\text{MTP}} = \sum_t \sum_{k=1}^K \alpha_k \text{CE}(q_{\phi,k}(\cdot | h_t), y_{t+k}), \quad (10)$$

the standard MTP loss with non-negative weights α_k .

Theorem A.3 (SFR $\equiv$ MTP under Bregman / KL geometry). *Let $\Psi(p) = \sum_i p_i \log p_i$ be the negative-entropy potential on Δ , with Bregman divergence $D_\Psi(p||q) = \Psi(p) - \Psi(q) - \langle \nabla \Psi(q), p - q \rangle$. For one-hot $p = e(y)$, $D_\Psi(e(y)||q) = -\log q_y = \text{KL}(e(y)||q)$. Under Definition A.2, the corresponding SFR loss*

$$\mathcal{L}_{\text{SFR}}^\Psi = \sum_t \sum_{k=1}^K \alpha_k D_\Psi(e(y_{t+k}) || q_{\phi,k}(\cdot | h_t)) \quad (11)$$

satisfies the strict identity $\mathcal{L}_{\text{SFR}}^\Psi \equiv \mathcal{L}_{\text{MTP}}$.

Proof. By Lemma 6.1 with $D = D_\Psi$, the velocity-matching loss reduces to the endpoint-matching loss (11). Each term equals $D_\Psi(e(y_{t+k})||q_{\phi,k}) = -\log q_{\phi,k}(y_{t+k} | h_t)$, the per-token CE term in (10). □

Theorem A.4 (Brier-MTP relative under Euclidean geometry). *Replacing D_Ψ in (11) by squared Euclidean distance, $D(p||q) := \|p - q\|_2^2$, gives a near-cousin $\mathcal{L}_{\text{SFR}}^{\ell_2} = \sum_t \sum_k \alpha_k \|q_{\phi,k} - e(y_{t+k})\|_2^2$ whose population minimiser coincides with that of $\mathcal{L}_{\text{MTP}}$: $q_{\phi,k}^*(v | h_t) = P^*(y_{t+k} = v | h_t)$. The two losses differ only in their gradient geometry; Euclidean penalizes confident mistakes substantially less aggressively than Bregman / log-loss.*

Proof. Both $-\log q_y$ and $\|q - e(y)\|_2^2$ are strictly proper scoring rules on Δ (Gneiting and Raftery, 2007), so each is uniquely minimized at $q = p$ in expectation under any conditional distribution $P^*(\cdot | c)$. The gradient remark follows from $\partial_q(-\log q_y) = -e_y \oslash q$ (unbounded as $q_y \rightarrow 0$) versus $\partial_q\|q - e(y)\|_2^2 = 2(q - e(y))$ (bounded by $2\sqrt{2}$ on Δ). □

A.4 Inclusion chain

Corollary A.5 (Inclusion chain). *Let $\mathcal{F}_{\text{NT}} = \{\mathcal{L}_{\text{AR}}\}$, $\mathcal{F}_{\text{MTP}} = \{\mathcal{L}_{\text{AR}} + \lambda \mathcal{L}_{\text{MTP}}\}$, $\mathcal{F}_{\Delta\text{-FM}}$ the family obtained by varying Definition A.2 only over the head’s parameterisation, and $\mathcal{F}_{\text{SFR}}$ the general continuous-target family of the body. Then*

$$\mathcal{F}_{\text{NT}} \subsetneq \mathcal{F}_{\text{MTP}} = \mathcal{F}_{\Delta\text{-FM}} \subsetneq \mathcal{F}_{\text{SFR}}. \quad (12)$$

Proof. Setting $K = 0$ in Definition A.2 recovers $\mathcal{L}_{\text{AR}}$, giving the first strict inclusion. The middle equality is Theorem A.3. For the last strict inclusion we exhibit two structural gaps that $\mathcal{F}_{\Delta\text{-FM}}$ cannot bridge:

1. *Source randomness:* $\mathcal{F}_{\Delta\text{-FM}}$ fixes $p_0 = \delta_{[\pi, \dots, \pi]}$, so its velocity field is deterministic given h_t ; it models token-level uncertainty in the discrete simplex but lacks stochastic continuous source modeling and segment-level semantic targets. $\mathcal{F}_{\text{SFR}}$ with $z_0 \sim \text{Unif}(\mathbb{S}^{d_z-1})$ preserves all modes (Proposition 3.2).
2. *Target dimensionality:* Δ^K is a $K(|V|-1)$ -dimensional bounded polytope that encodes only the identity of the next K tokens; it cannot represent continuous semantic attributes (style intensity, syntactic structure, or discourse coherence) that live in directions orthogonal to the one-hot vertices. The contin-

uous target $z_1 \in \mathbb{R}^{d_z}$ of $\mathcal{F}_{\text{SFR}}$ has no such restriction.

