

000 VGPO: FINE-TUNING SPEECH AUTOREGRESSIVE 001 DIFFUSION MODELS WITH VALUE GUIDED POLICY 002 OPTIMIZATION

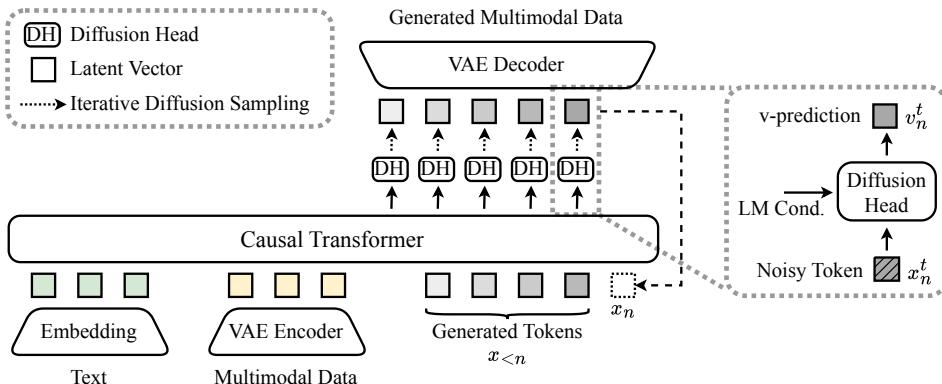
003 **Anonymous authors**

004 Paper under double-blind review

010 ABSTRACT

011 Autoregressive diffusion models (ARDMs), which generate continuous latent se-
012 quences, have recently achieved state-of-the-art zero-shot text-to-speech (TTS)
013 performance. However, fine-tuning these models with reinforcement learn-
014 ing (RL) to directly optimize user-defined reward functions remains an open chal-
015 lenge. In this work, we propose Value-Guided Policy Optimization (VGPO), an
016 actor-critic RL algorithm tailored to ARDMs. We train a causal value model to
017 predict expected future rewards and update the ARDM using gradients from this
018 value model. To validate VGPO, we fine-tune the recently introduced DiTAR
019 model and evaluate it on two tasks: improving F0 variance to enhance expres-
020 siveness; and optimizing text log-probability to improve the model’s robustness
021 to challenging long text. VGPO can achieve significant improvement in zero-shot
022 TTS expressiveness and robustness, while maintaining naturalness and speaker
023 similarity.

026 1 INTRODUCTION



042 Figure 1: Illustration of a widely adopted autoregressive diffusion model architecture. It is used in
043 state-of-the-art models for speech (Sun et al., 2024; Jia et al., 2025), and image generation (Kou
044 et al., 2025; Sun et al., 2024; NextStep Team et al., 2025). A causal transformer encodes the gen-
045 eration conditions and previously generated tokens, while a diffusion head predicts the next token x_n
046 via diffusion modeling. This architecture is highly efficient, as only the diffusion head is repeatedly
047 evaluated when generating each next token.

048 Autoregressive diffusion models (ARDMs) represent continuous modalities using latent vectors
049 (continuous tokens) and generate data sequentially by predicting the next token through diffusion
050 modeling. This approach has been adopted in an increasing number multimodal generation models,
051 including audio (Liu et al., 2024; Jia et al., 2025; Yang et al., 2025), image (Li et al., 2024a; Hu et al.,
052 2024), and video generation (Yin et al., 2025; Deng et al., 2024). Unlike the approach of tokenizing
053 data into discrete symbols for next-token prediction, ARDMs offer two key advantages with next-
token diffusion (Sun et al., 2024): they preserve fine-grained details while avoiding prohibitively

long sequences, thanks to the compactness of latent continuous representations. This allows for more precise control over the generated content, yielding superior performance on tasks that require high-fidelity details. In particular, several recent works have introduced ARDMs for speech generation (Jia et al., 2025; Sun et al., 2024), achieving state-of-the-art performance in zero-shot text-to-speech.

Reinforcement fine-tuning (RFT) (Ouyang et al., 2022; Rafailov et al., 2023; Xu et al., 2023) is a key stage in the post-training of multimodal generative models. It improves model performance by directly optimizing the expected reward of generated samples under a user-specified reward function $r(\mathbf{x})$. Given a pre-trained reference model $\mu(\mathbf{x})$, RFT optimizes a policy $\rho(\mathbf{x})$ to maximize the expected reward while remaining close to the reference via a divergence regularizer $d(\cdot\|\cdot)$, which helps mitigate reward hacking (Weng, 2024):

$$\max_{\rho} \mathbb{E}_{\mathbf{x} \sim \rho(\mathbf{x})} [r(\mathbf{x})] - d(\rho\|\mu) \quad (1)$$

For speech generation, prior work has instantiated $r(\mathbf{x})$ with a variety of reward models, including predicted mean opinion score (MOS) (Chen et al., 2025; 2024; Tian et al., 2025), speaker similarity from speaker encoders (Du et al., 2025; Li et al., 2025), emotion classification accuracy (Anastassiou et al., 2024), and text accuracy measured by automatic speech recognition (ASR) (Zhang et al., 2025). RFT with such rewards has been shown to improve intelligibility, naturalness, speaker identity preservation, and controllability of speech generative models.

Although recent work has extensively explored the network architectures and training techniques of ARDMs, relatively little research has investigated RFT algorithms for these models (NextStep Team et al., 2025; Liu et al., 2025b), leaving this area largely unexplored. In this work, we introduce **Value-Guided Policy Optimization** (VGPO), a novel actor–critic RFT algorithm tailored for ARDMs.

The core component of VGPO is a learned (soft) value function, initialized from a pre-trained ARDM. Given the previously generated tokens and the current partially denoised token, it predicts the expected final reward. Extending prior results in maximum-entropy (MaxEnt) RL (Levine, 2018) and exact energy guidance (Lu et al., 2023) to ARDMs, we prove that the optimal ARDM prediction equals the sum of the reference model’s prediction and the gradient of the value function. To train the model, VGPO samples trajectories online and regresses the model’s intermediate predictions toward this theoretical optimum.

The contributions of this paper are as follows:

- We propose VGPO, a value-based reinforcement fine-tuning algorithm specifically designed for ARDMs. VGPO trains a value model to predict the total reward from partial trajectories and updates the diffusion score using the value model gradient.
- We apply VGPO to fine-tune DiTAR (Jia et al., 2025), a recently proposed state-of-the-art zero-shot TTS model based on ARDMs. We evaluate VGPO on two benchmark tasks: Task A, which aims to enhance the expressiveness of generated speech by optimizing the F0 variance (Quatieri, 2002). And task B, which focuses on improving TTS robustness when handling challenging long texts that are difficult for autoregressive TTS models.
- We propose to regularize VGPO through adversarial distribution matching (ADM). Unlike the common KL regularization (Fan et al., 2023), which struggles to correct errors accumulated during reinforcement fine-tuning, ADM effectively mitigates this issue.

The audio samples are available at <https://vgpo-web.github.io/>.

2 PRELIMINARIES

In this section, we describe the related background formulations necessary for describing our algorithm, including formulations of diffusion models (DM) and ARDMs.

2.1 NOTATIONS

Without loss of generality, we omit input conditions \mathbf{c} , such as text or prompt speech, and we assume that the generated token sequence has a fixed length N . Each sample trajectory $\mathbf{x} \in \mathbb{R}^{N \times d}$

108 consists of N tokens (x_1, \dots, x_N) , where each token resides in \mathbb{R}^d . We define $x_{\leq n}$ as the sequence
 109 (x_1, \dots, x_n) , with similar definitions for $x_{<n}$, $x_{>n}$, and $x_{\geq n}$.
 110

111 To prevent ambiguity between the diffusion time index, t , and the sequence index, n , we use a
 112 superscript on variables to denote the diffusion time when both t and n appear simultaneously
 113 (e.g., x_n^t). For conciseness, all t -conditioned models are written without the explicit time argument:
 114 $g(x_n^t) := g(x_n^t, t)$. This convention applies to all score models and value models.
 115

2.2 AUTOREGRESSIVE DIFFUSION MODELS

117 **Diffusion Models.** One view that unifies many diffusion model formulations (Song et al., 2021b;
 118 Liu et al., 2023) is to view DMs as multiscale score estimators. For each diffusion time t , define a
 119 Gaussian transition distribution $q(x_t|x) := \mathcal{N}(x_t; \alpha_t x, \sigma_t^2 I_d)$, where $\alpha_t, \sigma_t > 0$. Suppose $p(x)$ is
 120 the clean token distribution. Then the marginal distribution of noisy tokens at time t is:
 121

$$p(x_t) := \int p(x)q(x_t|x)dx. \quad (2)$$

124 A diffusion model trained with denoising score matching can be interpreted as a score estimator
 125 $s_\theta(x_t)$ that approximates the true score $\nabla \log p(x_t)$. Given the true score of all diffusion time t ,
 126 one can draw samples from $p(x)$ with various diffusion model samplers such as DDPM (Ho et al.,
 127 2020), DDIM (Song et al., 2021a), SDE and ODE solvers (Song et al., 2021b).

