

000 001 002 003 004 005 FARI: ROBUST ONE-STEP INVERSION FOR WATER- 006 MARKING IN DIFFUSION MODELS 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030

ABSTRACT

031 Inversion-based watermarking is a promising approach to authenticate diffusion-
032 generated images, yet practical use is bottlenecked by inversion that is both slow
033 and error-prone. While the primary challenge in the watermarking setting is ro-
034 bustness against external distortions, existing approaches over-optimize internal
035 truncation error, and because that error scales with the sampler step size, they are
036 inherently confined to high-NFE (number of function evaluations) regimes that
037 cannot meet the dual demands of speed and robustness. In this work, we have two
038 key observations: (i) the inversion trajectory has markedly lower curvature than
039 the forward generation path does, making it highly compressible and amenable
040 to low-NFE approximation; and (ii) in inversion for watermark verification, the
041 trade-off between speed and truncation error is less critical, since external distor-
042 tions dominate the error. A faster inverter provides a dual benefit: it is not only
043 more efficient, but it also enables end-to-end adversarial training to directly target
044 robustness, a task that is computationally prohibitive for the original, lengthy in-
045 version trajectories. Building on this, we propose **FARI** (Fast Asymmetric Robust
046 Inversion), a one-step inversion framework paired with lightweight adversarial
047 LoRA fine-tuning of the denoiser for watermark extraction. While consolidation
048 slightly increases internal error, FARI delivers large gains in both speed and ro-
049 bustness: with 20 minutes of fine-tuning on a single NVIDIA RTX A6000 GPU,
050 it surpasses 50-step DDIM inversion on watermark-verification robustness while
051 dramatically reducing inference time.
052
053

1 INTRODUCTION

034 The rapid proliferation of diffusion models (Ho et al., 2020; Song et al., 2021) has led to an explosion
035 of AI-generated content, simplifying creative production but also fueling the spread of synthetic
036 misinformation and raising concerns about intellectual property protection for model providers. In
037 response, inversion-based watermarking (Yang et al., 2024; Wen et al., 2023; Huang et al., 2024;
038 Gunn et al., 2024) has shown remarkable promise for authenticating and tracing diffusion-generated
039 images. By embedding a watermark in the initial noise, the mark becomes deeply integrated with
040 the image’s semantics during the iterative generation process, ensuring minimal visual impact. To
041 extract the watermark, the image typically needs to be reconstructed back to noise via inversion
042 techniques (Song et al., 2021; Hong et al., 2024). This inversion step, however, is the method’s
043 critical bottleneck. It is computationally expensive, time-consuming, and introduces substantial
044 errors, all of which hinder the practical, large-scale deployment of inversion-based watermarks. This
045 bottleneck motivates the development of a fast and accurate inversion method tailored for watermark
046 extraction.

047 Although many inversion techniques (Mokady et al., 2023; Hong et al., 2024; Wallace et al., 2023)
048 exist, most primarily aim to reduce internal inversion error (caused by discretization truncation
049 and Classifier-Free Guidance) via iteration (Pan et al., 2023; Garibi et al., 2024; Samuel et al.,
050 2023), optimization (Hong et al., 2024; Li et al., 2024; Mokady et al., 2023), or analytical control of
051 truncation-error bounds (Wallace et al., 2023; Zhang et al., 2024a; Wang et al., 2024). While effec-
052 tive for diffusion-based image editing (Hertz et al., 2022), this strategy is ill-suited for watermarking.
053 Prior to extraction, images may be subjected to diverse distortions (e.g., JPEG compression, blur)
054 that induce substantial initial-condition shifts; given the denoiser’s sensitivity, these perturbations

054 compound rapidly along the inversion trajectory and become the dominant bottleneck to extraction
 055 accuracy.
 056

057 This shift in the error bottleneck for watermark extraction leads us to question the necessity of
 058 traditional high-NFE inversion. In prior methods, which primarily address clean-image scenarios,
 059 performance is limited by the discretization truncation error of the ODE sampler, which is directly
 060 related to the step size. Consequently, a high NFE is required to maintain precision. In the wa-
 061 termarking context, however, this internal error is dwarfed by the accumulated error from external
 062 distortions, a factor that is not explicitly mitigated by a larger number of steps. Furthermore, the
 063 natural solution to instill this robustness, adversarial training, is blocked by the high NFE of these
 064 traditional inverters. A powerful end-to-end training regime is rendered computationally infeasible
 065 by the prohibitive memory costs of backpropagating through a long iterative process. Meanwhile,
 066 the more computationally feasible, factorized objective, similar to that used in diffusion pretraining
 067 (Ho et al., 2020), proves insufficient for learning the global robustness required to counter complex
 068 distortions (see Appendix E.3). These facts indicate that first finding a low-NFE solution is not only
 069 beneficial for speed, but also enables a breakthrough in enhancing robustness.
 070

071 Motivated by this, we propose **FARI**: Fast Asymmetric Robust Inversion, a framework that achieves
 072 fast and robust inversion tailor-made for watermarking at a minimal fine-tuning cost. FARI is based
 073 on a key insight into the geometric asymmetry between generation and inversion trajectories: while
 074 the estimation error in inversion makes the reconstructed noise inaccurate, it also indirectly en-
 075 dows the inversion path with a significantly lower curvature than its generation counterpart. A
 076 lower-curvature trajectory is inherently easier to approximate with fewer steps. This enables a step-
 077 distillation approach that collapses multi-step inversion into a single efficient step. This reduction
 078 in NFE unlocks efficient end-to-end adversarial training. While this distillation-based estimation
 079 slightly sacrifices precision on clean, distortion-free inversion, the direct speed-up and the indirect
 080 enhancement in robustness are substantial, and we find that the downside of this trade-off has a
 081 negligible effect on the performance of the downstream watermarking task (Yang et al., 2024; Wen
 082 et al., 2023). Furthermore, our use of LoRA (Hu et al., 2022) for fine-tuning elegantly avoids the
 083 degradation of image quality. By storing the learned robustness knowledge externally in the LoRA
 084 parameters, we can simply deactivate the LoRA branch during generation, ensuring that the orig-
 085 inal model’s generation quality remains unchanged. Our experiments demonstrate that with just
 086 20 minutes of fine-tuning on a single NVIDIA RTX A6000 GPU, the one-step FARI surpasses the
 087 robustness of the 50-step DDIM baseline in watermark verification tasks.
 088

089 2 BACKGROUND

090 2.1 DIFFUSION MODELS

091 Diffusion models (Ho et al., 2020; Song et al., 2021) are a class of generative models that operate
 092 by iteratively transforming a pure Gaussian noise vector $\mathbf{z}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ into a real data sample
 093 $\mathbf{z}_0 \sim q(\mathbf{z})$ through T denoising steps. The process is defined by two Markov chains. The forward
 094 process gradually diffuses a data sample \mathbf{z}_0 by adding Gaussian noise over T timesteps according
 095 to a fixed variance schedule $\{\beta_t\}_{t=1}^T$:

$$096 q(\mathbf{z}_t | \mathbf{z}_{t-1}) = \mathcal{N}(\mathbf{z}_t; \sqrt{1 - \beta_t} \mathbf{z}_{t-1}, \beta_t \mathbf{I}), \quad (1)$$

097 A key property of this process is that we can sample \mathbf{z}_t at any arbitrary timestep t directly from \mathbf{z}_0 :

$$098 \mathbf{z}_t = \sqrt{\bar{\alpha}_t} \mathbf{z}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}, \quad (2)$$

099 where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$, and $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The reverse process learns to denoise these
 100 corrupted samples to recover the original data. This is achieved by training a neural network ϵ_θ to
 101 predict the added noise $\boldsymbol{\epsilon}$ from the noisy input \mathbf{z}_t . The objective function is typically a simplified
 102 version of the evidence lower bound:

$$103 \mathcal{L}(\theta) = \mathbb{E}_{\mathbf{z}_0, t \sim \text{Uniform}(1, T), \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\|\boldsymbol{\epsilon} - \epsilon_\theta(\mathbf{z}_t, t)\|_2^2], \quad (3)$$

104 2.2 DDIM SAMPLING AND INVERSION

105 The denoising diffusion implicit model (Song et al., 2021) (DDIM) provides a deterministic sam-
 106 pling process by defining a non-Markovian forward process that leads to the same marginal distri-

108 butions. Given a noisy latent \mathbf{z}_t , DDIM computes the subsequent latent \mathbf{z}_{t-1} by first predicting an
 109 estimate of the clean image, $\hat{\mathbf{z}}_0$, and then stepping towards it:
 110

$$\hat{\mathbf{z}}_0 = \frac{\mathbf{z}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{z}_t, t)}{\sqrt{\bar{\alpha}_t}}, \quad (4)$$

$$\mathbf{z}_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \hat{\mathbf{z}}_0 + \sqrt{1 - \bar{\alpha}_{t-1}} \epsilon_\theta(\mathbf{z}_t, t). \quad (5)$$

114 The deterministic nature of DDIM is crucial as it allows for an invertible generation process, which
 115 iteratively computes \mathbf{z}_t from \mathbf{z}_{t-1} by reversing the sampling steps. This unique invertible charac-
 116 teristic allows us to recover the initial noise representation \mathbf{z}_T from any generated image \mathbf{z}_0 , which
 117 serves as a powerful tool for inversion-based watermarking.

118 2.3 INVERSION-BASED WATERMARKING FOR DIFFUSION MODELS

120 We categorize these methods into three classes. The first class, epitomized by Tree-Ring (Wen et al.,
 121 2023), embeds a robust pattern into the Fourier domain of the initial noise to enable detection. Sub-
 122 sequent works have focused on enhancing its practical applications or extending its capabilities. For
 123 instance, RingID (Ci et al., 2024) extends it to a multi-bit watermark, ROBIN (Huang et al., 2024)
 124 improves its imperceptibility, and ZoDiac (Zhang et al., 2024b) generalizes it as a post-processing
 125 watermark, all without altering the core embedding and extraction logic. The second class, rep-
 126 resented by Gaussian Shading (Yang et al., 2024), embeds a multi-bit watermark into the spatial
 127 domain of the noise through distribution-preserving sampling. Follow-up research has concentrated
 128 on improving its key reuse problem, as seen in PRC-Watermark (Gunn et al., 2024) and Gaussian
 129 Shading++ (Yang et al., 2025), and on functional extensions; for example, TAG-WM (Chen et al.,
 130 2025) and VideoShield (Hu et al., 2025) provide functionality for detecting tampered regions. The
 131 third class, such as GaussMarker (Li et al., 2025), combines the first two approaches to compensate
 132 for their weakness against geometric distortions.