Either gap alone suffices for strict inclusion. $\square$

Table 5 of the body provides the empirical instantiation of this chain on Qwen3-4B / MBPP.

A.5 ELBO interpretation (informal motivation)

The purpose of this subsection is to provide an *informal* variational motivation for adding $\lambda\mathcal{L}_{\text{SFR}}$ to $\mathcal{L}_{\text{AR}}$. The argument is heuristic rather than a rigorous bound, since the actual training procedure does not instantiate a proper latent variable model at inference time.

Setup. Consider the hypothetical generative model of (7): draw $z \sim P_{\text{SFR}}(z \mid x, s)$ then $y \sim P_{\text{AR}}(y \mid z, x, s)$. Let $r_\phi(z \mid h)$ be the variational distribution implied by the auxiliary head—under SFR, the law of the time-1 marginal of the ODE $\dot{z}_\tau = v_\phi(z_\tau, \tau; h)$ with $z_0 \sim p_0$; under deterministic alignment, the Dirac $\delta_{f_\phi(h)}$. Applying Jensen’s inequality to (7) and following the standard ELBO derivation,

$$\log P(y \mid x, s) \geq \mathbb{E}_{z \sim r_\phi(z \mid h)} \log P_{\text{AR}}(y \mid z, x, s) - \text{KL}(r_\phi(z \mid h) \parallel P_{\text{SFR}}(z \mid x, s)). \quad (13)$$

Informal identification with training losses. In our actual model, h_t does not condition on an explicit z at inference; rather, the backbone is *shaped* by the FM gradient so that h_t implicitly encodes future-aware information. Under this lens:

- The reconstruction term $\mathbb{E}_{z \sim r_\phi} \log P_{\text{AR}}(y \mid z, x, s)$ plays the role analogous to $-\mathcal{L}_{\text{AR}}$: maximizing the likelihood of the response given the latent plan.
- The KL term $\text{KL}(r_\phi \parallel P_{\text{SFR}})$ is related to $\mathcal{L}_{\text{SFR}}$: under regularity conditions the FM loss controls the distance between the learned push-forward r_ϕ and the data distribution $p_1(\cdot \mid c)$, since minimizing the velocity-matching objective drives the generated marginal toward the target (Lipman et al., 2023).

Combining, $\mathcal{L}_{\text{AR}} + \lambda\mathcal{L}_{\text{SFR}}$ can be viewed as an approximate β -ELBO surrogate of $\log P(y \mid x, s)$ with $\lambda = \beta$. We emphasize that this is a *motivational analogy*—the formal contributions of this appendix rest on Propositions 3.1 and 3.2, theorems A.3 and A.4, and corollary A.5, which do not depend on this ELBO interpretation.

A.6 Why CFM and not diffusion as the auxiliary loss

The substantive question is not whether CFM and diffusion can each be used, but which one is better suited as a *small auxiliary* on top of an AR loss. We give three reasons; the same argument is summarized in Section 6.1.

Single-step, simulation-free, $O(1)$ -magnitude target. The CFM target $z_1 - z_0$ is a constant vector with $O(1)$ norm under L2-normalized z_1 , sampled via one straight-line interpolation. The DSM-style diffusion loss $\mathbb{E}_t[w(t)\|\epsilon_\theta(z_t, t) - \epsilon\|^2]$ requires choosing the noise schedule, the weighting $w(t)$, and (for stable training) the parameterisation among ϵ, x_0, v , with established literature on Min-SNR weighting (Hang et al., 2023), EDM preconditioning (Karras et al., 2022), and v -prediction (Salimans and Ho, 2022). When the diffusion loss is the *primary* objective, this engineering is well-understood; when it is added as a *small fraction* of an AR loss, the per-batch loss variance under typical schedules is orders of magnitude higher than $\mathcal{L}_{\text{AR}}$, making the effective λ extremely sensitive.

Algebraic naturalness of the MTP correspondence. Theorem A.3 relies on three structural properties of straight-line CFM: a constant target velocity, a translation-invariant divergence, and a constant-velocity head. The diffusion equivalent of this argument requires reparameterizing MTP as a noise-conditional softmax with a learned variance schedule, which is mathematically possible but obscures the connection.

Empirical separation from JEPA / alignment. Proposition 3.1 formalizes the failure mode that JEPA-style EMA-teacher predictors share with MSE/cosine alignment: their optima collapse to the conditional mean. The randomness of z_0 in CFM, by contrast, is what makes the objective preserve all modes of $P(z_1 \mid c)$. This is the precise sense in which FM-based SFR is not a relabelling of JEPA.