128 **Autoregressive Models.** Let $p(\mathbf{x})$ denote the data distribution. By the chain rule of probability,
 129

$$p(\mathbf{x}) = \prod_{n=1}^N p(x_n|x_{<n}), \quad (3)$$

133 An autoregressive (AR) model parameterizes these conditionals and approximates $p(x_n | x_{<n})$.
 134 Sampling proceeds ancestrally: first draw $x_1 \sim p(x_1)$, then $x_2 \sim p(x_2 | x_1)$, and, in general,
 135 $x_n \sim p(x_n | x_{<n})$ until all N tokens are generated.

136 **Autoregressive Diffusion Models.** ARDMs sample from each conditional $p(x_n|x_{<n})$ with a diffu-
 137 sion model. Let $q(x_n^t|x_n) := \mathcal{N}(x_n^t; \alpha_t x_n, \sigma_t^2 I_d)$. For each diffusion time t , define the conditional
 138 marginal distribution $p(x_n^t|x_{<n})$ as:
 139

$$p(x_n^t|x_{<n}) := \int p(x_n|x_{<n})q(x_n^t|x_n)dx_n \quad (4)$$

142 An ARDM learns a conditional multiscale score estimator $s_\theta(x_n^t|x_{<n})$ that estimates the conditional
 143 score $\nabla_{x_n^t} \log p(x_n^t|x_{<n})$, given clean tokens generated previously $x_{<n}$, and the token currently
 144 being denoised x_n^t . Let $\pi_\theta(\mathbf{x})$ be the sample distribution of the ARDM with the score model s_θ .
 145 Assuming $s_\theta(x_n^t|x_{<n}) = \nabla_{x_n^t} \log p(x_n^t|x_{<n})$, and that the diffusion sampler is an exact solver, we
 146 have $\pi_\theta(\mathbf{x}) = p(\mathbf{x})$.
 147

2.3 KL REGULARIZED REWARD MAXIMIZATION

150 **Definition 1 (KL RFT).** Suppose that the distribution of the reference model is $\mu(\mathbf{x})$ and the reward
 151 function is $r(\mathbf{x}) : \mathbb{R}^{N \times d} \rightarrow \mathbb{R}$. Suppose that we pick the Kullback–Leibler (KL) divergence as our
 152 divergence regularizer, the goal of RFT now becomes the following:
 153

$$\max_{\rho} \mathbb{E}_{\mathbf{x} \sim \rho(\mathbf{x})} [r(\mathbf{x})] - D_{\text{KL}} (\rho(\mathbf{x}) \parallel \mu(\mathbf{x})). \quad (5)$$

155 **Theorem 1** (Solution of KL RFT). The closed form solution π for the KL RFT problem in Defini-
 156 tion (1) is given by:
 157

$$\pi(\mathbf{x}) = \frac{\mu(\mathbf{x}) \exp r(\mathbf{x})}{Z}, \quad Z = \int \mu(\mathbf{x}) \exp r(\mathbf{x}) d\mathbf{x}. \quad (6)$$

160 *Proof.* The proof can be found in previous works (Rafailov et al., 2023; Peng et al., 2019). A short
 161 proof is provided in Appendix A.1. \square

162

3 METHODS

163

164

3.1 VALUE GUIDANCE FOR ARDMs

165

166 **Definition 2** (Soft Value Functions of ARDMs). Given a reward function $r(\mathbf{x}) : \mathbb{R}^{N \times d} \rightarrow \mathbb{R}$ and
167 an ARDM with distribution $\mu(\mathbf{x})$. Soft value function $V(\cdot)$ estimates the future reward given partial
168 information of the complete sample trajectory.

169
$$V(x_{\leq n}) := \log(\mathbb{E}_{\mu(x_{>n}|x_{\leq n})} [\exp r(\mathbf{x})]); \quad (7)$$
170

171
$$V(x_{<n}, x_n^t) := \log(\mathbb{E}_{\mu(x_n|x_n^t, x_{<n})} [\exp V(x_{\leq n})]) = \log(\mathbb{E}_{\mu(x_{\geq n}|x_n^t, x_{<n})} [\exp r(\mathbf{x})]). \quad (8)$$
172

173 **Theorem 2** (Solution of KL RFT for ARDMs). Given a reward function $r(\mathbf{x}) : \mathbb{R}^{N \times d} \rightarrow \mathbb{R}$ and a
174 reference ARDM with distribution $\mu(\mathbf{x})$. Suppose $\pi(\mathbf{x}) \propto \mu(\mathbf{x}) \exp r(\mathbf{x})$ is the optimal solution of
175 the KL RFT problem in Eq. (5). Then the conditional distribution $\pi(x_n^t|x_{<n})$ is given by:

176
$$\pi(x_n^t|x_{<n}) = \mu(x_n^t|x_{<n}) \exp(V(x_{<n}, x_n^t) - V(x_{<n})). \quad (9)$$
177

178 where the soft value function $V(\cdot)$ is defined in Definition 2.

179 *Proof.* We provide two proofs for this result in A.2 and A.3. Proof in A.2 is based on directly
180 solving for $\pi(x_n^t|x_{<n})$. We can view the ARDM sampling process as a Markov decision process
181 (MDP), and apply MaxEnt RL (Levine, 2018) to the MDP. This leads to the proof in A.3. We show
182 that $V(\cdot)$ is indeed the soft value function of the optimal policy. \square

183 **Corollary 1** (Value Guidance). By applying $\nabla_{x_n^t} \log$ on both sides of Eq. (9), we can show that
184 the conditional score of the optimal policy $\pi(x_n^t|x_{<n})$ is the sum of the reference score and value
185 gradient:

186
$$\nabla_{x_n^t} \log \pi(x_n^t|x_{<n}) = \nabla_{x_n^t} \log \mu(x_n^t|x_{<n}) + \nabla_{x_n^t} V(x_{<n}, x_n^t). \quad (10)$$
187

188

3.2 ESTIMATING THE VALUE FUNCTION

189

190 Apply \exp on both sides of Equation (8) gives:

191
$$\exp V(x_{<n}, x_n^t) = \mathbb{E}_{\mu(x_{\geq n}|x_{<n}, x_n^t)} [\exp r(\mathbf{x})]. \quad (11)$$
192

193 Therefore, given a parameterized soft value model $V_\phi(x_{<n}, x_n^t)$, we can approximate $V(x_{<n}, x_n^t)$
194 by minimizing the following Exp-MSE loss(Lu et al., 2023; Uehara et al., 2024):

195
$$\mathcal{L}_V^{n,t}(\phi) := \mathbb{E}_{\mu(\mathbf{x}), q(x_n^t|x_n)} [\exp V_\phi(x_{<n}, x_n^t) - \exp r(\mathbf{x})]^2. \quad (12)$$
196

197 The Exp-MSE loss in Eq. (12) is numerically unstable (Lu et al., 2023) due to the $\exp(\cdot)$ functions.
198 We propose an alternative loss $\tilde{\mathcal{L}}_V^{n,t}(\phi)$ that shares the same global minimum, while providing better
199 numerical stability. We leave the analysis in A.4. Previous works (Li et al., 2024b; Lu et al., 2023)
200 observed that replacing the *soft* value function in Eq. (10) with the value function also provides good
201 results, at the cost of losing the theoretical distribution guarantee provided by Theorem 1. We can
202 minimize the following MSE loss to learn a value model $V_\phi(\cdot)$ approximating the value function.

203
$$\tilde{\mathcal{L}}_V^{n,t}(\phi) := \mathbb{E}_{\mu(\mathbf{x}), q(x_n^t|x_n)} [V_\phi(x_{<n}, x_n^t) - r(\mathbf{x})]^2. \quad (13)$$
204

205 Algorithm 1 learns a (soft) value
206 model given a reference ARDM
207 model. As a result of Corollary 1,
208 the value model can be used to guide
209 ARDM sampling via Algorithm 2.
210 Note that Algorithm 2 scales the
211 value gradient by the parameter λ .
212 Similar to classifier guidance (Ho &
213 Salimans, 2022), we observe that setting
214 $\lambda > 1$ can further improve the
215 rewards of the samples.