133 2.4 INVERSION METHODS

135 There is a substantial body of work on diffusion model inversion. Methods such as BELM (Wang
 136 et al., 2024), BDIA (Zhang et al., 2024a), and EDICT (Wallace et al., 2023) directly modify the
 137 sampling process to make it invertible. Others, including AIDI (Pan et al., 2023), ExactDPM (Hong
 138 et al., 2024), and ReNoise (Garibi et al., 2024), employ iteration or gradient descent to obtain better
 139 intermediate values for trajectory alignment. A third category, which includes NTI (Mokady et al.,
 140 2023) and NPI (Miyake et al., 2025), focuses on optimizing a better null-text embedding to guide
 141 the regeneration process. As we have previously mentioned, these methods are primarily designed
 142 for training-free image editing. Consequently, they may fail in adversarial watermark extraction
 143 scenarios, a point we will demonstrate in our experiments section.

144 2.5 DIFFUSION MODEL ACCELERATION

146 The acceleration of diffusion models can be broadly categorized into two paths. The first path in-
 147 volves using solvers with lower truncation error (Lu et al., 2022a;b; Zhang & Chen, 2022). While
 148 these methods can reduce the number of inference steps to between 20 and 30, the quality of image
 149 generation in extreme few-step scenarios (e.g., < 10) remains unsatisfactory. A noteworthy method
 150 in this category is the AMED-Solver (Zhou et al., 2024a), which is based on the mean value the-
 151 orem. It uses a small model to predict the timestep where the mean value occurs, thereby estimating
 152 the average velocity and enabling generation in as few as two steps. The second path is distilla-
 153 tion, where a student model is trained to replicate the output of multiple teacher steps in a single
 154 step. Techniques such as progressive distillation (Salimans & Ho, 2022), consistency distillation
 155 (Song et al., 2023), and distribution matching distillation (Yin et al., 2024b;a) follow this paradigm.
 156 However, these methods typically demand a substantial amount of pre-generated training data, GPU
 157 memory, and time.

158 3 OUR PROPOSED METHOD

159 In this section, we propose FARI, a robust one-step inversion method designed for watermark ex-
 160 traction. It is based on our key finding that the inversion trajectory exhibits a significantly lower

162 curvature than the generation path does, enabling efficient one-step distillation, which in turn makes
 163 the adversarial fine-tuning computationally feasible.
 164

165 3.1 THE INVERSION TRAJECTORY EXHIBITS LOWER CURVATURE 166

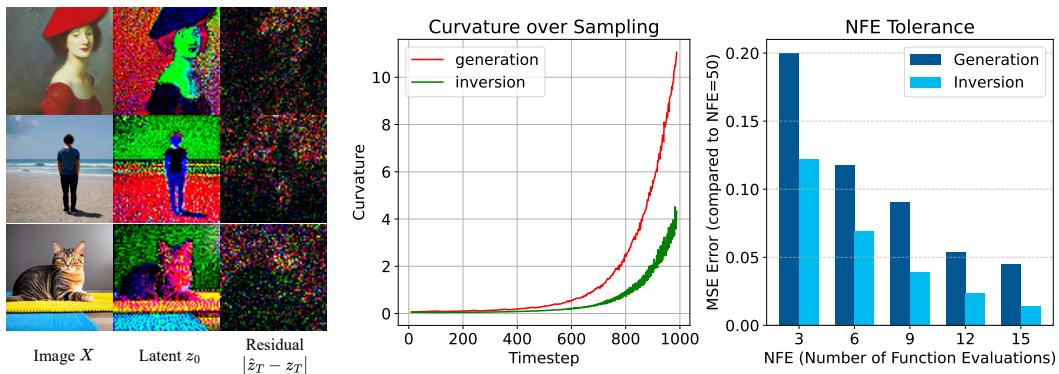
167 We begin with the inherent systematic error in DDIM inversion (Song et al., 2021). For deterministic
 168 DDIM sampling, the denoising process, which computes z_{t-1} from z_t , can be written in a single
 169 recurrence relation(the conditioning terms are omitted for simplicity):
 170

$$171 z_{t-1} = \sqrt{\frac{\bar{\alpha}_{t-1}}{\bar{\alpha}_t}} z_t + \left(\sqrt{1 - \bar{\alpha}_{t-1}} - \sqrt{\frac{\bar{\alpha}_{t-1}(1 - \bar{\alpha}_t)}{\bar{\alpha}_t}} \right) \epsilon_\theta(z_t, t). \quad (6)$$

174 The inversion step, which solves for z_t based on z_{t-1} , is derived as:
 175

$$176 z_t = \sqrt{\frac{\bar{\alpha}_t}{\bar{\alpha}_{t-1}}} z_{t-1} + \left(\sqrt{1 - \bar{\alpha}_t} - \sqrt{\frac{\bar{\alpha}_t(1 - \bar{\alpha}_{t-1})}{\bar{\alpha}_{t-1}}} \right) \epsilon_\theta(z_t, t). \quad (7)$$

177 However, since our goal is to solve for z_t , the term $\epsilon_\theta(z_t, t)$ on the right-hand side of Eq. 7 cannot be
 178 explicitly calculated. Generally, this is addressed by making a piecewise linear assumption, approximating
 179 $\epsilon_\theta(z_t, t) \approx \epsilon_\theta(z_{t-1}, t)$. The validity of this assumption, however, requires a sufficiently
 180 small step size, a condition that practical settings often fail to meet. This becomes a significant
 181 source of inversion error, even for clean images. While many works (Lin et al., 2024; Wang et al.,
 182 2024; Staniszewski et al., 2024) have recognized that changes in the trajectory direction $\epsilon_\theta(\cdot)$ cause
 183 an offset of the reconstructed noise \hat{z}_T and have attempted to mitigate this asymmetry, we further
 184 point out that under the combined effect of directional and positional offsets, curvature—a higher-
 185 order property of the trajectory—also exhibits a profound asymmetry. Specifically, the curvature of
 186 the inversion trajectory is substantially lower than that of the denoising trajectory.
 187



201 Figure 1: **Left:** Visualization of the inversion error, where latent vectors are projected down to 3
 202 channels via PCA for display. **Middle:** Curvature of generation and inversion trajectories across
 203 diffusion timesteps. **Discrete curvature estimated using 100 unconditionally generated images from**
 204 **Stable Diffusion v2.1.** **Right:** The resulting error for generation and inversion when reducing the
 205 NFE, compared with a 50-step NFE baseline.

206 In Figure 1(middle), we illustrate the curvature differences between the 1000-step denoising (generation)
 207 and inversion trajectories. We observe that trajectories exhibit greater curvature near the noise
 208 end of the process (as $t \rightarrow T$). This is because the denoising network is trained on the forward diffu-
 209 sion process, where different images can diffuse to the same noise point, causing trajectory crossing
 210 (Lee et al., 2023). Consequently, in the early stages of denoising, the model must constantly correct
 211 its direction, leading to high curvature (Lee et al., 2023). This effect is particularly pronounced at the
 212 very beginning when the latent variable is nearly pure noise. Once the fundamental semantics of the
 213 image have formed, the direction of progress becomes relatively fixed, and the trajectory’s curvature
 214 decreases significantly. However, this high-curvature phenomenon is substantially less pronounced
 215 during the inversion process.

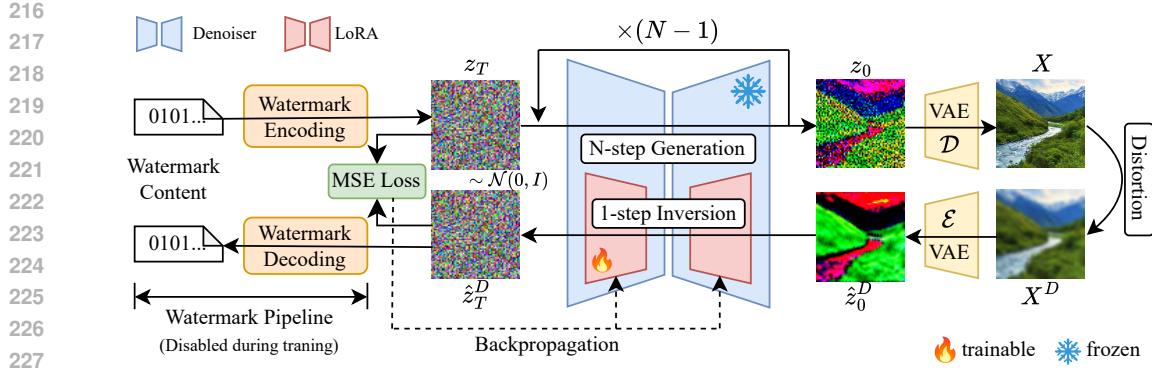


Figure 2: The framework of FARI. FARI simultaneously performs one-step distillation and adversarial training through a unified, end-to-end, LoRA-based fine-tuning process, enhancing both the efficiency and robustness of the inversion. The LoRA adapters are injected into the denoiser network and are deactivated during generation but activated for inversion. This strategy prevents any degradation of the original model’s generation quality and eliminates the need to deploy a second, complete denoiser, making it highly memory-efficient.

We attribute this partly to the fact that the accumulated error during inversion retains low-frequency information from the source image. As shown in Figure 1(left), partial outlines of the image remain visible in the noise reconstruction error, an observation consistent with prior work (Lin et al., 2024; Staniszewski et al., 2024; Nguyen et al., 2025). This residual semantic information helps to more accurately determine the correct direction of progress as the inversion approaches the noise end.