We remark that under straight-line probability paths, CFM and score-based diffusion are mathematically equivalent at the level of marginal probability paths (Liu et al., 2023b; Lipman et al., 2023); the distinction in this appendix is at the level of the loss form and its hyperparameter sensitivity, not the hypothesis class.

B Data construction

This appendix details the training and evaluation data for each experiment. Code and configurations to reproduce all experiments are available at <https://github.com/WeAgentAI/SFR>; the proprietary stylized-dialogue data itself is subject to internal compliance review, and we will release de-identified evaluation subsets separately once that review is complete.

B.1 Stylized dialogue (Experiment 1)

Source. We use proprietary data from a widely-used industrial official-accounts platform (as of April 2026). The platform contains over one million accounts; the raw data includes published articles and authentic user-comment-author-reply pairs.

Training set construction. Following the style-extraction pipeline of Fan et al. (2025), we identify nine representative stylistic personas from the platform corpus. We then prompt an LLM to cross-generate stylized QA pairs: for each (context, query, answer) triple, the model produces a response conditioned on each of the nine styles. The prompt templates for both style extraction and cross-style generation follow the designs of Fan et al. (2025). We score every generated sample on four dimensions—context coherence (Ctx), query relevance (Rel), style strength (SS), and fluency (Flu)—using a judge LLM, following prior practice (Wang et al., 2024; Ostheimer et al., 2024). Only top-scoring samples are retained, yielding approximately 170k training examples for full-parameter fine-tuning.

Evaluation set. To the best of our knowledge, no public dataset simultaneously provides articles, user queries, and large-scale stylized replies from millions of authors. We construct a held-out set of ~ 150 examples: one portion consists of human-curated stylized QA pairs drawn directly from real platform interactions, and the remaining portion is synthesized and filtered with the same pipeline as the training set (but on disjoint queries).

B.2 LiveCodeBench (Experiment 2)

We use the open-source rStar-Code dataset (Liu et al., 2025), which contains long chain-of-thought code solutions. The cleaning pipeline is:

1. Retain only samples marked both `verified` and `is_passed`.

2. Remove samples whose chain-of-thought exceeds 16k tokens.
3. Deduplicate problems by adjacent-hash matching on the problem statement.

After filtering, approximately 380k training examples remain.

B.3 MBPP (Experiment 3)

We use the open-source OpenCodeInstruct dataset (Ahmad et al., 2025) (solutions without a reasoning trace). The filtering criteria are:

1. LLM-judged `requirement_conformance` = 5 and `logical_correctness` = 5.
2. LLM-judged `edge_case_consideration` ≥ 4 .
3. All unit tests pass (`tests_execution_status`).
4. Exact-match deduplication on the problem input (retaining the highest-scoring, most concise solution per problem).

After filtering, approximately 500k training examples remain.

C Implementation notes

C.1 Why we avoid runtime `requires_grad` toggling

Under DeepSpeed ZeRO-3 (Rajbhandari et al., 2020), the optimizer’s parameter groups are bound once at `accelerator.prepare()` and the gradient-allreduce graph is materialized at the first backward call. Switching individual parameters’ `requires_grad` between steps causes the actual gradient state to drift out of sync with the parameter-group view, which manifests as silent NCCL hangs at the very step the toggle happens. We therefore implement Phase A by setting `h_for_fm = h.detach()` inside the FM branch only; the backbone parameters’ `requires_grad` are left unchanged for the entire training run.

C.2 Submodule-order consistency under ZeRO-3

ZeRO-3 tracks the per-rank order in which submodules are visited during forward to coordinate parameter gather/scatter. If a rank skips the FM-head forward (e.g. a sample with no response tokens), the order diverges and NCCL deadlocks. We therefore execute a tiny “dummy forward” through the FM head (single dummy token, no gradient) on any rank that would otherwise skip it.

C.3 Cross-rank synchronization of stochastic decisions

Three quantities are synchronized via collectives: (i) the gradient-gating decision (`all_reduce_MAX` so any rank that wants to gate causes all ranks to gate); (ii) the running EMA of `fm_loss` (`all_reduce_SUM`, divided by world size); and (iii) the EMA-teacher update step is performed identically on every rank because the student parameters are already synchronized after the optimizer step.

C.4 Online suffix encoding under multi-GPU training

For the token-level suffix mode we load a frozen encoder onto each local GPU outside the ZeRO-3 scope (see above). To minimize encoder calls we use sparse anchors with stride s ; non-anchor positions inherit the embedding of their nearest anchor. The encoder’s parameters are explicitly excluded from the optimizer and from checkpointing.