Algorithm 1 Value Training

Require: Reference ARDM model with distribution $\mu(\mathbf{x})$;
initialized value model $V_\phi(\cdot)$

1: **while** V_ϕ has not converged **do**
2: Sample a trajectory $\mathbf{x} \sim \mu(\mathbf{x})$
3: Compute reward $r(\mathbf{x})$
4: Update ϕ by minimizing the loss $\mathbb{E}_{n,t} [\tilde{\mathcal{L}}_V^{n,t}(\phi)]$ or
 $\mathbb{E}_{n,t} [\tilde{\mathcal{L}}_V^{n,t}(\phi)]$
5: **end while**

216 **Algorithm 2** Value Guided Sampling
217
218 **Require:** Reference ARDM model with distribution $\mu(\mathbf{x})$ and score estimator $s_\mu(x_n^t | x_{<n})$; pre-
219 trained value model $V_\phi(\cdot)$; guidance scale $\lambda \in \mathbb{R}$
220 1: **for** $n = 1$ to N **do**
221 2: Sample initial noise from $\mathcal{N}(0, I_d)$
222 3: Sample x_n by running diffusion sampler with score $s_\mu(x_n^t | x_{<n}) + \lambda \cdot \nabla_{x_n^t} V_\phi(x_{<n}, x_n^t)$ at t
223 4: **end for**

224
225 3.3 VALUE GUIDED POLICY OPTIMIZATION
226

227 Given an ARDM model with distribution $\pi_\theta(\mathbf{x})$ and score estimator $s_\theta(x_n^t | x_{<n})$, VGPO fine-tunes
228 its score prediction to match the optimal solution in Eq. (10). Suppose the reference model $\mu(\mathbf{x})$
229 has score estimator $s_\mu(x_n^t | x_{<n}) = \nabla_{x_n^t} \log \mu(x_n^t | x_{<n})$, and let $V_\phi(x_{<n}, x_n^t)$ be a trained value
230 model. The MSE loss $\mathbb{E}_{n,t} [\mathcal{L}_{\text{VD}}^{n,t}(\theta)]$ updates the prediction of s_θ to match the optimal solution
231 $\nabla_{x_n^t} \log \mu(x_n^t | x_{<n}) + \nabla_{x_n^t} V_\phi(x_{<n}, x_n^t)$.

232
$$\mathcal{L}_{\text{VD}}^{n,t}(\theta) := \mathbb{E}_{\pi_{\text{sg}[\theta]}(\mathbf{x}), q(x_n^t | x_n)} \|s_\theta(x_n^t | x_{<n}) - (s_\mu(x_n^t | x_{<n}) + \nabla_{x_n^t} V_\phi(x_{<n}, x_n^t))\|_2^2. \quad (14)$$

233 Note that the loss $\mathcal{L}_{\text{VD}}^{n,t}$ can be decomposed into two components that pull $s_\theta(x_n^t | x_{<n})$ in different
234 directions. The first component is a KL regularization term (Liu et al., 2025a):
235

236
$$\mathcal{L}_{\text{KL}}^{n,t}(\theta) := \mathbb{E}_{\pi_{\text{sg}[\theta]}(\mathbf{x}), q(x_n^t | x_n)} \|s_\theta(x_n^t | x_{<n}) - s_\mu(x_n^t | x_{<n})\|_2^2, \quad (15)$$

237 and the second is a value-guidance term:
238

239
$$\mathcal{L}_{\text{VG}}^{n,t}(\theta) := \mathbb{E}_{\pi_{\text{sg}[\theta]}(\mathbf{x}), q(x_n^t | x_n)} \|s_\theta(x_n^t | x_{<n}) - \text{sg}[s_\theta(x_n^t | x_{<n})] - \nabla_{x_n^t} V_\phi(x_{<n}, x_n^t)\|_2^2. \quad (16)$$

240 VGPO is described in Algorithm 3. In Algorithm 3 we multiply the term $\mathcal{L}_{\text{KL}}^{n,t}$ by w_{KL} to control
241 the KL regularization strength. Additionally, we can choose to update the value model online with
242 Algorithm 1, turning the algorithm into a variant of online policy mirror decent (Tomar et al., 2020;
243 Kimi Team et al., 2025; Ma et al., 2025).

244 **Algorithm 3** Value Guided Policy Optimization (VGPO)
245

246 **Require:** Reference ARDM model with distribution $\mu(\mathbf{x})$ and score estimator $s_\mu(x_n^t | x_{<n})$; pre-
247 trained value model $V_\phi(\cdot)$; KL loss weight w_{KL} ; target ARDM model $\pi_\theta(\mathbf{x})$ initialized from μ with
248 score estimator $s_\theta(x_n^t | x_{<n})$

249 **Optionally Require:** Discriminator $D_\psi(\cdot)$; weight of adversarial gradient w_A

250 1: **while** π_θ has not converged **do**
251 2: Sample a trajectory $\mathbf{x} \sim \pi_\theta(\mathbf{x})$
252 3: Update θ by minimizing $w_{\text{KL}} \mathcal{L}_{\text{KL}}^{n,t}(\theta) + \mathcal{L}_{\text{VG}}^{n,t}(\theta)$ on randomly sampled n, t pairs
253 4: **if** enabled online value update **then**
254 5: Invoke Algorithm 1 to sample from π_θ and update V_ϕ
255 6: **end if**
256 7: **if** enabled adversarial distribution matching (Section 3.4) **then**
257 8: Update θ by minimizing $w_A \mathcal{L}_{\text{G}}^{n,t}(\theta)$ on randomly sampled n, t pairs
258 9: Sample a trajectory $\mathbf{x}' \sim \mu(\mathbf{x})$
259 10: Update ψ by minimizing $\mathcal{L}_{\text{D}}^{n,t}(\psi)$ with \mathbf{x} and \mathbf{x}' .
260 11: **end if**
261 12: **end while**

262
263 3.4 REGULARIZING VGPO WITH ADVERSARIAL DISTRIBUTION MATCHING
264

265 There are several sources of gradient noise in Algorithm 3, including errors in value model prediction
266 and noise from Monte Carlo loss estimation. We observed error accumulation running Algorithm 3,
267 and find that the KL loss in Eq. (15) alone cannot fully rectify these errors. Tuning the KL weight
268 w_{KL} does not fully resolve this issue. When the KL weight is high, the original suboptimal behaviors

270 of the reference policy tend to be preserved in the target policy, leading to slow optimization. And
271 when the KL weight is low, it fails to correct the error accumulation.

272 We propose an alternative regularization method based on token-level adversarial distribution match-
273 ing (Goodfellow et al., 2014; Ho & Ermon, 2016; Huang et al., 2025). A discriminator network
274 $D_\psi(x_n^t)$ is trained to distinguish between true noisy tokens and fake noisy tokens by minimizing
275 $\mathbb{E}_{n,t}[\mathcal{L}_D^{n,t}(\psi)]$, where

$$278 \mathcal{L}_D^{n,t}(\psi) := \mathbb{E}_{\pi_{\text{sg}[\theta]}(\mathbf{x}), q(x_n^t | x_n)} [D_\psi(x_n^t) + 1]^2 + \mathbb{E}_{\mu(\mathbf{x}), q(x_n^t | x_n)} [D_\psi(x_n^t) - 1]^2. \quad (17)$$

280 And the generator is updated by minimizing $\mathbb{E}_{n,t}[\mathcal{L}_G^{n,t}(\psi)]$, where

$$283 \mathcal{L}_G^{n,t}(\theta) := \mathbb{E}_{\pi_{\text{sg}[\theta]}(\mathbf{x}), q(x_n^t | x_n)} \|s_\theta(x_n^t | x_{<n}) - \text{sg}[s_\theta(x_n^t | x_{<n})] - \nabla_{x_n^t} D_\psi(x_n^t)\|_2^2. \quad (18)$$

285 \mathcal{L}_G can be obtained from replacing V_ϕ with D_ψ in Eq. (16).

288 4 EXPERIMENTS

291 4.1 COMMON SETUP

293 **Base Model.** We fine-tuned a DiTAR model comprising approximately 0.4 billion parameters,
294 trained on an internal dataset containing about 280k hours of Chinese and English speech. The
295 language model (LM) within this architecture consists of 24 layers, while the diffusion head (DiT)
296 includes 4 layers. All Transformer layers in the model have a hidden dimension of 1024, 16 attention
297 heads, and a feed-forward network (FFN) dimension of 4096.

298 **Diffusion Sampler and CFG.** For all online and offline sampling in the experiments, we used the
299 DDPM sampler with 16 sampling steps, with LM Guidance (Jia et al., 2025) weight $w = 2$. We
300 always enable LM Guidance with $w = 2$ during training. See Appendix B.2 for more discussion
301 about classifier-free guidance (CFG) in VGPO.