In general, a trajectory with lower curvature can be more accurately approximated with fewer linear steps (i.e., a lower NFE), as curvature is strongly correlated with the truncation error of the numerical solver. In the simple case where the curvature is zero, a single sampling step is sufficient. This key finding motivates us to explore the change in precision as the NFE is reduced for both generation and inversion. We decreased the NFE from 15 to 3, observing the deviation from the results of a standard 50-step NFE. The results in Figure 1(right) confirm that the inversion trajectory can indeed tolerate a much lower NFE, which provides the foundational premise for our proposed method by offering a significant increase in processing speed and, crucially, by enabling efficient adversarial training.

3.2 FARI

Guided by the geometric intuition that the inversion trajectory is highly compressible, our method is simple and effective. In essence, our strategy is to first find a low-NFE approximation of the DDIM inversion (Song et al., 2021) trajectory and then perform adversarial training upon this condensed path to achieve both speed and robustness. We use distillation, a common technique for accelerating diffusion models, to achieve the first step. While standard trajectory distillation (Salimans & Ho, 2022; Song et al., 2023; Yin et al., 2024a) for the generation process often requires days or even dozens of GPU-days and substantial memory, the favorable geometric properties of the DDIM inversion trajectory allow us to obtain a reasonably accurate low-NFE estimate with minimal effort. Although this initial estimate has some error (see Figure 3), we find this trade-off is acceptable in exchange for the immense gains in robustness and speed, and it has a negligible effect on the performance of the downstream watermarking tasks.

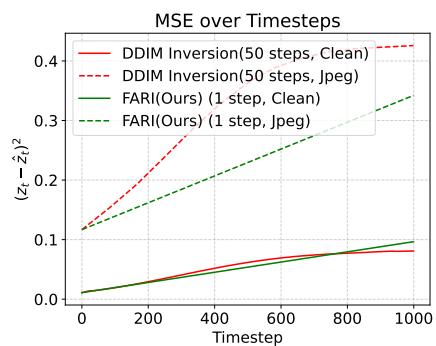


Figure 3: Inversion trajectory error of FARI and DDIM on clean and JPEG-compressed images.

For simplicity and efficiency, we do not explicitly separate the distillation and adversarial fine-tuning into two distinct stages, as they share a consistent optimization objective and converge rapidly. It is crucial to note our departure from common distillation practices for the generation process. We do not start with a real image and train our one-step model to mimic the output of a 50-step DDIM inversion. Instead, we sample a ground-truth Gaussian noise vector, perform the full generation process to obtain an image, and then learn a direct one-step mapping from this generated image back to the ground-truth initial noise. This approach avoids the performance ceiling imposed by the inherent inaccuracies of the 50-step DDIM inversion itself and is better aligned with the generative nature of the watermarking scenario.

Specifically, we fine-tune the denoising network of the diffusion model using Low-Rank Adaptation (LoRA) (Hu et al., 2022), a parameter-efficient fine-tuning technique that updates pretrained weight matrices through low-rank decomposition. Given a weight matrix $\mathbf{W}_0 \in \mathbb{R}^{d \times k}$, the update is represented as $\mathbf{W}_0 + \Delta \mathbf{W} = \mathbf{W}_0 + \mathbf{B} \mathbf{A}$, where $\mathbf{B} \in \mathbb{R}^{d \times r}$, $\mathbf{A} \in \mathbb{R}^{r \times k}$, and the rank $r \ll \min(d, k)$. During training, \mathbf{W}_0 is frozen, and gradient updates are applied only to \mathbf{A} and \mathbf{B} . The modified forward pass for an input \mathbf{z} becomes:

$$\mathbf{h} = \mathbf{W}_0 \mathbf{z} + \Delta \mathbf{W} \mathbf{z} = \mathbf{W}_0 \mathbf{z} + \mathbf{B} \mathbf{A} \mathbf{z}. \quad (8)$$

By decomposing the full-rank matrix into the product of two low-rank matrices, LoRA significantly reduces the number of trainable parameters, thereby lowering memory usage.

As illustrated in Figure 2, each training loop proceeds as follows. We randomly sample an initial noise vector $\mathbf{z}_T \sim \mathcal{N}(0, \mathbf{I})$ and a condition c from a dataset \mathcal{C} to generate an image X . During this generation phase, the LoRA branch is deactivated. The resulting image is then subjected to a distortion $D(\cdot)$ randomly selected from a predefined set \mathcal{T} , yielding a distorted image X^D , which is then encoded by the VAE encoder into a latent representation \mathbf{z}_0^D . Subsequently, the LoRA branch is activated to perform a one-step inversion, which reconstructs the noise according to the following formula:

$$\hat{\mathbf{z}}_T^D = \sqrt{\frac{\bar{\alpha}_T}{\bar{\alpha}_0}} \mathbf{z}_0^D + \left(\sqrt{1 - \bar{\alpha}_T} - \sqrt{\frac{\bar{\alpha}_T(1 - \bar{\alpha}_0)}{\bar{\alpha}_0}} \right) \epsilon_\theta(\mathbf{z}_0^D, 0, \emptyset; \psi). \quad (9)$$

Since $\bar{\alpha}_0 = 1$, this formula can also be equivalently written in the form of Eq.2:

$$\hat{\mathbf{z}}_T^D = \sqrt{\bar{\alpha}_T} \mathbf{z}_0^D + \sqrt{1 - \bar{\alpha}_T} \epsilon_\theta(\mathbf{z}_0^D, 0, \emptyset; \psi), \quad (10)$$

where ψ represents the LoRA parameters. For the inversion process, we use an unconditional setting (guidance scale = 1.0 and a null prompt). Prior works (Mokady et al., 2023; Wallace et al., 2023) have demonstrated that for standard DDIM inversion, an unconditional setting is often more precise because of the lack of invertibility in Classifier-Free Guidance (Ho & Salimans, 2022) (CFG). It is also important to note that we set the timestep $t = 0$ in the formula, rather than $t = T$ as expected from Eq. 7. This is because in a single-step scenario, the piecewise linear assumption is clearly violated. **Empirically, we find that any other small timestep value ($t \approx 0$) can achieve performance comparable to $t = 0$, providing a much better match for the latent \mathbf{z}_0^D , reducing the initial error and improving convergence.** Finally, our training objective is defined as:

$$\min_{\psi} \mathbb{E}_{\mathbf{z}_T, c \in \mathcal{C}, D \in \mathcal{T}} [\|\mathbf{z}_T - \hat{\mathbf{z}}_T^D\|_2^2] \quad (11)$$

Similarly, after training is complete, we deactivate the LoRA branch during denoising inference to preserve the original generation quality and enable it only for watermark extraction. This strategy is memory-efficient, eliminating the need to deploy two separate, largely identical denoisers. **It is worth noting that the LoRA component can be regarded as a plug-and-play enhancement module. Even when it is removed, DDIM in principle allows inversion with arbitrary step counts, but the error may be very large.** Further discussions, including details on the fine-tuned modules, training strategies, and hyperparameter selection, are provided in the ablation studies (Section4.4) and the Appendix E.3.

4 EXPERIMENTS

4.1 IMPLEMENTATION DETAILS

Diffusion Models. We selected Stable Diffusion v1.5 and v2.1 (Rombach et al., 2022) to cover the requirements of both the inversion baselines and the downstream watermarking task. For genera-

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

378

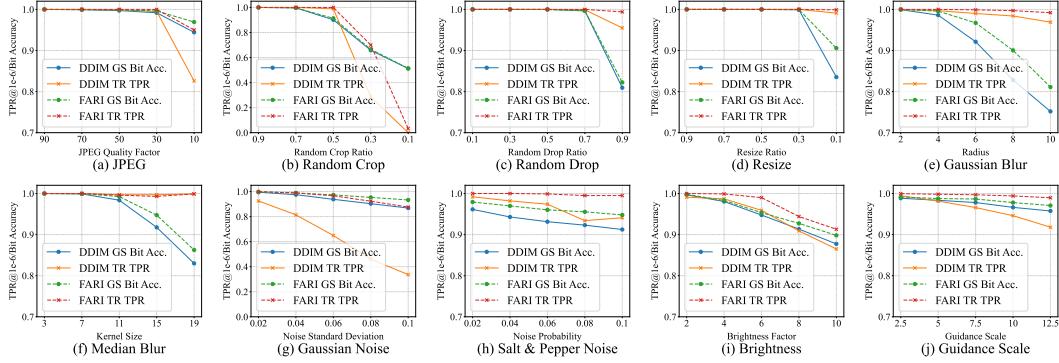
379
380
Table 1: Comparison of inversion methods on downstream watermarking methods under various
image distortions.