C.5 FM-head detachment at `save_model`

At checkpoint time we temporarily pop the FM-head out of the model’s submodule dict on every rank, run `super().save_model`, then re-attach. This keeps the published checkpoint structurally identical to a plain AR LM, ensuring that downstream tools (`from_pretrained`, `vLLM`, LoRA merging, GP-TQ/AWQ quantization) work without modification.

D Hyper-parameters and configurations

Table 8 lists the recommended FM hyper-parameters for each model scale used in our experiments.

	4B-base	7B-instr.	32B
λ_0	0.2	0.1	0.1
learning rate	3.0e-5	3.0e-5	1.0e-5
global batch size	128	128	128
epochs	2	4	2
schedule	cosine	cosine	cosine
warmup steps	200	400	200
T_ℓ (loss ramp)	300	300	300
suffix k	32	64	32
stride s	1	4	2
EMA μ	0.999	0.999	0.999
grad gating	$\gamma=1.2$	same	same
inserted index	-1	-1	-1
fm head width	1024	2048	2048
fm head depth	3	3	5
target encoder	BGE-Large	same	same

Table 8: Recommended FM configurations by scale.

E Design space table

Table 9 situates our method within the broader design space of future-target supervision approaches.

Method	space	Source p_0	Horizon
Next-token CE	Δ	—	1
MTP-CE	Δ^K	δ_π	K
MTP-Brier	Δ^K	δ_π	K
Cosine align.	$\mathbb{R}^{d_z}$	δ_0	full
MSE align.	$\mathbb{R}^{d_z}$	δ_0	full
FM (ours)	$\mathbb{R}^{d_z}$	$\text{Unif}(\mathbb{S}^{d_z-1})$	suffix- k

Table 9: Future-target flow design space. Existing methods occupy restricted corners; ours fills the remaining cell.

Reading the table. Moving down a row generally enlarges the function class. Moving from “ δ_π ” to “ $\text{Unif}(\mathbb{S}^{d_z-1})$ ” is what unlocks multi-modal target distributions (Proposition 3.2). Moving from “Bregman” to “Euclidean” on the same row preserves the optimum but changes the gradient geometry (Theorem A.4). Moving from “ Δ^K ” to “ $\mathbb{R}^{d_z}$ ” is what lets a single auxiliary head encode segment-level rather than token-level future information (Corollary A.5).

F Training-time pseudocode

The complete training procedure is given in Algorithm 1.

Algorithm 1: AR + Flow Matching joint training (one step).

Input: Backbone F_θ , FM head v_ϕ (zero-init output proj.), frozen target encoder E , dataset $\mathcal{D}$, suffix length k , stride s , schedule parameters $(\lambda_0, T_h, T_\ell, T_{\max})$, EMA decay μ , gating EMA μ_g , gating sensitivity γ .

```

1 for step  $n = 1, \dots, T_{\max}$  do
2    $(x, y) \sim \mathcal{D}$ ;
3    $h_{1:N} \leftarrow F_\theta(x \parallel y_{<N})$ ; // causal
4    $\mathcal{L}_{\text{AR}} \leftarrow \sum_i \text{CE}(\text{softmax}(W_{\text{AR}} h_i), y_{i+1})$ ;
   // Phase A ( $n < T_h$ ): only  $v_\phi$  trains; backbone frozen for FM
5   if  $n < T_h$  then
6      $h^{\text{fm}} \leftarrow \text{detach}(h)$ 
7   end
8   else
9      $h^{\text{fm}} \leftarrow h$ 
10  end
11  Compute anchor positions on response w/ stride  $s$ ;
12   $z_1^{(i)} \leftarrow E(y_{i+1:i+k})$  at each anchor; nearest-anchor fill;
13  Sample  $z_0^{(i)} \sim \text{Unif}(\mathbb{S}^{d_z-1})$ ,  $\tau \sim \mathcal{U}[0, 1]$ ;
14   $z_\tau \leftarrow (1 - \tau)z_0 + \tau z_1$ ;
15   $\hat{v} \leftarrow v_\phi(z_\tau, \tau; h^{\text{fm}})$ ;
16   $\ell_i \leftarrow \|\hat{v} - (z_1 - z_0)\|_2^2$  (mean over  $d_z$ );
17   $\widetilde{\mathcal{L}}_{\text{SFR}} \leftarrow \frac{1}{|A|} \sum_{i \in A} \ell_i$ ;
   //  $\lambda$  schedule: Phase A trains  $v_\phi$  only (backbone detached)
18  if  $n < T_h$  then
19     $\lambda(n) \leftarrow \lambda_0$ ; // gradient only reaches  $v_\phi$ 
20  else
21     $\lambda(n) \leftarrow \lambda_0 \cdot \min(1, (n - T_h)/T_\ell)$ ; // Phase B: ramp into backbone
22  end
   // Adaptive grad gating
23  Update  $\bar{L} \leftarrow \mu_g \bar{L} + (1 - \mu_g) \widetilde{\mathcal{L}}_{\text{SFR}}$ ;
24  if  $\widetilde{\mathcal{L}}_{\text{SFR}} < \bar{L}/\gamma$  (synchronized via all_reduce MAX) then
25     $h^{\text{fm}} \leftarrow \text{detach}(h^{\text{fm}})$ ; recompute  $\widetilde{\mathcal{L}}_{\text{SFR}}$ ;
26  end
27  if  $n < T_h$  then
28     $\mathcal{L}_{\text{total}} \leftarrow \lambda(n) \widetilde{\mathcal{L}}_{\text{SFR}}$ ; // Phase A: no  $\mathcal{L}_{\text{AR}}$ , only  $v_\phi$  updates
29  else
30     $\mathcal{L}_{\text{total}} \leftarrow \mathcal{L}_{\text{AR}} + \lambda(n) \widetilde{\mathcal{L}}_{\text{SFR}}$ ;
31  end
32  Backprop  $\mathcal{L}_{\text{total}}$ ; optimizer step on  $(\theta, \phi)$ ;
   // EMA teacher: only after Phase A
33  if  $n \geq T_h$  then
34     $\phi^{(\text{ema})} \leftarrow \mu \phi^{(\text{ema})} + (1 - \mu) \phi$ ;
35     $\phi \leftarrow \phi^{(\text{ema})}$ ;
36  end
37 end