302 **VGPO Training.** All experiments were conducted on 32 A100 GPUs. We used the AdamW opti-
303 mizer, with learning rate fixed to 1×10^{-6} , $\beta_1 = 0.9$, $\beta_2 = 0.95$, weight decay 0.01.

305 **Value Models.** All value models used in our experiments are initialized from the base DiTAR model,
306 with the last linear layer of the diffusion head replaced by a zero initialized linear layer that outputs
307 a scalar. We refer to the modified diffusion head as the value head. We choose to initialize the value
308 head with diffusion head parameters, which brings better performance than random initialization in
309 our preliminary experiments. The inputs to the value models are the same as in the DiTAR TTS
310 model, including prompt speech, prompt text, and target text.

311 **Objective Evaluations.** We report the word error rate (**WER**) for speech intelligibility using
312 Whisper-large-v3. For speaker similarity (**SIM**), we report the cosine similarity of speaker em-
313 beddings between prompts and generated audios, using a WavLM-TDCNN model. We computed
314 WER and SIM using the same models and evaluation code¹ as in Seed-TTS (Anastassiou et al.,
315 2024). For all evaluations, we run tests eight times and report the average. We also report the av-
316 erage KL loss (**KL**) on the test sets to report the divergence between the fine-tuned model and the
317 base model. For evaluation, we used the KL loss definition in Eq. (17) in Liu et al. (2025b).

318 **Subjective Evaluations.** We evaluate the subjective quality of fine-tuned models using comparative
319 mean opinion score (CMOS) for speech naturalness (N-CMOS), speaker similarity (S-CMOS), and
320 speech expressiveness (E-CMOS). Human listeners compare the generated audio against base model
321 response, and assign a score from -2 to 2. See Appendix C for more details.

322 **Baselines.** We compare VGPO against the following baselines: (1) Guided sampling results using
323 Algorithm 2. (2) ARDM-DPO, we evaluated the Liu et al. (2025b) (3) Base model sampling and
best-of-K (BoK) sampling.

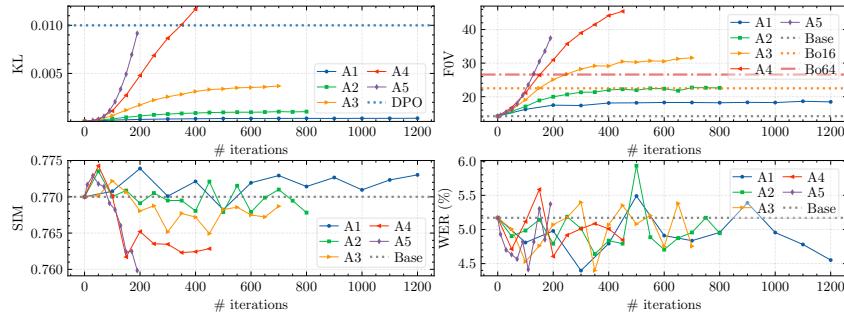


Figure 2: Evolution of KL, F0V, SIM, WER during VGPO training for task A.

4.2 TASK A: IMPROVING F0 VARIANCE

Motivation. We choose the fundamental frequency variance (F0V) as the reward for task A because it is a robust and label-free proxy of perceived expressiveness and directly counteracts monotony. F0V provides a simple and reproducible benchmark for testing the RFT algorithms of TTS models.

Dataset. We used the LibriTTS corpus (Zen et al., 2019) for both speech prompts and target texts, which includes 555 hours of recordings from 2,311 speakers. Evaluations were performed with 38 test cases, with prompts and target texts from 38 different speakers within the LibriTTS test-clean subset.

Reward Function. For a given utterance, we extract the F0 track of the voiced regions with an off-the-shelf F0 tracker. Then we apply a band-pass filter to the F0 track to focus on phrase-level variations. The reward is computed as the standard deviation of the filtered F0 track.

Value Training. We trained the value model directly with audios and transcripts from the LibriTTS corpus. The value model takes a transcript and an audio as input, and is trained to predict the F0V with the MSE loss. This is justified if we assume $\mu(x)$ is significantly similar to the dataset distribution. The value model is trained for 60k steps with a dynamic batch size of approximately 2 hours of speech per batch. The results of the value-guided sampling with this value model can be found in Table 1. The WER and SIM of the guided models are close to the base model, indicating good prior distribution preservation. We also report Best-of-N (BoN) results with $N \in \{16, 64\}$ in Table 1. We see that the value guidance can match the performance of Bo64 without reducing SIM.

VGPO Training. During VGPO training, we sampled 32 pairs of prompts and target texts. And for each pair, we generate about 2 minutes synthesized speech with a dynamic number of rollouts per iteration. The evolution of KL, F0V, and SIM during training can be found in Figure 2. Experiments A1, A2, A3, A4, A5 are trained with KL weight $2^{-4}, 2^{-5}, 2^{-6}, 2^{-7}, 2^{-8}$, respectively. When achieving similar F0V as the ARDM-DPO baseline, VGPO (600 steps, A3) obtains lower KL divergence, and higher SIM compared to DPO (200 steps, $\beta = 200$).

Table 1: Selected objective evaluation results on task A.

Method	F0V \uparrow	WER \downarrow	SIM \uparrow	KL \downarrow
Base	14.2	5.17	0.770	—
Best-of-16	22.5	4.74	0.770	—
Best-of-64	26.6	4.93	0.770	—
Guided $s = 2^2$	15.6	5.41	0.769	—
Guided $s = 2^4$	19.3	5.24	0.771	—
Guided $s = 2^5$	27.3	4.93	0.771	—
Guided $s = 2^6$	52.6	5.76	0.760	—
DPO $^{200 \text{ steps}}$ $\beta = 200$	29.2	3.73	0.765	0.010
VGPO $^{600 \text{ steps}}$ A_3	30.5	4.75	0.767	0.003

Table 2: Subjective evaluation results for task A.

Method	E-CMOS \uparrow	N-CMOS \uparrow	S-CMOS \uparrow
DPO $^{200 \text{ steps}}$ $\beta = 200$	1.70 ± 0.36	-0.14 ± 0.12	-0.05 ± 0.16
VGPO $^{600 \text{ steps}}$ A_3	1.65 ± 0.34	0.05 ± 0.13	-0.03 ± 0.20

¹<https://github.com/BytedanceSpeech/seed-tts-eval>

378 **Subjective Evaluations.** For each prompt in the test set, we generated three samples from the base
 379 model, DPO model, and VGPO model. We conducted N-CMOS, E-CMOS, and S-CMOS tests, for
 380 each we collected 114 scores. Evaluation results can be found in Table 2.
 381

382 4.3 TASK B: ENHANCING ROBUSTNESS TO DIFFICULT TEXTS

383 Table 3: Selected objective evaluation results on task B. Table contains results of VGPO with
 384 different hyperparameters. O.V. is abbreviation for online value update. w_{KL} is the weight of the KL
 385 loss. w_A is the weight of ADM loss.
 386

ID (steps)	O.V.	w_{KL}	w_A	NLL \downarrow	CER \downarrow	SIM \uparrow	KL \downarrow
Base Model	—	—	—	0.55	8.37	0.711	0
Best-of-8 (WER)	—	—	—	0.39	4.99	0.713	—
Best-of-8 (NLL)	—	—	—	0.27	6.79	0.712	—
DPO _{9000 steps} $\beta = 1600$	—	—	—	0.32	6.32	0.712	0.009
B1	No	2^{-11}	0	—	diverged	—	—
B2 (6k)	No	2^{-10}	0	0.34	7.56	0.698	0.143
B3 (6k)	No	2^{-9}	0	0.33	7.07	0.705	0.008
B4 (6k)	Yes	2^{-12}	0	0.29	6.52	0.696	0.037
B5 (6k)	Yes	2^{-11}	0	0.31	6.65	0.705	0.017
B6 (6k)	Yes	2^{-10}	0	0.33	6.58	0.709	0.007
B7 (6k)	Yes	2^{-9}	0	0.39	7.18	0.710	0.002
B4 (15k)	Yes	2^{-12}	0	0.26	6.18	0.689	0.660
B5 (15k)	Yes	2^{-11}	0	0.27	6.20	0.700	0.299
B6 (15k)	Yes	2^{-10}	0	0.28	6.48	0.707	0.099
B7 (15k)	Yes	2^{-9}	0	0.36	7.06	0.708	0.040
B8 (15k)	Yes	0	2^{-6}	0.29	6.27	0.725	0.027
B9 (15k)	Yes	0	2^{-3}	0.32	6.36	0.732	0.016

409 **Motivation.** Autoregressive TTS models often struggle to accurately read complex texts containing
 410 repetitive words or phrases. When evaluating our base model on such sentences, it frequently fails
 411 to correctly handle repetitions, either by omitting some repetitions, adding additional ones, or be-
 412 coming trapped in a loop of repetitive generation, unable to terminate. In task B, our objective is to
 413 enhance the robustness of the DiTAR model when processing these challenging texts.