381	DM	Methods	NFE	Clean	Adv.	Jpeg	R.Crop	R.Drop	Resize	G.Blur	M.Blur	G.Noise	S&P	Bright
382 Bit Accuracy of Gaussian Shading Watermark														
383	SD v1.5	DDIM	50	0.9999	0.9777	0.9889	0.9781	0.9736	0.9975	0.9873	0.9983	0.9609	0.9354	0.9567
384			1	0.9999	0.9376	0.9703	0.8859	0.8808	0.9906	0.9585	0.9934	0.9398	0.9105	0.9085
385		EDICT	50	1.0000	0.9637	0.9786	0.9656	0.9568	0.9969	0.9807	0.9985	0.9390	0.9124	0.9450
386		BELM	50	0.9991	0.9465	0.9847	0.8960	0.8958	0.9939	0.9617	0.9956	0.9355	0.9275	0.9278
387		AMED [†]	2	1.0000	0.9656	0.9807	0.9528	0.9462	0.9970	0.9808	0.9989	0.9587	0.9346	0.9410
388		LCM-LoRA	2	0.9999	0.9541	0.9819	0.9352	0.9308	0.9924	0.9668	0.9955	0.9311	0.9030	0.9504
389	SD v2.1	DMD2	1	0.9988	0.9287	0.9760	0.8446	0.8241	0.9792	0.9209	0.9836	0.9252	0.9007	0.9336
390		FARI(Ours)	1	1.0000	0.9834	0.9935	0.9777	0.9761	0.9990	0.9957	0.9992	0.9836	0.9649	0.9612
391		DDIM	50	1.0000	0.9755	0.9892	0.9752	0.9669	0.9980	0.9860	0.9991	0.9590	0.9373	0.9447
392		EDICT	50	1.0000	0.9585	0.9755	0.8841	0.8637	0.9748	0.9173	0.9819	0.9280	0.9069	0.9284
393		BELM	50	0.9990	0.9411	0.9847	0.8923	0.8938	0.9933	0.9602	0.9952	0.9333	0.9269	0.9334
394		AMED [†]	2	1.0000	0.9662	0.9813	0.9559	0.9495	0.9970	0.9805	0.9989	0.9585	0.9357	0.9384
395	TPR@1e-3 of Tree-Ring Watermark	ExactDPM	> 150	1.0000	0.9670	0.9831	0.9675	0.9599	0.9974	0.9815	0.9987	0.9653	0.9241	0.9354
396		FARI(Ours)	1	1.0000	0.9824	0.9941	0.9771	0.9700	0.9992	0.9956	0.9994	0.9815	0.9659	0.9588
397		DDIM	50	1.000	0.949	0.989	1.000	1.000	0.999	0.996	1.000	0.636	0.946	0.972
398		EDICT	50	1.000	0.863	0.905	0.602	0.649	1.000	0.994	1.000	0.891	0.990	0.737
399		BELM	50	0.933	0.592	0.768	0.032	0.054	0.889	0.873	0.865	0.384	0.852	0.608
400		AMED [†]	2	1.000	0.909	0.939	0.947	0.936	0.999	0.995	0.999	0.618	0.912	0.835
401	SD v2.1	LCM-LoRA	2	1.000	0.875	0.914	0.996	0.991	0.999	0.987	0.999	0.331	0.812	0.850
402		DMD2	1	1.000	0.760	0.709	0.116	0.473	0.996	0.971	0.995	0.913	0.985	0.678
403		FARI(Ours)	1	1.000	0.997	1.000	1.000	1.000	1.000	1.000	1.000	0.980	1.000	0.992
404		DDIM	50	1.000	0.962	0.993	1.000	1.000	1.000	0.997	1.000	0.726	0.982	0.960
405		EDICT	50	1.000	0.896	0.896	0.709	0.845	1.000	0.993	0.999	0.903	0.991	0.729
406		BELM	50	0.882	0.543	0.721	0.001	0.000	0.803	0.787	0.801	0.417	0.812	0.541
407		AMED [†]	2	1.000	0.926	0.966	0.957	0.983	1.000	0.998	1.000	0.656	0.983	0.795
408		ExactDPM	> 150	1.000	0.906	0.991	0.571	0.847	1.000	0.999	1.000	0.835	0.992	0.915
409		FARI(Ours)	1	1.000	0.997	0.999	1.000	1.000	1.000	0.999	1.000	0.979	0.999	0.993

406 posit that this is because predicting the reverse ODE direction from a highly structured image is a
407 fundamentally different challenge than predicting it from pure noise.

409 4.3 GENERALIZATION

411 In this section, we explore FARI’s generalization ability to different generation conditions. By
412 default, we use the SD v2.1 model (Rombach et al., 2022) for these experiments and employ the
413 standard 50-step DDIM inversion (Song et al., 2021) as the baseline. Experiments with different
414 samplers and NFE in generation process can be found in Appendix E.
415



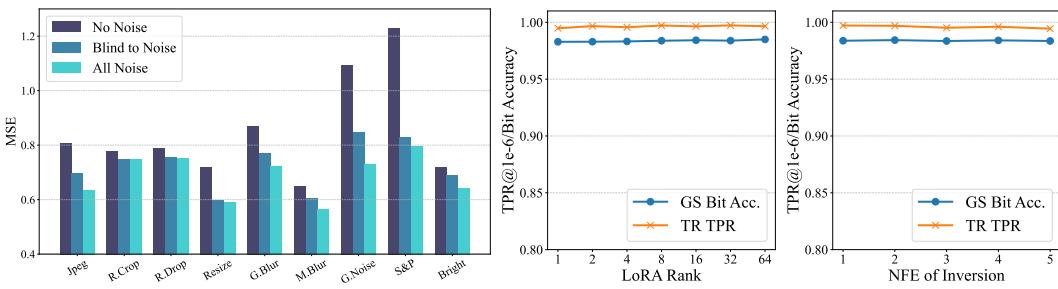
422 Figure 4: The results of experiments on various guidance scales and noise intensities.
423

424 **Guidance Scales.** Given diverse user preferences for prompt adherence, higher guidance scales
425 enforce the original prompt more strictly, while lower scales allow greater creative freedom. Our

432 experiments span a wide range of 2.5 to 12.5. As shown in Figure 4(j), FARI’s performance degrades
 433 only marginally under these settings.

434 **Noise Intensities.** To further test the robustness, we conduct experiments using different intensities
 435 of distortions. The results are shown in Figure 4(a-i). FARI consistently outperforms the stan-
 436 dard DDIM baseline, and its advantage becomes even more pronounced as the distortion intensity
 437 increases.

438 **Noise types.** Although our model is trained with a rich set of augmentations, real-world distortions
 439 may include types unseen during training. To evaluate FARI’s generalization ability against such
 440 unseen distortions, we compare the mean squared error (MSE) of the reconstructed noise under three
 441 distinct conditions: 1) **No Noise**, where the model is trained without any adversarial augmentations,
 442 serving as a baseline; 2) **Blind to Noise**, where FARI is trained on all distortions except for the one
 443 being tested; and 3) **All Noise**, our standard training procedure. The results, presented in Figure 5
 444 (left), show that FARI achieves a notable improvement in robustness even against distortions it has
 445 never encountered during training.



446 Figure 5: **Left:** Performance of FARI when trained under different noise settings, measured in mean
 447 squared error (MSE), where lower values are better. **Middle:** The effect of different LoRA ranks on
 448 FARI’s performance. **Right:** The impact of using a higher NFE for end-to-end training on FARI’s
 449 final performance.

461 4.4 ABLATION STUDY

462 In this section, we present ablation studies of FARI’s training settings.

463 **LoRA Ranks.** We experimented with different LoRA ranks. Figure 5 (middle) shows that even with
 464 a rank of 1, the performance is considerable. Increasing the rank yields only marginal performance
 465 gains, confirming that a low rank is sufficient for our method.

466 **More NFE for Training.** We also attempted end-to-end training on a multi-step (NFE > 1) in-
 467 version. As shown in Figure 5 (right), using more steps not only sacrifices speed but also fails to
 468 improve performance, instead causing a slight degradation. We posit two reasons for this: first, a
 469 single step may already be sufficient to accurately approximate the low-curvature inversion trajec-
 470 tory; second, the non-trivial accumulation of errors during multiple forward passes may harm the
 471 final reconstruction. This confirms that our choice of a one-step inversion is optimal.

475 5 DISCUSSION

476 **Extended Application and Future Work.** We also briefly investigated FARI’s performance on im-
 477 age reconstruction, given its nature as an inversion method. The specific results are available in the
 478 Appendix D. Furthermore, the reduction in NFE opens up another possibility. The adversarial re-
 479 moval of inversion-based watermarks has traditionally been difficult or extremely resource-intensive
 480 due to the need for gradient propagation through the entire multi-step inversion process (Müller et al.,
 481 2024). Our one-step method may alleviate this, as it dramatically shrinks the computational graph.
 482 This will be a direction for our future work.

483 **Limitations.** Despite its strong performance in speed and robustness, FARI has two main limita-
 484 tions. First, its sacrificed precision on clean image inversions makes it unsuitable for direct appli-
 485 cation in image editing (Hertz et al., 2022). Second, like all inversion-based watermarks, FARI is

486 dependent on ODE-based sampling (Song et al., 2021; Lu et al., 2022a) and will fail if an SDE
 487 sampler (Ho et al., 2020; Song et al., 2020) is employed.
 488

489 6 CONCLUSION 490

491 Starting from the demands of watermark extraction, this paper identifies the critical bottleneck as
 492 the cumulative error arising from external distortions in the transmission channel. This leads us to
 493 question the necessity of traditional high-NFE inversion methods; they are not only slow and inef-
 494 fective at mitigating these external errors but also computationally prohibit end-to-end adversarial
 495 training. Subsequently, we propose FARI, a framework built upon our key discovery of a geom-
 496 etric asymmetry: the inversion trajectory possesses a significantly lower curvature than its generation
 497 counterpart. This inherent compressibility allows us to efficiently distill the entire multi-step pro-
 498 cess into a single step. By doing so, FARI unlocks a robust adversarial training regime, creating
 499 an inverter that achieves state-of-the-art robustness and speed for practical, large-scale watermark
 500 verification.

501 502 REPRODUCIBILITY STATEMENT

503 The resources required to reproduce the experiments of this paper are provided in the supplementary
 504 materials. This includes the complete source code for our proposed method, FARI, along with de-
 505 tailed instructions for setting up the environment and running the training and evaluation scripts. The
 506 implementation details for all baseline methods and experimental settings are described in Section
 507 4.1 and Appendix B.

508 509 REFERENCES 510

511 Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational
 512 image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*, 2018.

513 Yuzhuo Chen, Zehua Ma, Han Fang, Weiming Zhang, and Nenghai Yu. Tag-wm: Tamper-
 514 aware generative image watermarking via diffusion inversion sensitivity. *arXiv preprint
 515 arXiv:2506.23484*, 2025.

516 Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with
 517 discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF
 518 conference on computer vision and pattern recognition*, pp. 7939–7948, 2020.

519 Hai Ci, Pei Yang, Yiren Song, and Mike Zheng Shou. Ringid: Rethinking tree-ring watermarking
 520 for enhanced multi-key identification. In *Proceedings of the European Conference on Computer
 521 Vision (ECCV)*, pp. 338–354, 2024.

522 Trung Dao, Thuan Hoang Nguyen, Thanh Le, Duc Vu, Khoi Nguyen, Cuong Pham, and Anh Tran.
 523 Swiftbrush v2: Make your one-step diffusion model better than its teacher. In *European Confer-
 524 ence on Computer Vision*, pp. 176–192. Springer, 2024.