```

G Single-factor ablation on Qwen3-4B-Base / MBPP

All ablations below are evaluated on MBPP (Base) and MBPP+ (Plus) at $T=0.6$, reporting pass@1, pass@5, pass@10. Unless otherwise noted, the default configuration for the suffix and stride sweeps is $\lambda=0.1$, suffix tokens $k=32$, stride $s=1$. The main-text result (Table 5) uses the best configuration $\lambda=0.2$. SFT baseline (no FM): 77.2 / 67.0 / 85.6 / 74.4 / 88.1 / 76.7 on Base/Plus pass@1/5/10.

λ	pass@1		pass@5		pass@10	
	Base	Plus	Base	Plus	Base	Plus
0.01	77.9	66.5	87.0	75.3	89.5	78.0
0.05	75.5	64.7	84.4	74.3	86.8	77.1
0.1	75.6	65.4	84.8	73.7	87.9	76.8
0.2	78.3	67.7	87.4	76.8	89.2	79.0
0.5	77.1	65.8	85.8	74.3	87.7	76.1
1.0	76.7	66.3	85.7	74.7	88.6	78.0

Table 10: Loss-weight λ sweep (fixed $k=32$, $s=1$).

k	pass@1		pass@5		pass@10	
	Base	Plus	Base	Plus	Base	Plus
2	73.8	63.6	83.8	73.3	86.5	76.2
4	76.3	65.3	86.1	74.6	88.7	77.1
8	75.3	64.9	84.7	74.4	87.3	77.0
16	76.2	65.7	85.5	74.8	87.9	77.2
32	75.6	65.4	84.8	73.7	87.9	76.8

Table 11: Suffix-token horizon k sweep (fixed $\lambda=0.1$, $s=1$).

s	pass@1		pass@5		pass@10		Time
	Base	Plus	Base	Plus	Base	Plus	
1	75.6	65.4	84.8	73.7	87.9	76.8	20h40m
2	77.1	66.0	86.0	75.2	88.5	77.7	18h24m
4	75.9	65.0	84.0	73.7	86.5	76.6	16h56m
16	75.9	65.4	85.8	75.3	88.4	78.3	15h47m

Table 12: Stride s sweep (fixed $\lambda=0.1$, $k=32$). Training wall-clock time is also reported.

The λ sweep (Table 10) shows that $\lambda=0.2$ performs best overall; larger values ($\lambda \in [0.5, 1.0]$) still improve over SFT on most metrics, but very small or mismatched λ (e.g. 0.05, 0.1) can underperform SFT on individual splits, indicating that the FM gradient must be sufficiently strong to reshape the backbone. The suffix-token sweep (Table 11) shows that $k=4$ already recovers much of the high- k pass@5/pass@10 gain, although pass@1 remains below SFT at small horizons; diminishing returns set in beyond $k=16$. The stride

sweep (Table 12) confirms that quality is roughly flat from $s=1$ to $s=16$ while training time drops from 20h40m to 15h47m (compared with SFT at 14h43m), making $s=16$ a strong efficiency–quality trade-off that adds roughly $\sim 10\%$ over vanilla SFT.