414 **Dataset.** We used the DiDiSpeech-2 (Guo et al., 2021) dataset as the source of speech prompts,
 415 which comprises 227 hours of recordings from 1,500 speakers. We excluded all speakers included
 416 in the hard subset of the SEED-TTS-Eval test set. For the training text set, we used a corpus of
 417 100,000 long Chinese sentences. These sentences contain randomly repeated words, phrases, and
 418 clauses, making them difficult to synthesis correctly for autoregressive TTS models. All evaluations
 419 were performed on the test-hard subset of SEED-TTS-Eval. We excluded 2 test cases with the
 420 longest target texts, as they significantly exceed the context-length limit of our base model.

421 **Reward Modeling.** Following previous work (Du et al., 2025; Li et al., 2025), we utilize the likeli-
 422 hood of speech in automatic speech recognition (ASR) models as a proxy reward function. An alter-
 423 native reward choice is CER, but computing CER is significantly slower than evaluating likelihood.
 424 For all experiments on task B, we employ a phoneme-based CTC model trained on DidiSpeech-2 as
 425 the reward model.

426 **Group Reward Normalization.** We find it necessary to normalize the reward for each pair of
 427 prompts and target texts to train the value model. Otherwise, the value model failed to capture
 428 the relative differences between trajectories with the same prompt and text. We perform a reward
 429 normalization similar to GRPO (Shao et al., 2024). Suppose that for each pair of prompt and target
 430 text, the CTC likelihoods are $\mathbf{r} = (r_1, \dots, r_G) \in (0, 1)^G$ for G samples. We normalize the reward
 431 to $\tilde{\mathbf{r}} \in \mathbb{R}^G$ as $\tilde{\mathbf{r}} := \frac{\mathbf{r} - \text{mean}(\mathbf{r})}{\text{std}(\mathbf{r})}$. This normalization is applied in all experiments of task B.

432 **Value Training.** To train the value model, we randomly generated approximately 430k pairs of prompts and target texts.
 433 For each pair, we produced 16 samples using the base DiTAR
 434 model, resulting in an offline corpus containing 27k hours of
 435 speech. The value model was trained on this dataset for 150k
 436 steps with MSE loss in Eq. (13), with a dynamic batch size of
 437 approximately 3 hours of generated speech per batch.
 438

439 **VGPO Training.** During VGPO training, each iteration involved
 440 sampling 32 pairs of prompts and target texts, followed by generating approximately 8 rollouts for each pair.
 441 This process yielded approximately 2 hours of synthesized
 442 speech per training iteration. The value model (when updated online) and the discriminator (when
 443 ADM enabled) is trained 10 steps per iteration. The ARDM π_θ is always updated once per iteration.
 444

445 **Results.** From the results in Table 4 and Tabel 3, we see that VGPO leads to better performance than
 446 value guided sampling in task B. We observe that enabling online value update leads to more stable
 447 training and better performance. VGPO with more iterations not only leads to a lower NLL and CER
 448 but also causes an increase in the accumulation of errors, as reflected in the decrease of SIM. For
 449 example, model B4 (15k) sometimes generates audio with audible artifacts such as sudden change
 450 of volume. We find that regularization with ADM can significantly mitigate this issue. Experiments
 451 B8 and B9 in Table 3 are initialized from B6 (10k), and further trained for 5k iterations with ADM
 452 enabled. For discriminator training, we used only the first 7 seconds of speech generated from the
 453 base model, as they include less distribution drift caused by exposure bias. As a result, B8 and B9
 454 can beat the base model in SIM.

455 **Subjective Evaluations.** We randomly sampled
 456 40 test cases from the test set and generated 3
 457 random outputs each from the base, DPO, and
 458 VGPO models. For Task B, we conducted N-
 459 CMOS and S-CMOS evaluations, collecting 120
 460 scores for each test. Evaluation results can be
 461 found in Table 5.

462 5 RELATED WORK

463 RFT algorithms for non-autoregressive (NAR) diffusion models (Uehara et al., 2024; 2025) and rein-
 464forcement learning methods with diffusion policies (Zhu et al., 2023) are closely related to our work.
 465 However, it is unclear whether these approaches can be effectively applied to fine-tune ARDMs. Ex-
 466isting research in both directions shares several core ideas: (1) Policy gradient (PG) methods (Black
 467 et al., 2024; Fan et al., 2023; Liu et al., 2025a; Xue et al., 2025) cast the diffusion sampling process
 468 as a Gaussian MDP. (2) Reward-weighted regression (RWR) methods (Lee et al., 2023; Zhang et al.,
 469 2024; Dong et al., 2023) iteratively maximize the likelihood of high-reward samples. (3) Direct
 470 gradient (DG) methods (Xu et al., 2023; Clark et al., 2024; Li et al., 2025) differentiate through
 471 the diffusion sampling process to optimize the reward. (4) Multiple works adapt DPO to diffusion
 472 models (Wallace et al., 2024; Yang et al., 2024; Liu et al., 2025b).

473 6 CONCLUSION

474 We introduce Value-Guided Policy Optimization (VGPO), an online actor-critic reinforcement fine-
 475 tuning algorithm tailored for autoregressive diffusion models. We derive VGPO from the exact
 476 solution to KL-regularized policy optimization, thereby providing strong theoretical guarantees. We
 477 apply VGPO to fine-tune the DiTAR TTS model and evaluate it on two benchmarks. Empirically,
 478 VGPO exhibits stable training and achieves better results than DPO. For future work, VGPO can be
 479 applied to fine-tuning image and video ARDMs. In this work, we focus on an ARDM architecture
 480 with a causal transformer and a diffusion head. It would be interesting to investigate whether VGPO
 481 generalizes to other variants of autoregressive diffusion models, including masked autoregressive
 482 diffusion models (Li et al., 2024a) and diffusion-forcing models (Song et al., 2025).

Table 4: Results of value guided sampling with different guidance scale s on task B.

s	NLL \downarrow	CER \downarrow	SIM \uparrow
0	0.55	8.37	0.711
32	0.47	7.54	0.712
64	0.40	7.18	0.712
128	0.38	7.40	0.711
192	0.36	8.23	0.712

Table 5: Subjective evaluation results for task B.

Method	N-CMOS \uparrow	S-CMOS \uparrow
DPO _{$\beta = 1600$} 9000 steps	-0.03 ± 0.11	0.05 ± 0.12
VGPO _{B8} 15k steps	-0.05 ± 0.15	0.19 ± 0.23

486 REFERENCES
487

488 Philip Anastassiou, Jiawei Chen, Jitong Chen, Yuanzhe Chen, Zhuo Chen, Ziyi Chen, Jian Cong,
489 Lelai Deng, Chuang Ding, Lu Gao, et al. Seed-TTS: A family of high-quality versatile speech
490 generation models. *arXiv preprint arXiv:2406.02430*, 2024.

491 Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion
492 models with reinforcement learning. In *ICLR*, 2024.

493 Chen Chen, Yuchen Hu, Wen Wu, Helin Wang, Eng Siong Chng, and Chao Zhang. Enhancing
494 zero-shot text-to-speech synthesis with human feedback. *arXiv preprint arXiv:2406.00654*, 2024.

495 Jingyi Chen, Ju Seung Byun, Micha Elsner, Pichao Wang, and Andrew Perrault. Fine-tuning text-
496 to-speech diffusion models using reinforcement learning with human feedback. In *Interspeech*,
497 2025.

498 Kevin Clark, Paul Vicol, Kevin Swersky, and David J Fleet. Directly fine-tuning diffusion models
499 on differentiable rewards. In *ICLR*, 2024.

500 Haoge Deng, Ting Pan, Haiwen Diao, Zhengxiong Luo, Yufeng Cui, Huchuan Lu, Shiguang Shan,
501 Yonggang Qi, and Xinlong Wang. Autoregressive video generation without vector quantization.
502 *arXiv preprint arXiv:2412.14169*, 2024.

503 Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao,
504 Jipeng Zhang, KaShun SHUM, and Tong Zhang. RAFT: Reward rAnked FineTuning for genera-
505 tive foundation model alignment. *TMLR*, 2023.

506 Zhihao Du, Changfeng Gao, Yuxuan Wang, Fan Yu, Tianyu Zhao, Hao Wang, Xiang Lv, Hui Wang,
507 Chongjia Ni, Xian Shi, et al. CosyVoice 3: Towards in-the-wild speech generation via scaling-up
508 and post-training. *arXiv preprint arXiv:2505.17589*, 2025.