525 Wenkai Dong, Song Xue, Xiaoyue Duan, and Shumin Han. Prompt tuning inversion for text-driven
 526 image editing using diffusion models. In *Proceedings of the IEEE/CVF International Conference
 527 on Computer Vision*, pp. 7430–7440, 2023.

528 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
 529 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers
 530 for high-resolution image synthesis. In *Forty-first international conference on machine learning*,
 2024.

531 Daniel Garibi, Or Patashnik, Andrey Voynov, Hadar Averbuch-Elor, and Daniel Cohen-Or. Renoise:
 532 Real image inversion through iterative noising. In *European Conference on Computer Vision*, pp.
 533 395–413. Springer, 2024.

534 Sam Gunn, Xuandong Zhao, and Dawn Song. An undetectable watermark for generative image
 535 models. *arXiv preprint arXiv:2410.07369*, 2024.

540 Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or.
 541 Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*,
 542 2022.

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

545 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances*
 546 *in Neural Information Processing Systems*, pp. 6840–6851. Curran Associates, Inc., 2020.

547 Seongmin Hong, Kyeonghyun Lee, Suh Yoon Jeon, Hyewon Bae, and Se Young Chun. On exact
 548 inversion of dpm-solvers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
 549 *Pattern Recognition*, pp. 7069–7078, 2024.

550 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
 551 Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

552 Runyi Hu, Jie Zhang, Yiming Li, Jiwei Li, Qing Guo, Han Qiu, and Tianwei Zhang. Videoshield:
 553 Regulating diffusion-based video generation models via watermarking, 2025. URL <https://arxiv.org/abs/2501.14195>.

554 Huayang Huang, Yu Wu, and Qian Wang. Robin: Robust and invisible watermarks for diffusion
 555 models with adversarial optimization. *Advances in Neural Information Processing Systems*, 37:
 556 3937–3963, 2024.

557 Xuan Ju, Ailing Zeng, Yuxuan Bian, Shaoteng Liu, and Qiang Xu. Direct inversion: Boosting
 558 diffusion-based editing with 3 lines of code. *arXiv preprint arXiv:2310.01506*, 2023.

559 Sangyun Lee, Beomsu Kim, and Jong Chul Ye. Minimizing trajectory curvature of ode-based
 560 generative models. In *International Conference on Machine Learning*, pp. 18957–18973. PMLR,
 561 2023.

562 Kecen Li, Zhicong Huang, Xinwen Hou, and Cheng Hong. Gaussmarker: Robust dual-domain
 563 watermark for diffusion models. *arXiv preprint arXiv:2506.11444*, 2025.

564 Ruibin Li, Ruihuang Li, Song Guo, and Lei Zhang. Source prompt disentangled inversion for
 565 boosting image editability with diffusion models. In *European Conference on Computer Vision*,
 566 pp. 404–421. Springer, 2024.

567 Haonan Lin, Yan Chen, Jiahao Wang, Wenbin An, Mengmeng Wang, Feng Tian, Yong Liu, Guang
 568 Dai, Jingdong Wang, and Qianying Wang. Schedule your edit: A simple yet effective diffu-
 569 sion noise schedule for image editing. *Advances in Neural Information Processing Systems*, 37:
 570 115712–115756, 2024.

571 Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr
 572 Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer*
 573 *vision–ECCV 2014: 13th European conference, zurich, Switzerland, September 6–12, 2014, pro-*
 574 *ceedings, part v 13*, pp. 740–755. Springer, 2014.

575 Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on
 576 manifolds. *arXiv preprint arXiv:2202.09778*, 2022.

577 Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast
 578 ode solver for diffusion probabilistic model sampling in around 10 steps. In *Advances in Neural*
 579 *Information Processing Systems*, volume 35, pp. 5775–5787, 2022a.

580 Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast
 581 solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*,
 582 2022b.

583 Nils Lukas, Abdulrahman Diaa, Lucas Fenaux, and Florian Kerschbaum. Leveraging optimization
 584 for adaptive attacks on image watermarks. *arXiv preprint arXiv:2309.16952*, 2023.

594 Simian Luo, Yiqin Tan, Suraj Patil, Daniel Gu, Patrick von Platen, Apolinário Passos, Longbo
 595 Huang, Jian Li, and Hang Zhao. Lcm-lora: A universal stable-diffusion acceleration module.
 596 *arXiv preprint arXiv:2311.05556*, 2023.

597

598 Daiki Miyake, Akihiro Iohara, Yu Saito, and Toshiyuki Tanaka. Negative-prompt inversion: Fast
 599 image inversion for editing with text-guided diffusion models. In *2025 IEEE/CVF Winter Con-
 600 ference on Applications of Computer Vision (WACV)*, pp. 2063–2072. IEEE, 2025.

601 Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for
 602 editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF conference
 603 on computer vision and pattern recognition*, pp. 6038–6047, 2023.

604

605 Andreas Müller, Denis Lukovnikov, Jonas Thietke, Asja Fischer, and Erwin Quiring. Black-box
 606 forgery attacks on semantic watermarks for diffusion models. *arXiv preprint arXiv:2412.03283*,
 607 2024.

608 Trong-Tung Nguyen, Quang Nguyen, Khoi Nguyen, Anh Tran, and Cuong Pham. Swiftedit: Light-
 609 ning fast text-guided image editing via one-step diffusion. In *Proceedings of the Computer Vision
 610 and Pattern Recognition Conference*, pp. 21492–21501, 2025.

611

612 Zhihong Pan, Riccardo Gherardi, Xiufeng Xie, and Stephen Huang. Effective real image editing
 613 with accelerated iterative diffusion inversion. In *Proceedings of the IEEE/CVF International
 614 Conference on Computer Vision*, pp. 15912–15921, 2023.

615 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe
 616 Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image
 617 synthesis. *arXiv preprint arXiv:2307.01952*, 2023.

618

619 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 620 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Con-
 621 ference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, June 2022.

622

623 Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv
 preprint arXiv:2202.00512*, 2022.

624

625 Dvir Samuel, Barak Meiri, Haggai Maron, Yoad Tewel, Nir Darshan, Shai Avidan, Gal Chechik, and
 626 Rami Ben-Ari. Lightning-fast image inversion and editing for text-to-image diffusion models.
 627 *arXiv preprint arXiv:2312.12540*, 2023.

628

629 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Inter-
 630 national Conference on Learning Representations*, 2021.

631

632 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 633 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint
 634 arXiv:2011.13456*, 2020.

635

636 Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint
 637 arXiv:2303.01469*, 2023.

638

639 Łukasz Staniszewski, Łukasz Kuciński, and Kamil Deja. There and back again: On the relation
 640 between noise and image inversions in diffusion models. *arXiv preprint arXiv:2410.23530*, 2024.

641

642 Bram Wallace, Akash Gokul, and Nikhil Naik. Edict: Exact diffusion inversion via coupled trans-
 643 formations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recog-
 644 nition*, pp. 22532–22541, 2023.

645

646 Fangyikang Wang, Hubery Yin, Yue-Jiang Dong, Huminhao Zhu, Hanbin Zhao, Hui Qian, Chen
 647 Li, et al. Belm: Bidirectional explicit linear multi-step sampler for exact inversion in diffusion
 648 models. *Advances in Neural Information Processing Systems*, 37:46118–46159, 2024.

649

650 Yuxin Wen, John Kirchenbauer, Jonas Geiping, and Tom Goldstein. Tree-ring watermarks: Fin-
 651 gerprints for diffusion images that are invisible and robust. *arXiv preprint arXiv:2305.20030*,
 652 2023.

648 Zijin Yang, Kai Zeng, Kejiang Chen, Han Fang, Weiming Zhang, and Nenghai Yu. Gaussian shading-
 649 ing: Provable performance-lossless image watermarking for diffusion models. In *Proceedings*
 650 *of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 12162–
 651 12171, June 2024.

652 Zijin Yang, Xin Zhang, Kejiang Chen, Kai Zeng, Qiyi Yao, Han Fang, Weiming Zhang, and Nenghai
 653 Yu. Gaussian shading++: Rethinking the realistic deployment challenge of performance-lossless
 654 image watermark for diffusion models. *arXiv preprint arXiv:2504.15026*, 2025.

655 Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and
 656 Bill Freeman. Improved distribution matching distillation for fast image synthesis. *Advances in*
 657 *neural information processing systems*, 37:47455–47487, 2024a.

658 Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman,
 659 and Taesung Park. One-step diffusion with distribution matching distillation. In *Proceedings of*
 660 *the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6613–6623, 2024b.

661 Guoqiang Zhang, Jonathan P Lewis, and W Bastiaan Kleijn. Exact diffusion inversion via bidi-
 662 rectional integration approximation. In *European Conference on Computer Vision*, pp. 19–36.
 663 Springer, 2024a.

664 Lijun Zhang, Xiao Liu, Antoni V Martin, Cindy X Bearfield, Yuriy Brun, and Hui Guan. Attack-
 665 resilient image watermarking using stable diffusion. *Advances in Neural Information Processing*
 666 *Systems*, 37:38480–38507, 2024b.

667 Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator.
 668 *arXiv preprint arXiv:2204.13902*, 2022.

669 Wenliang Zhao, Lujia Bai, Yongming Rao, Jie Zhou, and Jiwen Lu. Unipc: A unified predictor-
 670 corrector framework for fast sampling of diffusion models. *Advances in Neural Information*
 671 *Processing Systems*, 36:49842–49869, 2023.

672 Xuandong Zhao, Kexun Zhang, Zihao Su, Saastha Vasan, Ilya Grishchenko, Christopher Kruegel,
 673 Giovanni Vigna, Yu-Xiang Wang, and Lei Li. Invisible image watermarks are provably removable
 674 using generative ai. *Advances in neural information processing systems*, 37:8643–8672, 2024.

675 Zhenyu Zhou, Defang Chen, Can Wang, and Chun Chen. Fast ode-based sampling for diffusion
 676 models in around 5 steps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
 677 *Pattern Recognition*, pp. 7777–7786, 2024a.

678 Zhenyu Zhou, Defang Chen, Can Wang, Chun Chen, and Siwei Lyu. Simple and fast distillation of
 679 diffusion models. *Advances in Neural Information Processing Systems*, 37:40831–40860, 2024b.