G.1 FM versus deterministic alignment

Propositions 3.1 and 3.2 predict that replacing the stochastic FM source with a deterministic regression target (cosine or MSE alignment to z_1 , i.e. the same auxiliary head and target but with z_0 fixed) should collapse the auxiliary objective toward the conditional mean and erase most of SFR’s gain. We test this directly under the exact matched setting of Tables 10 to 12 (Qwen3-4B-Base, same head, same λ, k, s), swapping only the loss function.

Loss	pass@1		pass@5		pass@10	
	Base	Plus	Base	Plus	Base	Plus
Cosine	75.6	65.5	85.3	74.5	87.5	76.8
MSE	76.2	64.9	85.0	73.6	87.5	76.4
FM (SFR)	78.3	67.7	87.4	76.8	89.2	79.0

Table 13: Same auxiliary head and target, loss swapped (Qwen3-4B-Base, pass@ k on MBPP Base/Plus). 95% CIs (paired bootstrap) of each variant against SFR are non-overlapping with zero at 10 of 12 entries.

Neither deterministic variant approaches SFR at any k on either split (95% CIs of Cosine/MSE – SFR exclude zero at 10 of the 12 entries, with the two exceptions at pass@10 where the overall margin is smallest), confirming that the gain is attributable to the flow-matching objective itself—specifically its stochastic source—rather than to the auxiliary head’s added capacity.

H Per-query Cross-Style Self-BLEU analysis

Rather than listing all 148 queries, we highlight the queries where SFR most outperforms SFT in terms of relative CS-SB reduction, and discuss the qualitative pattern.

Pattern. The queries where SFR yields the largest diversity gains share three characteristics: (i) they are *open-ended* and admit many structurally distinct valid answers; (ii) they require *opinion or analytical reasoning* rather than factual recall; and (iii) they span diverse domains (relationships, culture, medicine, management), giving each persona ample room to exhibit a distinctive voice.

#	Query (de-identified, with English gloss)	Δ
1	在一段真正健康的亲密关系中，最重要的情感支持应该具备哪些特征？ <i>What should the most important emotional support look like in a truly healthy intimate relationship?</i>	33.4%
2	青年在传承历史文脉和推动中华文化创新发展过程中，如何将航天领域取得的重要成就与赓续历史文脉相结合？ <i>How can young people link major aerospace achievements to carrying forward historical and cultural heritage while advancing Chinese culture?</i>	23.7%
3	某些历史科学研究成果是如何被曲解并影响公众对特定生物学特征的认知的？ <i>How have certain historical scientific findings been misinterpreted, shaping public perception of specific biological traits?</i>	22.2%
4	某作者在文中指出左翼媒体生态发生了哪些关键性转变？ <i>What key shifts does the author identify in the left-leaning media ecosystem?</i>	19.3%
5	某品牌最新400cc摩托车系列在2023年的市场策略对消费者选购有何促进作用？ <i>How did a brand's 2023 marketing strategy for its newest 400cc motorcycle series drive consumer purchases?</i>	19.2%
6	九十年代中期香港电影产业在类型片发展上有何特点？ <i>What characterized genre-film development in Hong Kong cinema in the mid-1990s?</i>	18.1%
7	被匿名水军攻击时如何固定符合诉讼标准的证据？ <i>How can one preserve litigation-grade evidence when targeted by anonymous coordinated online attacks?</i>	16.5%
8	从管理学效率的角度看，如何通过优化资源分配和精力管理来实现在相同工作时间内产出更多价值？ <i>From a management-efficiency view, how can optimizing resource allocation and energy management yield more value in the same working hours?</i>	16.3%
9	游客在前往某些热门旅游目的地时可能会遇到哪些需要特别警惕的情况？ <i>What situations should travelers be especially wary of at certain popular destinations?</i>	15.5%
10	在临床实践中如何通过呼气峰流量和呼气时间常数的异常值来区分不同病理状态？ <i>In clinical practice, how can abnormal peak expiratory flow and expiratory time-constant values distinguish pathological states?</i>	15.2%

Table 14: Top-10 queries by average relative CS-SB reduction of SFR over SFT. “ Δ ” is the mean relative improvement across CS-SB₁₋₄. Queries shown are de-identified excerpts from the proprietary platform; the English gloss is placed under the Chinese original in the same cell.