509 Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel,
510 Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. DPOK: Reinforcement learning for
511 fine-tuning text-to-image diffusion models. *NeurIPS*, 2023.

512 Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,
513 Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014.

514 Tingwei Guo, Cheng Wen, Dongwei Jiang, Ne Luo, Ruixiong Zhang, Shuaijiang Zhao, Wubo Li,
515 Cheng Gong, Wei Zou, Kun Han, et al. DiDiSpeech: A large scale mandarin speech corpus. In
516 *ICASSP*, 2021.

517 Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *NeurIPS*, 2016.

518 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint*
519 *arXiv:2207.12598*, 2022.

520 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*,
521 2020.

522 Jinyi Hu, Shengding Hu, Yuxuan Song, Yufei Huang, Mingxuan Wang, Hao Zhou, Zhiyuan Liu,
523 Wei-Ying Ma, and Maosong Sun. ACDiT: Interpolating autoregressive conditional modeling and
524 diffusion transformer. *arXiv preprint arXiv:2412.07720*, 2024.

525 Xun Huang, Zhengqi Li, Guande He, Mingyuan Zhou, and Eli Shechtman. Self Forcing: Bridging
526 the train-test gap in autoregressive video diffusion. *arXiv preprint arXiv:2506.08009*, 2025.

527 Dongya Jia, Zhuo Chen, Jiawei Chen, Chenpeng Du, Jian Wu, Jian Cong, Xiaobin Zhuang, Chumin
528 Li, Zhen Wei, Yuping Wang, et al. DiTAR: Diffusion transformer autoregressive modeling for
529 speech generation. In *ICML*, 2025.

530 Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun
531 Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1.5: Scaling reinforcement learning with
532 LLMs. *arXiv preprint arXiv:2501.12599*, 2025.

540 Siqi Kou, Jiachun Jin, Zhihong Liu, Chang Liu, Ye Ma, Jian Jia, Quan Chen, Peng Jiang, and Zhijie
 541 Deng. Orthus: Autoregressive interleaved image-text generation with modality-specific heads. In
 542 *ICML*, 2025.

543

544 Kimin Lee, Hao Liu, Moonkyung Ryu, Olivia Watkins, Yuqing Du, Craig Boutilier, Pieter Abbeel,
 545 Mohammad Ghavamzadeh, and Shixiang Shane Gu. Aligning text-to-image models using human
 546 feedback. *arXiv preprint arXiv:2302.12192*, 2023.

547 Sergey Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review.
 548 *arXiv preprint arXiv:1805.00909*, 2018.

549

550 Tianhong Li, Yonglong Tian, He Li, Mingyang Deng, and Kaiming He. Autoregressive image
 551 generation without vector quantization. In *NeurIPS*, 2024a.

552 Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scialia, Gokcen Eraslan, Surag Nair, Tommaso Bian-
 553 calani, Shuiwang Ji, Aviv Regev, Sergey Levine, et al. Derivative-free guidance in continuous
 554 and discrete diffusion models with soft value-based decoding. *arXiv preprint arXiv:2408.08252*,
 555 2024b.

556 Yinghao Aaron Li, Rithesh Kumar, and Zeyu Jin. DMOSpeech: Direct metric optimization via
 557 distilled diffusion model in zero-shot speech synthesis. In *ICML*, 2025.

558

559 Jie Liu, Gongye Liu, Jiajun Liang, Yangguang Li, Jiaheng Liu, Xintao Wang, Pengfei Wan,
 560 Di Zhang, and Wanli Ouyang. Flow-GRPO: Training flow matching models via online RL. *arXiv*
 561 *preprint arXiv:2505.05470*, 2025a.

562 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and
 563 transfer data with rectified flow. In *ICLR*, 2023.

564

565 Zhijun Liu, Shuai Wang, Sho Inoue, Qibing Bai, and Haizhou Li. Autoregressive diffusion trans-
 566 former for text-to-speech synthesis. *arXiv preprint arXiv:2406.05551*, 2024.

567

568 Zhijun Liu, Dongya Jia, Xiaoqiang Wang, Chenpeng Du, Shuai Wang, Zhuo Chen, and Haizhou
 569 Li. Direct preference optimization for speech autoregressive diffusion models. *arXiv preprint*
 570 *arXiv:2509.18928*, 2025b.

571

572 Cheng Lu, Huayu Chen, Jianfei Chen, Hang Su, Chongxuan Li, and Jun Zhu. Contrastive energy
 573 prediction for exact energy-guided diffusion sampling in offline reinforcement learning. In *ICML*,
 2023.

574

575 Haitong Ma, Tianyi Chen, Kai Wang, Na Li, and Bo Dai. Efficient online reinforcement learning
 576 for diffusion policy. In *ICML*, 2025.

577

578 NextStep Team, Chunrui Han, Guopeng Li, Jingwei Wu, Quan Sun, Yan Cai, Yuang Peng, Zheng
 579 Ge, Deyu Zhou, Haomiao Tang, et al. NextStep-1: Toward autoregressive image generation with
 continuous tokens at scale. *arXiv preprint arXiv:2508.10711*, 2025.

580

581 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
 582 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow
 583 instructions with human feedback. *NeurIPS*, 2022.

584

585 Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression:
 Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*, 2019.

586

587 Thomas F Quatieri. *Discrete-Time Speech Signal Processing: Principles and Practice*. Pearson,
 2002.

588

589 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 590 Finn. Direct preference optimization: Your language model is secretly a reward model. *NeurIPS*,
 2023.

591

592 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 593 Mingchuan Zhang, YK Li, Yang Wu, et al. DeepSeekMath: Pushing the limits of mathematical
 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

594 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*,
 595 2021a.

596

597 Kiwahn Song, Boyuan Chen, Max Simchowitz, Yilun Du, Russ Tedrake, and Vincent Sitzmann.
 598 History-guided video diffusion. In *ICML*, 2025.

599

600 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 601 Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*,
 602 2021b.

603

604 Yutao Sun, Hangbo Bao, Wenhui Wang, Zhiliang Peng, Li Dong, Shaohan Huang, Jianyong Wang,
 605 and Furu Wei. Multimodal latent language modeling with next-token diffusion. *arXiv preprint*
 606 *arXiv:2412.08635*, 2024.

607

608 Jinchuan Tian, Chunlei Zhang, Jiatong Shi, Hao Zhang, Jianwei Yu, Shinji Watanabe, and Dong Yu.
 609 Preference alignment improves language model-based TTS. In *ICASSP*, 2025.

610

611 Manan Tomar, Lior Shani, Yonathan Efroni, and Mohammad Ghavamzadeh. Mirror descent policy
 612 optimization. *arXiv preprint arXiv:2005.09814*, 2020.

613

614 Masatoshi Uehara, Yulai Zhao, Tommaso Biancalani, and Sergey Levine. Understanding rein-
 615 forcement learning-based fine-tuning of diffusion models: A tutorial and review. *arXiv preprint*
 616 *arXiv:2407.13734*, 2024.

617

618 Masatoshi Uehara, Yulai Zhao, Chenyu Wang, Xiner Li, Aviv Regev, Sergey Levine, and Tom-
 619 maso Biancalani. Inference-time alignment in diffusion models with reward-guided generation:
 620 Tutorial and review. *arXiv preprint arXiv:2501.09685*, 2025.

621

622 Bram Wallace, Meihua Dang, Rafael Rafailov, Linqi Zhou, Aaron Lou, Senthil Purushwalkam,
 623 Stefano Ermon, Caiming Xiong, Shafiq Joty, and Nikhil Naik. Diffusion model alignment using
 624 direct preference optimization. In *CVPR*, 2024.

625

626 Lilian Weng. Reward hacking in reinforcement learning. *lilianweng.github.io*, 2024.

627

628 Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao
 629 Dong. ImageReward: Learning and evaluating human preferences for text-to-image generation.
 630 In *NeurIPS*, 2023.

631

632 Zeyue Xue, Jie Wu, Yu Gao, Fangyuan Kong, Lingting Zhu, Mengzhao Chen, Zhiheng Liu, Wei
 633 Liu, Qiushan Guo, Weilin Huang, et al. DanceGRPO: Unleashing GRPO on visual generation.
 634 *arXiv preprint arXiv:2505.07818*, 2025.

635

636 Chenyu Yang, Shuai Wang, Hangting Chen, Wei Tan, Jianwei Yu, and Haizhou Li. SongBloom:
 637 Coherent song generation via interleaved autoregressive sketching and diffusion refinement. *arXiv*
 638 *preprint arXiv:2506.07634*, 2025.