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 A DISCUSSION ABOUT RELATED WORKS

704 The body of work related to our method is extensive, primarily encompassing inversion techniques
 705 and diffusion model acceleration. We discuss these in separate categories below.

707 A.1 INVERSION TECHNIQUES

709 A significant portion of inversion research focuses on improving the quality of image editing by
 710 enforcing trajectory symmetry. NTI (Mokady et al., 2023) first performs a standard DDIM inver-
 711 sion and then optimizes a null-text embedding for each step. This embedding sequence is then used
 712 during regeneration to ensure the new path closely follows the reverse of the inversion trajectory,
 713 achieving high-fidelity reconstruction. PTI (Dong et al., 2023) extends this by linking the embed-
 714 ding to the target prompt to enhance editing quality. Furthermore, NPI (Miyake et al., 2025) replaces
 715 the optimized null-text sequence with the prompt embedding itself, eliminating the need for opti-
 716 mization and greatly improving efficiency. Direct Inversion (Ju et al., 2023) also follows this path
 717 by decoupling the source and target diffusion branches. Notably, these methods do not actually
 718 change the result of the initial inversion; their focus is on the reconstruction phase. Therefore, their
 719 performance in a watermarking scenario is nearly identical to that of standard DDIM inversion, and
 720 we do not include them in our quantitative comparisons.

721 Another representative class of inversion methods, including AIDI (Pan et al., 2023), SPDInv (Li
 722 et al., 2024), ReNoise (Garibi et al., 2024), GNRI (Samuel et al., 2023), and ExactDPM (Hong
 723 et al., 2024), directly optimizes the inversion trajectory using fixed-point iteration or gradient
 724 descent. While these can indeed reduce the inversion error, they are often unstable. This instability is
 725 exacerbated by the initial value offset common in watermarking, which can cause the optimization
 726 to proceed in an incorrect direction or fail entirely. We found that while ExactDPM is the slowest, it
 727 is the most stable among them and also targets watermarking as a downstream task. Consequently,
 728 we select it as a representative of this category for comparison.

729 A third class of methods addresses the inherent lack of invertibility in the sampling process itself.
 730 Techniques like EDICT (Wallace et al., 2023), BDIA (Zhang et al., 2024a), and BELM (Wang
 731 et al., 2024) modify the sampler to achieve a smaller theoretical error bound. However, their design
 732 carries an implicit assumption that the entire process is free from external error. This holds true
 733 for reconstruction but is clearly violated in watermarking, where the image is subject to various
 734 distortions after generation that can push it out of the method’s convergence domain. Furthermore,
 735 these methods can be more sensitive when handling images generated with a high guidance scale.

736 Recently, an interesting work in this area is SwiftEdit (Nguyen et al., 2025), which also achieves
 737 one-step inversion via model fine-tuning. However, it still considers the problem from an editing
 738 perspective, focusing on adapting to real images and enabling subsequent edits. While it achieves
 739 fast, high-quality results on editing tasks, it relies on base models that are already one-step genera-
 740 tors (e.g., SwiftBrush v2 (Dao et al., 2024)) and has not demonstrated generalization to multi-step
 741 generators. This implies the trajectory it learns to fit is already a straight line, which is a simpler
 742 task. Additionally, its code and weights are not publicly available, precluding a direct comparison.

743 A.2 DIFFUSION MODEL ACCELERATION

745 Given the lack of fast and stable inversion methods, we also include several diffusion model acceler-
 746 ation techniques in our baseline comparison. For the task of acceleration, we focus on methods that
 747 retain capabilities similar to the original model, rather than fully retraining a new one for faster gen-
 748 eration. These methods can be divided into two categories. The first is based on higher-order solvers
 749 with lower truncation error, such as DPMsampler (Lu et al., 2022a). However, their performance ceiling
 750 is limited in extreme few-step scenarios. A noteworthy exception is the AMED-Solver (Zhou
 751 et al., 2024a), which is based on the mean value theorem for vector fields. It trains an additional
 752 small model to predict the median point of the trajectory, whose velocity can be used as the average
 753 velocity for the entire path, enabling sampling in as few as two steps.

754 The second category is distillation, where a student model learns to replicate the output of multiple
 755 teacher steps in a single step. This, however, often requires substantial training resources and time,
 which we argue is excessive if used solely for inversion. We select two representative baselines from

756 this rich field. LCM-LoRA (Luo et al., 2023) is similar to our method in that it stores the distilled
 757 parameters in a LoRA (Hu et al., 2022) module. However, its purpose is different: it is designed
 758 to adapt to various user-personalized models, enabling accelerated sampling without requiring a
 759 separate distillation for each fine-tuned model. DMD2 (Yin et al., 2024a), on the other hand, uses
 760 distribution matching distillation. It does not directly learn the teacher’s output but rather its target
 761 distribution, allowing the student’s performance to surpass the ceiling of the original teacher model.
 762 Another relevant work is SFD (Zhou et al., 2024b), which observes the smooth modification of the
 763 gradient field and, inspired by this, focuses training resources on essential timesteps. While known
 764 for its fast training speed for a distillation method, it is still slower to train than FARI and its one-step
 765 performance is inferior to our chosen baseline, DMD2. Our results show that distillation methods
 766 do not perform exceptionally well on the inversion task. As we posited, while the prediction target
 767 is similar, the input is fundamentally different: predicting a direction from a highly structured image
 768 is a distinct challenge compared with predicting it from pure noise.

769 B DETAILED EXPERIMENTS SETTINGS

770 In this section, we provide the specific details of our experimental setup.

771 **772 Models and Generation.** We conduct experiments on Stable Diffusion (SD)
 773 v1.5 and v2.1 (Rombach et al., 2022). The generation process utilizes the
 774 diffusers library in Python, with pretrained weights sourced from the Hug-
 775 ging Face Hub repositories `runwayml/stable-diffusion-v1-5` and
 776 `stabilityai/stable-diffusion-2-1-base`, respectively. All images are gener-
 777 ated at a resolution of 512×512 pixels. For testing, we use 50-step DDIM sampling with a fixed
 778 guidance scale of 7.5, which are common settings for the downstream watermarking tasks.

779 **780 Watermarking Methods.** For both the Tree-Ring (Wen et al., 2023) and Gaussian Shading (Yang
 781 et al., 2024) watermarks, we use the official open-source code provided by the authors on GitHub.
 782 We use Tree-Ring in its `rand` mode, with the watermark embedded in the fourth channel of the
 783 latent space. For Gaussian Shading, we adopt the default settings ($f_{ch} = 4, f_h = 8, f_w = 8$, for a
 784 256-bit capacity) and use a stream cipher to encrypt the watermark message.

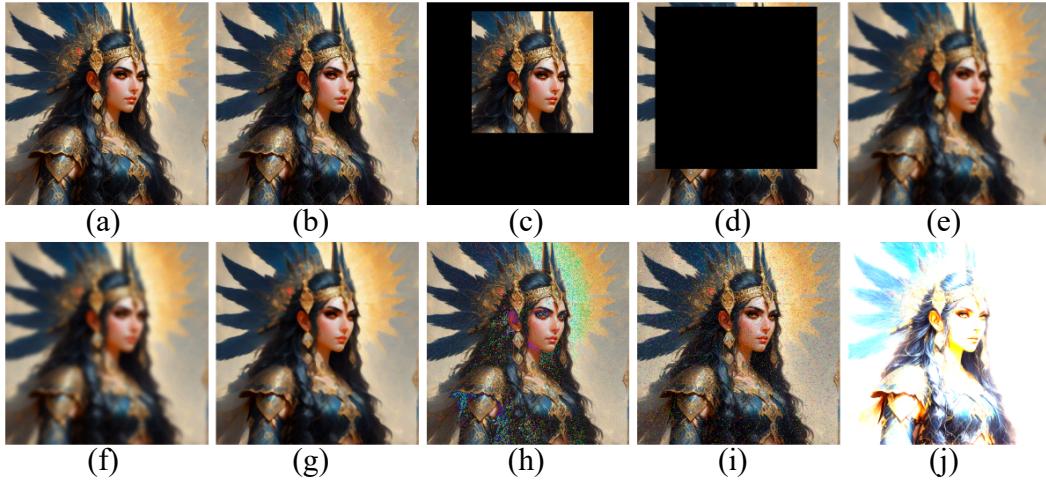
785 **786 Baselines.** For DDIM (Song et al., 2021), we use our own implementation, which is equiv-
 787 alent to using the `DDIMInverseScheduler` from `diffusers`. It is critical to note that
 788 the `timestep_spacing` must be set to "trailing" instead of the default "leading".
 789 The "leading" setting introduces a significant error at low NFEs because it sets the first
 790 `prev_timestep` to 0. Since the scheduler calculates the current `timestep` by subtracting a
 791 step size from `prev_timestep`, this can result in a negative timestep, which is then clipped to 0.
 792 This effectively makes the first step a non-operation (from $t = 0$ to $t = 0$). The "trailing"
 793 setting avoids this error. While both settings yield similar results at high NFEs, the difference is
 794 substantial at very few steps, especially when NFE=1. Furthermore, our one-step DDIM inversion
 795 baseline, similar to FARI, uses $t = 0$ when predicting epsilon, as using $t = T$ introduces a large
 796 error.

797 For EDICT (Wallace et al., 2023) and BELM (Wang et al., 2024), we use the implementation pro-
 798 vided by BELM. The NFE is set to 50 steps, and EDICT’s hyperparameter p is set to 0.93 as recom-
 799 mended in the original paper. For ExactDPM (Hong et al., 2024), we use the official implementa-
 800 tion with the Backward Euler method for inversion. To ensure a fair and efficient comparison, we disable
 801 the additional decoder inversion. We set its inversion step count to 10, but note that it requires an
 802 iterative optimization process that often takes over 150 iterations, depending on the distortion type
 803 and intensity.