Conversely, the queries where SFT matches or slightly outperforms SFR tend to be *closed-form* or *factual* (e.g. “Which documents must be submitted for drug procurement?”, “Which year is my Hong-Luan star year if my zodiac is Horse?”). In such cases the answer space is inherently narrow, leaving little room for stylistic diversification regardless of method.

This confirms that SFR’s multi-modal preservation mechanism is most impactful precisely where Cross-Style Collapse is most severe: open, subjective queries whose valid response space is large and style-sensitive.

I Full per-style LLM-judge scores

Table 2 in the main text reports four representative styles (style-0–3) and the nine-style average. Table 15 below gives the complete breakdown for all nine stylistic personas.

	R1-prompt	SFT	SFR
style-0 Ctx	4.5	4.568	4.622
style-0 Rel	4.662	4.764	4.77
style-0 SS	4.654	4.588	4.905
style-0 Flu	4.838	4.831	4.905
style-1 Ctx	4.175	4.216	4.257
style-1 Rel	4.562	4.581	4.628
style-1 SS	3.117	2.858	3.135
style-1 Flu	4.781	4.838	4.791
style-2 Ctx	4.117	4.169	4.128
style-2 Rel	4.285	4.311	4.223
style-2 SS	4.766	4.73	4.824
style-2 Flu	4.964	4.98	4.993
style-3 Ctx	4.511	4.568	4.696
style-3 Rel	4.759	4.818	4.851
style-3 SS	3.891	3.932	4.73
style-3 Flu	4.796	4.831	4.98
style-4 Ctx	4.182	4.297	4.372
style-4 Rel	4.416	4.527	4.534
style-4 SS	4.657	4.689	4.662
style-4 Flu	4.942	4.919	4.926
style-5 Ctx	4.456	4.52	4.48
style-5 Rel	4.699	4.764	4.743
style-5 SS	4.132	3.932	4.027
style-5 Flu	4.64	4.757	4.797
style-6 Ctx	4.387	4.419	4.514
style-6 Rel	4.628	4.703	4.777
style-6 SS	4.146	3.75	3.709
style-6 Flu	4.788	4.797	4.791
style-7 Ctx	4.522	4.554	4.541
style-7 Rel	4.706	4.73	4.736
style-7 SS	4.904	4.986	4.905
style-7 Flu	4.794	4.912	4.885
style-8 Ctx	4.467	4.534	4.561
style-8 Rel	4.723	4.757	4.73
style-8 SS	4.701	4.723	4.709
style-8 Flu	4.752	4.878	4.858
Average Ctx	4.368	4.427	4.463
Average Rel	4.604	4.661	4.666
Average SS	4.329	4.243	4.401
Average Flu	4.811	4.86	4.881

Table 15: Full stylized-dialogue LLM-judge scores at $T = 0.8$ on Qwen3-32B across all nine styles. Bold indicates the best score among the three methods.

J Use of AI Assistants

We used AI-based tools during the preparation of this work in the following capacities:

- **Code debugging:** AI coding assistants were used to identify and fix implementation bugs in the training pipeline (vibe coding / interactive debugging).
- **Writing assistance:** AI writing tools were used for translating between Chinese and English drafts and for polishing prose. All AI-generated text was reviewed and revised by the authors.

No AI tool was used to generate experimental results, training data, or core scientific content.

K Baseline, significance, judge, and reproducibility details

This appendix provides the full experimental setup and detailed tables underlying the robustness results summarized in Sections 4.1 and 4.3 (Table 1 and the MBPP significance intervals in the main text): additional diversity-enhancing baselines, significance-testing methodology, the cross-family judge protocol, and the public PersonaChat reproduction. Peak-memory measurements are also reported.

K.1 Additional baselines: SimCTG and Contrastive Decoding

Both baselines are evaluated on the same Qwen3-32B stylized-dialogue setup, with the identical 1,332 evaluation prompts, the same four LLM-judge dimensions (Ctx / Rel / SS / Flu), and CS-SB.

SimCTG (training-time). We add the contrastive token-representation loss $\mathcal{L}_{CL}$ (margin $\rho=0.5$, computed over the model’s own token hidden states, with no external encoder) to the SFT objective, i.e. $\mathcal{L} = \mathcal{L}_{MLE} + \mathcal{L}_{CL}$; all other settings match our main Qwen3-32B run.

Contrastive Decoding (CD, decoding-time). The expert is the untuned Qwen3-32B and the amateur is the same-family Qwen3-4B. At each step we build an adaptive plausibility set from the expert distribution (threshold $\alpha=0.1$), then sample within it by the contrastive score $(1+\beta) \log p_{\text{exp}} - \beta \log p_{\text{ama}}$ ($\beta=0.5$, $T=0.6$, $\text{top-}p=0.95$), using the same prompts and sampling hyper-parameters as the style-tuned SFT model.