639

640 Kai Yang, Jian Tao, Jiafei Lyu, Chunjiang Ge, Jiaxin Chen, Weihan Shen, Xiaolong Zhu, and Xiu
 641 Li. Using human feedback to fine-tune diffusion models without any reward model. In *CVPR*,
 642 2024.

643

644 Tianwei Yin, Qiang Zhang, Richard Zhang, William T Freeman, Fredo Durand, Eli Shechtman, and
 645 Xun Huang. From slow bidirectional to fast autoregressive video diffusion models. In *CVPR*,
 2025.

646

647 Heiga Zen, Viet Dang, Rob Clark, Yu Zhang, Ron J Weiss, Ye Jia, Zhifeng Chen, and Yonghui Wu.
 648 LibriTTS: A corpus derived from librispeech for text-to-speech. In *Interspeech*, 2019.

649

650 Shu Zhang, Xinyi Yang, Yihao Feng, Can Qin, Chia-Chih Chen, Ning Yu, Zeyuan Chen, Huan
 651 Wang, Silvio Savarese, Stefano Ermon, et al. HIVE: Harnessing human feedback for instructional
 652 visual editing. In *CVPR*, 2024.

653

654 Xueyao Zhang, Yuancheng Wang, Chaoren Wang, Ziniu Li, Zhuo Chen, and Zhizheng Wu. Ad-
 655 vancing zero-shot text-to-speech intelligibility across diverse domains via preference alignment.
 656 In *ACL*, 2025.

648 Zhengbang Zhu, Hanye Zhao, Haoran He, Yichao Zhong, Shenyu Zhang, Haoquan Guo, Tingting
 649 Chen, and Weinan Zhang. Diffusion models for reinforcement learning: A survey. *arXiv preprint*
 650 *arXiv:2311.01223*, 2023.

653 A ADDITIONAL DERIVATIONS

655 A.1 OPTIMAL SOLUTION FOR KL-CONSTRAINED RFT OBJECTIVE IN EQ. (5)

657 *Proof.* The objective in Eq. (5) can be written as

$$\begin{aligned}
 659 \quad \mathcal{J} &= \mathbb{E}_{\mathbf{x} \sim \rho(\mathbf{x})} \left[r(\mathbf{x}) - \log \frac{\rho(\mathbf{x})}{\mu(\mathbf{x})} \right] \\
 660 \\
 661 &= \mathbb{E}_{\mathbf{x} \sim \rho(\mathbf{x})} \left[-\log \frac{Z}{\exp r(\mathbf{x})} - \log \frac{\rho(\mathbf{x})}{\mu(\mathbf{x})} \right] + \text{const.} \\
 662 \\
 663 &= \mathbb{E}_{\mathbf{x} \sim \rho(\mathbf{x})} \left[-\log \frac{\rho(\mathbf{x})}{\pi(\mathbf{x})} \right] + \text{const.} = -D_{\text{KL}}(\rho(\mathbf{x}) \parallel \pi(\mathbf{x})) + \text{const.},
 \end{aligned} \tag{19}$$

666 where Z and the optimal policy $\pi(\mathbf{x})$ are defined in Eq. (6). \square

669 A.2 PROOF OF THEOREM 2 BASED ON SOLVING FOR $\pi(x_n^t | x_{<n})$

671 *Proof.* From Theorem 1 we know that $\pi(\mathbf{x}) = \mu(\mathbf{x}) \exp r(\mathbf{x}) / Z$. We can find the relation between
 672 $\pi(x_{\leq n})$ and $\mu(x_{\leq n})$ by marginalizing all $x_{>n}$.

$$\begin{aligned}
 673 \quad \pi(x_{\leq n}) &= \int \mu(\mathbf{x}) \frac{\exp r(\mathbf{x})}{Z} dx_{>n} = \int \mu(x_{>n} | x_{\leq n}) \mu(x_{\leq n}) \frac{\exp r(\mathbf{x})}{Z} dx_{>n} \\
 674 \\
 675 &= \mu(x_{\leq n}) \cdot \frac{1}{Z} \cdot \int \mu(x_{>n} | x_{\leq n}) \exp r(\mathbf{x}) dx_{>n} = \mu(x_{\leq n}) \frac{\exp V(x_{\leq n})}{Z}.
 \end{aligned} \tag{20}$$

678 Our goal is to obtain the relation between $\pi(x_n | x_{<n})$ and $\mu(x_n | x_{<n})$. Divide both sides of Eq. (20)
 679 with $\pi(x_{<n})$ gives:

$$\frac{\pi(x_{\leq n})}{\pi(x_{<n})} = \pi(x_n | x_{<n}) = \frac{\mu(x_{\leq n}) \exp V(x_{\leq n})}{\mu(x_{<n}) \exp V(x_{<n})} = \mu(x_n | x_{<n}) \frac{\exp V(x_{\leq n})}{\exp V(x_{<n})}. \tag{21}$$

684 Following Lu et al. (2023), we multiply both sides of Eq. (21) with $q(x_n^t | x_n)$ and marginalize x_n ,
 685 which gives:

$$\begin{aligned}
 686 \quad \pi(x_n^t | x_{<n}) &= \int \pi(x_n | x_{<n}) q(x_n^t | x_n) dx_n = \int q(x_n^t | x_n) \mu(x_n | x_{<n}) \frac{\exp V(x_{\leq n})}{\exp V(x_{<n})} dx_n \\
 687 \\
 688 &= \mu(x_n^t | x_{<n}) \cdot \frac{1}{\exp V(x_{<n})} \cdot \int \mu(x_n | x_n^t, x_{<n}) \exp V(x_{\leq n}) dx_n \\
 689 \\
 690 &= \mu(x_n^t | x_{<n}) \frac{\exp V(x_{<n}, x_n^t)}{\exp V(x_{<n})}.
 \end{aligned} \tag{22}$$

694 Now apply logarithm on both sides of Eq. (22) and take the gradient with respect to x_n^t gives our
 695 result in Eq. (10). \square

697 A.3 PROOF OF THEOREM 2 BASED ON MAXIMUM ENTROPY RL

699 In this section, we analyze diffusion models and ARDMs using the discrete-time sampler introduced
 700 in DDPM Ho et al. (2020). This choice does not limit our results to the DDPM sampler, as will
 701 become clear from the following derivation. The following derivation are based on the Markov
 chain view of ARDMs. See Liu et al. (2025b, Fig. 1).

702 A.3.1 BACKGROUND: MAXIMUM ENTROPY RL
703704 Consider a Markov Decision Process (MDP) in which all trajectories consist of T steps. Each
705 trajectory can be expressed as follows:

706
$$\tau := (s_0, a_0, s_1, a_1, \dots, s_T), \quad (23)$$

707

708 where s_T represents a terminal state. Given a reference policy $\mu(a|s)$ and a target policy $\pi(a|s)$,
709 Maximum Entropy Reinforcement Learning (MaxEnt RL) (Levine, 2018) optimizes the following
710 objective:

711
$$\pi^* := \arg \max_{\pi} \mathbb{E}_{\pi} \left[r(s_T) - \sum_{t=0}^{T-1} D_{\text{KL}}(\pi(\cdot|s_t) \| \mu(\cdot|s_t)) \right], \quad (24)$$

712

713 where D_{KL} denotes the Kullback-Leibler divergence.
714715 The soft state value function, V^{π} , is defined as:

716
$$V^{\pi}(s_t) := \mathbb{E}_{\pi} \left[r(s_T) - \sum_{u=t}^{T-1} D_{\text{KL}}(\pi(\cdot|s_u) \| \mu(\cdot|s_u)) \mid s_t \right]. \quad (25)$$

717

718 Similarly, the soft Q function, Q^{π} , is defined as:
719

720
$$Q^{\pi}(s_t, a_t) := \mathbb{E}_{s_{t+1} \sim p(s_{t+1}|s_t, a_t)} [V^{\pi}(s_{t+1})], \quad t < T, \quad (26)$$

721

722 where $p(s_{t+1}|s_t, a_t)$ represents the transition dynamics of the environment.
723724 For the optimal policy π^* , along with the corresponding optimal value functions V^* and Q^* , the
725 following expression holds (Levine, 2018):
726

727
$$\pi^*(a_t|s_t) := \mu(a_t|s_t) \cdot \frac{\exp(Q^*(s_t, a_t))}{\exp(V^*(s_t))}. \quad (27)$$

728

729 Additionally, the optimal soft value function V^* satisfies the following:
730

731
$$V^*(s_t) = \log \int \mu(a_t|s_t) \cdot \exp(Q^*(s_t, a_t)) da_t. \quad (28)$$