804 As the official weights for AMED-Solver (Zhou et al., 2024a) were not provided for our target
 805 models, we retrained it ourselves. Crucially, we performed an end-to-end adversarial training for
 806 the inversion task, similar to our FARI training, and aligned its parameter count with FARI. Since
 807 it requires a minimum of two steps for sampling, we set its NFE to 2. For LCM-LoRA (Luo et al.,
 808 2023) and DMD2 (Yin et al., 2024a), we use the officially provided weights. We set LCM-LoRA
 809 to use two steps (within its allowed 2-8 range) and use a single step for DMD2, which has one-step
 sampling capabilities.

810
Training and Evaluation. For FARI’s training, we use a batch size of 4 with the Adam opti-
 811 izer and a learning rate of 1e-4. When retraining the AMED-Solver from scratch, we adjusted the
 812 learning rate to 1e-3. For our training dataset, we use prompts associated with the MS-COCO-2017
 813 dataset (Lin et al., 2014); while the original dataset lacks prompts, we utilize the captions provided
 814 in the official Tree-Ring repository. For testing, we use the test split of the Stable-Diffusion-Prompts
 815 (SDP) dataset.

816
Distortions. The types of distortions used in our training and testing are illustrated in Figure 6.
 817 Geometric distortions, such as rotation and scaling, were excluded, as we consider their handling
 818 to be more related to the intrinsic robustness mechanism of the watermark itself rather than the
 819 inversion process.



838
 Figure 6: Visualization of the distortion set used in our experiments. (a) Clean image or identity
 839 transformation. (b) JPEG, $QF = 25$. (c) 60% area Random Crop (R.Crop). (d) 80% area Random
 840 Drop (R.Drop). (e) 25% Resize and restore (Resize). (f) Gaussian Blur, $r = 4$ (G.Blur). (g) Median
 841 Blur, $k = 7$ (M.Blur). (h) Gaussian Noise, $\mu = 0$, $\sigma = 0.05$ (G.Noise). (i) Salt and Pepper Noise, $p =$
 842 0.05 (S&P). (j) Brightness, $factor = 6$ (Bright).

845 C DETAILS ABOUT THE CURVATURE EVALUATION.

846
 In Section 3.1, we measure the curvature of the generation and inversion trajectories as a function
 847 of the timestep. The experiment is conducted on the Stable Diffusion v2.1 model (Rombach et al.,
 848 2022), generating 100 images. Both the generation and inversion processes are set to be uncondi-
 849 tional to eliminate interference from other factors. The curvature is approximated using a formal
 850 definition of discrete curvature.

851 Specifically, for a trajectory of latent variables $\{\mathbf{x}_t\}_{t=0}^T$, we consider three consecutive points \mathbf{x}_{t+1} ,
 852 \mathbf{x}_t , and \mathbf{x}_{t-1} . We first define the two corresponding velocity vectors as $\mathbf{v}_t = \mathbf{x}_{t-1} - \mathbf{x}_t$ and
 853 $\mathbf{v}_{t+1} = \mathbf{x}_t - \mathbf{x}_{t+1}$. We then approximate the arc length s_t as:

$$854 \quad s_t = \|\mathbf{v}_{t+1}\|_2 + \|\mathbf{v}_t\|_2 \quad (12)$$

855 The discrete curvature κ_t at timestep t is then defined as the angle θ_t between the velocity vectors,
 856 divided by the arc length s_t :

$$857 \quad \kappa_t = \frac{\theta_t}{s_t} \quad (13)$$

858 where the angle θ_t is given by:

$$859 \quad \theta_t = \arccos \left(\frac{\mathbf{v}_{t+1} \cdot \mathbf{v}_t}{\|\mathbf{v}_{t+1}\|_2 \|\mathbf{v}_t\|_2} \right) \quad (14)$$

We also conducted two additional experiments, with the results shown in the Figure 7. For conditional generation, we use prompts from the Stable-Diffusion-Prompts dataset. All other settings remain identical to our main experiments. First, we used conditional generation to create an image, and then compared the curvature of the original generation trajectory against two types of inversion: one that was conditionally-aligned and one that was unconditional. The results show that the curvature of both inversion trajectories is consistently lower than that of the generation trajectory. It is important to note, however, that lower curvature does not necessarily equate to higher inversion accuracy. Second, we measured the curvature of a regeneration trajectory (from the inverted noise). The results indicate that the regenerated path also has a lower curvature than the initial generation path. This further validates our conclusion that the residual low-frequency semantic information in the reconstructed noise partly helps to reduce trajectory curvature.

To strengthen the generalizability of our findings, we conduct additional on the COCO dataset and using Stable Diffusion v3.5 Medium (SD v3.5M) (Esser et al., 2024). The results are presented in the Figure 8.

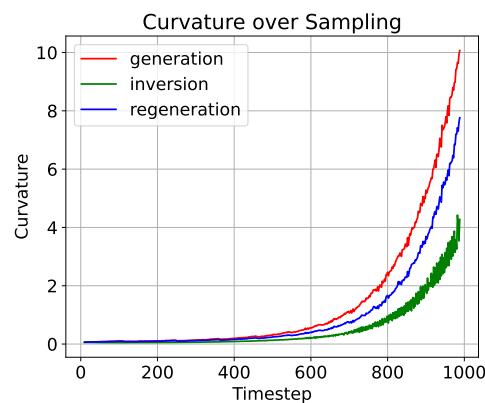
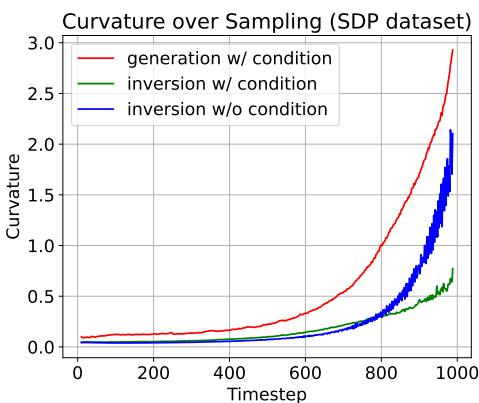


Figure 7: **Left:** The curvature of the conditional generation, unconditional inversion, and conditional inversion trajectories over the diffusion timesteps. **Right:** The curvature of the generation, inversion, and regeneration trajectories over the diffusion timesteps.

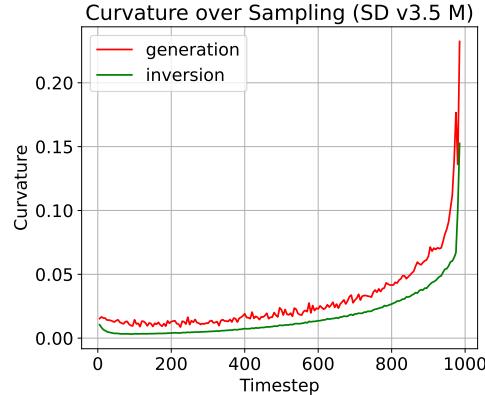
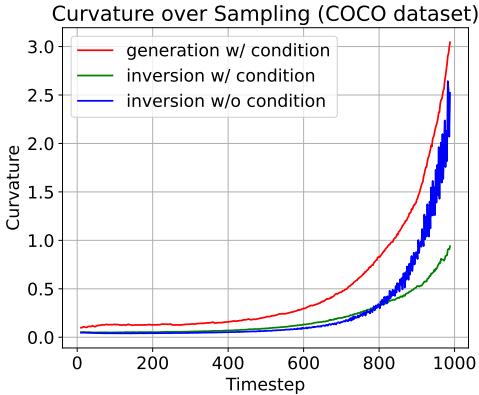


Figure 8: **Left:** The curvature of the the conditional generation, unconditional inversion, and conditional inversion trajectories over the diffusion timesteps, measured using SD v2.1 on the COCO dataset . **Right:** The curvature of the generation and inversion trajectories over the diffusion timesteps, measured using SD v3.5M.

D EXTENDED APPLICATION ON IMAGE RECONSTRUCTION AND EDITING

Although FARI is specifically designed for watermark extraction, its remarkable speed prompted us to investigate its potential for image reconstruction and editing. We therefore conducted an exper-

iment comparing the peak signal-to-noise ratio (PSNR) of images reconstructed by FARI against those reconstructed by DDIM (Song et al., 2021), EDICT (Wallace et al., 2023), BELM (Wang et al., 2024), and ExactDPM (Hong et al., 2024). The results are presented in Table 2. While disabling adversarial training ($FARI^{Cln}$) boosts FARI’s precision on clean inversions, a gap remains when compared to leading editing-oriented methods. Regarding image editing, many techniques require attention map or feature sharing during the intermediate steps of the inversion and regeneration process to ensure high consistency. As our one-step method lacks these intermediate steps for intervention, it is not directly compatible with mainstream plug-and-play diffusion-based editing frameworks (Hertz et al., 2022). We hope our work can serve as inspiration for subsequent research in this area.

Table 2: The performance of inversion methods on image reconstruction.

	DDIM	EDICT	BELM	ExactDPM	FARI	$FARI^{Cln}$
PSNR	17.33	25.26	24.04	20.02	10.61	14.60

E MORE EXPERIMENTAL RESULTS

E.1 DISCUSSION ABOUT TRAINING TIME AND INFERENCE TIME

Training Time Figure 9 illustrates the training dynamics of FARI, showing the exponentially weighted moving average (EWMA) of the loss curve alongside the performance on the Gaussian Shading (Yang et al., 2024) and Tree-Ring (Wen et al., 2023) watermark tasks. The entire training process, including the online construction of data pairs, takes approximately 70 minutes on a single NVIDIA RTX A6000 GPU. Notably, FARI surpasses the performance of the 50-step DDIM baseline on both watermarking tasks around the 300th training step (approximately 20 minutes), highlighting the efficiency of our method.

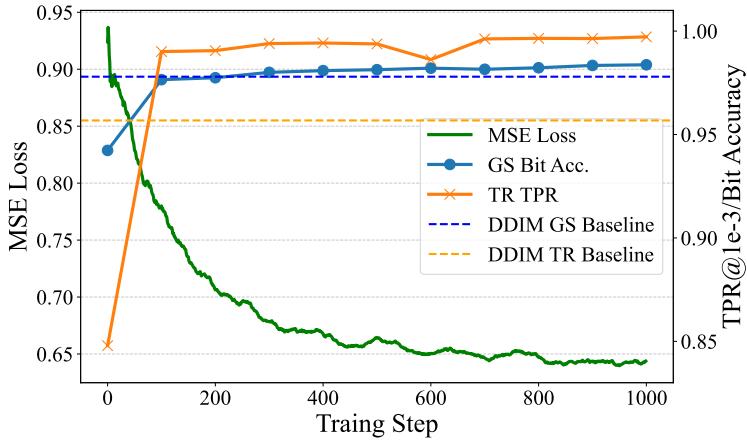


Figure 9: The loss tendency and tested performance on downstream watermarking tasks over the training step.