Table 16 reports the average LLM-judge scores for both baselines, alongside R1-prompt, SFT, and

	R1-p.	SFT	SimCTG	CD	SFR
Ctx	4.368	4.427	4.150	4.188	4.463
Rel	4.604	4.661	4.420	4.504	4.666
SS	4.329	4.243	4.381	3.748	4.401
Flu	4.811	4.860	4.811	4.660	4.881

Table 16: Average LLM-judge scores across the nine styles under the primary judge (Qwen3-235B), including the two additional baselines.

SFR, under our primary judge (Qwen3-235B); CS-SB results are in Table 1 in the main text. Neither baseline matches SFR: SimCTG’s scores stay close to SFT’s, indicating that its contrastive representation loss does not translate into stronger style-discriminative behavior; CD attains a competitive CS-SB (Table 1) but an abnormally low style-strength score (3.748, versus 4.401 for SFR), indicating that its diversity gain is a surface-level decoding artifact rather than genuine style-semantic diversification.

K.2 Significance testing

We compute paired-bootstrap 95% confidence intervals (10,000 resamples) on our existing per-example results, resampling queries (for CS-SB, 148 paired queries on Qwen3-32B) or problems (for MBPP pass@ k on Qwen3-4B-Base) with replacement and recomputing the metric on each resample. Results are reported directly in Table 1 and in §4.3.

K.3 Cross-family judge validation

To address the concern that scoring Qwen3-32B outputs with a same-family judge (Qwen3-235B) is circular, we re-score all dialogue outputs with a cross-family judge, **GLM-5.2**, and compare the two judges via a difference-in-differences (DiD) analysis over the $9 \text{ styles} \times 4 \text{ dimensions} = 36$ paired cells.

	R1-prompt	SFT	SFR
Ctx	4.113	4.105	4.188
Rel	4.509	4.497	4.528
SS	4.182	4.094	4.303
Flu	4.915	4.968	4.970

Table 17: Average scores under the cross-family validation judge (GLM-5.2), re-scoring the same outputs as Table 16.

SFR is rated best on every dimension under both judges. A DiD permutation test and Wilcoxon signed-rank test on each (dimension, model pair)

find one real same-family bias—on Ctx and Rel, Qwen gives SFT a significant boost relative to R1-prompt ($p < 0.05$)—but every cell involving SFR vs. SFT (our core contribution) is judged fair by both tests, so our central claim does not depend on judge family. Qwen3-235B is also systematically 0.10–0.15 points more lenient than GLM-5.2 across the board, but this constant offset cancels out in any cross-method comparison.

K.4 Public-data validation: PersonaChat

Following a reviewer’s suggestion, we fine-tune Qwen3-4B on the open bavard/personachat_truecased dataset (Zhang et al., 2018) to assess persona-controlled generation independently of our proprietary corpus. We split by dialogue to eliminate persona/history leakage (training on the official train split, 65,719 samples after expanding each utterance and cleaning __SILENCE__ placeholders; evaluation on the valid split). The evaluation set keeps only the inter-style axis: 9 personas $\times$ 100 neutral contexts = 900 responses, with one response per (context, persona) pair, and CS-SB computed per context (English tokenization, $T=0.6$, thinking disabled). Results are given in Table 4 in the main text.

K.5 Training-time memory overhead

Beyond the training-time overhead already reported in Table 12 ($\sim 10\%$ wall-clock at stride $s=1$, dropping to $\sim 8\%$ at $s=16$), we measure peak GPU memory for SFT vs. SFR at three scales, each at its own default parallelism and context length.

Scale	SFT	SFR	Increase
4B (ZeRO-1, ctx 8192)	793.6	811.9	+2.3%
7B (ZeRO-1, ctx 16384)	2423.3	2532.7	+4.5%
32B (ZeRO-3, ctx 8192)	2171.1	2362.8	+8.8%

Table 18: Peak GPU memory (GB), SFT vs. SFR.

The increase stays in the single-to-low-double-digit percentage range. Most of it comes from the frozen target encoder rather than the FM head itself, which is a lightweight MLP ($\sim 10^{-3} \times$ the backbone’s parameter count); in the token-level suffix mode this encoder is replicated outside the ZeRO-3 scope on every GPU, so its footprint grows linearly with GPU count. For settings where pre-processing is affordable (e.g. our Qwen3-4B bridge experiments), we instead precompute all target embeddings once before training and drop the encoder entirely, avoiding this cost.