732

733 With the above results from MaxEnt RL, we are ready to prove Theorem 2.
734735 *Proof.* The Markov sampling chain of ARDMs can be embedded in an MDP. The state space and
736 initial distribution of the MDP are the same as those of the Markov chain in Liu et al. (2025b, Fig. 1).
737 In state $s_n^t = (x_{1:n}^0, x_n^t)$, where $t > 1$, the action is x_n^{t-1} , determined by the ARDM. Given action
738 x_n^{t-1} in state s_n^t , the MDP deterministically transitions to $(x_{1:n}^0, x_n^{t-1})$. In state s_n^0 , where $n < N$,
739 the action is a randomly sampled Gaussian noise x_{n+1}^T . Given action x_{n+1}^T in state s_n^0 , the MDP
740 deterministically transitions to $(x_{1:n}^0, x_{n+1}^T)$. Intermediate states and actions in the MDP do not
741 receive a reward; only the terminal state receives a reward $r(x_{1..N}^0)$.
742743 Since we have mapped the ARDM sampling Markov chain as an MDP, we can analyze it using
744 MaxEnt RL. For ease of description, we will first flatten the sampling Markov chain of ARDM by
745 identifying state s_{nT-t} with state s_n^t . The Markov chain of ARDM sampling can now be written as
746 $s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_{NT}$. First, since the MDP has deterministic transitions, we have
747

748
$$V^*(s_{u+1}) = Q^*(s_u, a_u) \quad (29)$$

749

750 for all possible pairs (s_u, a_u, s_{u+1}) . Suppose that the reference policy is $\mu(s_{u+1}|s_u)$, the optimal
751 policy $\pi^*(s_{u+1}|s_u)$ satisfies:
752

753
$$\pi^*(s_{u+k}|s_u) \cdot \exp(V^*(s_u)) = \mu(s_{u+k}|s_u) \cdot \exp(V^*(s_{u+k})), \quad (30)$$

754

755 Rewriting Equation 30 by expanding the states as tokens, we have the following.
756

757
$$\pi^*(x_n^t|x_{1:n}^0) = \mu(x_n^t|x_{1:n}^0) \cdot \frac{\exp(V^*(x_{1:n}^0, x_n^t))}{\exp(V^*(x_{1:n}^0))}. \quad (31)$$

758

□

756 A.4 NUMERICAL STABILITY OF EXP-MSE LOSS
757

758 The Exp-MSE loss in Eq. (12) is numerically unstable during optimization. We propose an alterna-
759 tive loss that shares the same global minimum, while providing better numerical stability. Here, sg
760 is the stop gradient operator.

$$761 \quad \widehat{\mathcal{L}}_V^{n,t}(\phi) := \mathbb{E}_{\mu(\mathbf{x}), q(x_n^t | x_n)} \left[\exp(V_\phi(x_{<n}, x_n^t) - \text{sg}[V_\phi(x_{<n}, x_n^t)]) - \right. \\ 762 \quad \left. \exp(r(\mathbf{x}) - \text{sg}[V_\phi(x_{<n}, x_n^t)]) \right]^2. \quad (32)$$

766 *Proof.* To see why is this the case, consider the following simplified version of the Exp-MSE loss.
767 Here $f_\theta(x)$ is a neural network estimator, and $r(x, y)$ is an arbitrary bounded function depending on
768 both x, y . The conditional distribution of y given x is $p(y|x)$.
769

$$770 \quad \ell_\theta(x) := \int p(y|x) (\exp f_\theta(x) - \exp r(x, y))^2 dy. \quad (33)$$

772 The gradient has high variance since it contains multiple exponential functions:
773

$$774 \quad \nabla_\theta \ell_\theta(x) = \int p(y|x) \left[2(\exp f_\theta(x) - \exp r(x, y)) \exp f_\theta(x) \right] \nabla_\theta f_\theta(x) dy. \quad (34)$$

776 We propose the following surrogate loss:
777

$$778 \quad \hat{\ell}_\theta(x) := \int p(y|x) \left(\exp(f_\theta(x) - \text{sg}[f_\theta(x)]) - \exp(r(x, y) - \text{sg}[f_\theta(x)]) \right)^2 dy. \quad (35)$$

781 The gradient of this surrogate loss is

$$782 \quad \nabla_\theta \hat{\ell}_\theta(x) = 2 \left(1 - \int p(y|x) \exp(r(x, y) - f_\theta(x)) dy \right) \nabla_\theta f_\theta(x). \quad (36)$$

785 This loss has the desired unique global minimum:
786

$$787 \quad f_\theta^*(x) = \log \int p(y|x) \exp r(x, y) dy. \quad (37)$$

□

791 B MORE IMPLEMENTATION DETAILS
792793 B.1 DISCRIMINATOR ARCHITECTURE IN TASK B
794

795 Adversarial distribution matching is applied to Task B VGPO training. The discriminator in our
796 experiments is initialized from the base DiTAR model. The LM in DiTAR is adapted as a prompt
797 encoder, taking prompt speech as input and outputting a speaker embedding. The diffusion head is
798 adapted into a discriminator network that takes a speaker embedding, a noisy token, and its diffu-
799 sion time as input. Similar to the value model, the discriminator is first trained on offline sampled
800 trajectories from the base model.
801

802 B.2 CLASSIFIER-FREE GUIDANCE IN VGPO
803

804 Many ARDMs rely on classifier-free guidance (CFG) to obtain high-quality samples (Song et al.,
805 2025; Sun et al., 2024; Jia et al., 2025). However, CFG violates the assumption that $q(x_n^t | x_n)$ can
806 approximate the true marginal distribution $p(x_n^t | x_n)$, creating a gap between theory and practice. In
807 our implementation, we consistently enabled LM guidance with a fixed weight $w = 2$ in DiTAR,
808 which we found to work well. Consequently, the diffusion head is evaluated twice in each inference
809 and training step. Because the diffusion head in DiTAR is lightweight, the additional computational
cost remains acceptable.

810 **C SUBJECTIVE EVALUATIONS**
811

812 We conducted comparative mean opinion score (CMOS) tests to evaluate the performance of fine-
813 tuned models against the base model. The evaluation user interfaces for the different CMOS tests
814 are shown in Figures 3, 4, and 5. In all tests, listeners are presented with a set of instructions, the
815 target text, and audio generated by two models. For the speaker similarity test, the prompt audio is
816 provided as the reference.
817

818 **Speech Naturalness Listening Test**

819 **Instructions**

820 1. Listen to the reference audio first.
821 2. Then listen to Audio A and Audio B in any order.
822 3. Use headphones in a quiet environment.
823 4. Choose which audio sounds more natural overall.

824 **Text**

825 The queens had taken their seats upon a magnificent dias or platform, erected upon the borders of the lake, in a theater of wonderful elegance of construction.

826 **Audio A**

827 **Audio B**

828 **Your Preference**

829 Which audio sounds more natural? *

830 A +2 A +1 No Preference B +1 B +2

831 **Submit**

832 **Figure 3: Screen shot of CMOS evaluation interface for speech naturalness (N-CMOS).**
833

834 **Speaker Similarity Listening Test**

835 **Instructions**

836 1. Listen to the reference audio first.
837 2. Then listen to Audio A and Audio B in any order.
838 3. Use headphones in a quiet environment.
839 4. Judge which test audio sounds like it was spoken by the same person as the reference speaker.
840 5. Focus only on speaker identity (the voice itself), not on how the sentence is read.

841 **Reference Audio**

842 **Text**

843 The queens had taken their seats upon a magnificent dias or platform, erected upon the borders of the lake, in a theater of wonderful elegance of construction.

844 **Audio A**

845 **Audio B**

846 **Your Preference**

847 Which speaker sounds more similar to the reference? *

848 A +2 A +1 No Preference B +1 B +2

849 **Submit**

850 **Figure 4: Screen shot of CMOS evaluation interface for speaker similarity (S-CMOS).**
851852 **D LLM USAGE**
853

854 LLMs including GPT-5 and GPT-4o are used for correcting grammar errors.
855
856

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882

883 **Speech Expressiveness Listening Test**

884

885 **Instructions**

886 1. Listen to the reference audio first.
887 2. Then listen to Audio A and Audio B in any order.
888 3. Use headphones in a quiet environment.
889 4. Choose which audio sounds more expressive overall.

890

891 **Text**
892 The queens had taken their seats upon a magnificent dias or platform, erected upon the borders of the lake, in a theater of wonderful elegance of construction.

893

894 **Audio A** **Audio B**

895

896 **Your Preference**
897 Which audio sounds more expressive? *

898 **A +2** **A +1** **No Preference** **B +1** **B +2**

899

900 **Submit**

901 **Figure 5: Screen shot of CMOS evaluation interface for speech expressiveness (E-CMOS).**

902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917