Inference Time The use of a LoRA (Hu et al., 2022) modules affects the denoiser’s inference efficiency, with the average time per function evaluation on an NVIDIA RTX A6000 GPU increasing from 0.0247 s to 0.0358 s. While this introduces latency, the total time remains less than that of a two-step inversion. To eliminate this overhead during inference, a practical solution is to merge the LoRA weights by $\mathbf{W}' = \mathbf{W}_0 + \Delta\mathbf{W} = \mathbf{W}_0 + \mathbf{BA}$. This can be achieved by creating an extra fused copy of the LoRA-injected module, allowing the system to switch between the original module for image generation and the fine-tuned, robust module for one-step watermark inversion. Alternatively, for systems with sufficient VRAM, a simpler but more memory-intensive approach is to load an entirely separate model with the weights pre-merged.

972 E.2 MORE RESULTS ABOUT GENERALIZATION
973

974 **Samplers.** Users may employ samplers different from the one used during training. To test
975 FARI’s generalization to this possibility, we conducted sampling using a variety of common ODE
976 solvers (Song et al., 2021; Zhao et al., 2023; Liu et al., 2022; Zhang & Chen, 2022; Lu et al., 2022a).
977 The results, shown in Table 3, indicate that FARI outperforms the DDIM inversion baseline across
978 all tested samplers.

980
981 Table 3: Generalization to different samplers. For each cell, the values represent TR TPR / GS Bit
982 Acc.

	DDIM (Song et al., 2021)	UniPC (Zhao et al., 2023)	PNND (Liu et al., 2022)	DEIS (Zhang & Chen, 2022)	DPMsampler (Lu et al., 2022a)
DDIM	0.966 / 0.9780	0.886 / 0.9550	0.967 / 0.9757	0.963 / 0.9744	0.884 / 0.9549
FARI	0.997 / 0.9841	0.951 / 0.9637	0.997 / 0.9846	0.997 / 0.9826	0.951 / 0.9664

983
984
985
986
987
988 **Inference Steps of Generation.** To optimize the training speed, we used a fixed 20 steps for
989 image generation during our training process. While this may be a low number of inference steps
990 for the selected models, we present results with a greater number of generation steps in Table 4. The
991 results demonstrate that FARI, despite being trained under these efficient conditions, generalizes
992 well to higher-quality generation settings and consistently outperforms the baseline.

993
994
995 Table 4: Generalization to different NFEs of generation. For each cell, the values represent TR TPR
996 / GS Bit Acc.

	25	50	100
DDIM	0.965 / 0.9752	0.966 / 0.9780	0.964 / 0.9768
FARI	0.997 / 0.9826	0.997 / 0.9841	0.997 / 0.9843

1001
1002
1003
1004 **Other Models.** We evaluate FARI’s performance on SD v3.5M (Esser et al., 2024) and SDXL-
1005 Turbo (Podell et al., 2023). Since SD v3.5M does not support DDIM sampler, we implement inver-
1006 sion for this model via naive Euler Method. SDXL-Turbo inherently supports one-step generation,
1007 so we compare the performance of one-step DDIM inversion against FARI.

1008 Given that SD v3.5 has 16 latent channels, we set $f_{ch} = 2$ and $f_h = f_w = 8$ for Gaussian Shading,
1009 resulting in a capacity of 512 bits. For Tree-Ring Watermark, we use only the last four channels
1010 for watermark embedding. The corresponding results are shown in Table 5 and Table 6, demon-
1011 strating FARI’s strong generalization capability. FARI exhibits substantial improvements on both
1012 downstream watermarking approaches, especially for Tree-Ring. In the presence of distortions,
1013 naive inversion fails to maintain the viability of Tree-Ring watermarking, demonstrating that FARI
1014 significantly broadens the practical applicability of inversion-based watermarking techniques.

1015
1016
1017 Table 5: Performance of naive inversion and FARI on SD v3.5M.

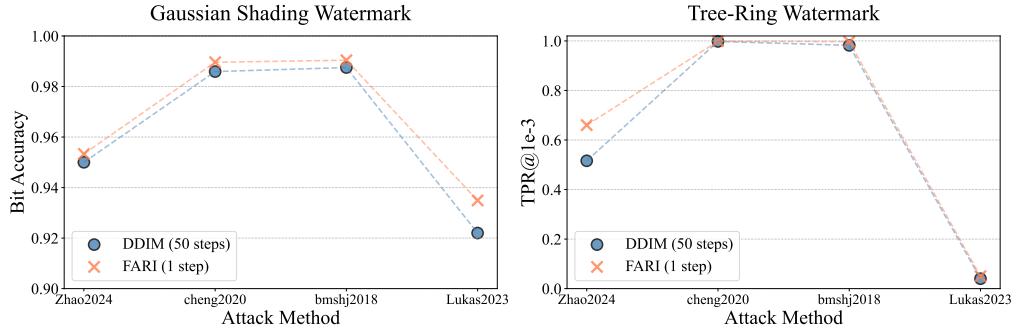
Methods	NFE	Clean	Adv.	JPEG	R.Crop	R.Drop	Resize	G.Blur	M.Blur	G.Noise	S&P	Bright
Bit Accuracy of Gaussian Shading Watermark												
Naive	20	0.9994	0.9178	0.9440	0.9654	0.9403	0.9437	0.8728	0.9609	0.7778	0.8736	0.9818
FARI	1	0.9995	0.9515	0.9671	0.9713	0.9469	0.9756	0.9384	0.9852	0.8765	0.9191	0.9833
TPR@1e-3 of Tree-Ring Watermark												
Naive	20	0.958	0.161	0.020	0.364	0.543	0.000	0.000	0.001	0.009	0.003	0.512
FARI	1	0.998	0.978	0.994	0.998	0.995	0.995	0.952	0.996	0.917	0.968	0.985

1026

1027 Table 6: Performance of naive inversion and FARI on SDXL-Turbo.

Methods	NFE	Clean	Adv.	JPEG	R.Crop	R.Drop	Resize	G.Blur	M.Blur	G.Noise	S&P	Bright
Bit Accuracy of Gaussian Shading Watermark												
DDIM	1	0.9755	0.8576	0.9142	0.7449	0.7552	0.9370	0.8841	0.9346	0.8935	0.8229	0.8315
FARI	1	0.9998	0.9511	0.9903	0.9115	0.9161	0.9871	0.9370	0.9922	0.9801	0.9354	0.9102
TPR@1e-3 of Tree-Ring Watermark												
DDIM	1	0.561	0.201	0.672	0.000	0.000	0.065	0.001	0.221	0.412	0.271	0.163
FARI	1	0.998	0.855	0.967	0.743	0.933	0.893	0.683	0.928	0.949	0.837	0.763

1034



1035 Figure 10: Performance of DDIM inversion and FARI under different attacks.

1046

1047 **More Datasets.** To further validate generalization across datasets, we additionally select 1,000
 1048 prompts from a new dataset (DiffusionDB²). Combined with the existing SDP and COCO datasets,
 1049 we conduct cross-dataset validation. The results are shown in Table 7, presented in the format of TR
 1050 TPR / GS Bit Acc., with 50-step DDIM as the baseline.

1051

1052

1053

1054

1055 Table 7: Cross-dataset generalization results. Performance is reported as TR TPR / GS Bit Acc.

Training Dataset \ Test Dataset	SDP	COCO	DiffusionDB
DDIM Baseline (50-step)	0.966 / 0.9780	0.964 / 0.9785	0.962 / 0.9763
SDP	—	0.997 / 0.9852	0.995 / 0.9839
COCO	0.997 / 0.9841	—	0.997 / 0.9846
DiffusionDB	0.995 / 0.9832	0.997 / 0.9861	—

1062

1063

1064

1065 **Removal Attacks.** To evaluate FARI’s robustness against watermark removal attacks, we test
 1066 three regeneration attacks and one adversarial optimization attack. We evaluate three different regen-
 1067 eration attacks. For Zhao et al. (2024), we set the noise steps to 300. For Cheng et al. (2020) and
 1068 Ballé et al. (2018), we set the quality parameter to 3 (representing the highest attack strength). For
 1069 the adversarial optimization attack (Lukas et al., 2023), we use the following settings: $\epsilon = 4/255$,
 1070 Adam optimizer, learning rate of 0.01, and 5 optimization steps per image, all of which are
 1071 consistent with the default configuration. Note that the attack implementation does not involve Gaus-
 1072 sian Shading (GS). Since the GS verification process is non-differentiable and relies on sign-based
 1073 verification, we optimized by minimizing the MSE distance between the reconstructed noise after
 1074 optimization and the negative of the original noise. This can cause more elements in the recon-
 1075 structed noise to flip their signs and proved to be effective (otherwise, the bitwise accuracy will be
 1076 1.0000, which means the attack is useless). The corresponding results are shown in Figure 10. FARI
 1077 provides improvements in robustness. However, the ability to resist attacks remains largely depen-
 1078 dent on the underlying design of the base watermarking method, as FARI serves as a plug-and-play
 1079 inversion approach to enhance their detection performance.

²<https://huggingface.co/datasets/poloclub/diffusiondb>

1080 E.3 MORE ABLATION STUDIES
1081

1082 **Trained Modules.** We experimented with training different modules of the denoiser. By default,
1083 we only fine-tune the linear layers within the attention-related modules. In an *Extended LoRA FT*
1084 setting, we fine-tune all linear and convolutional layers using LoRA. We also tested full fine-tuning
1085 as an upper bound. The results are shown in Table 8. Although training more modules yields a
1086 marginal improvement, the increase is not significant. We therefore opted for the more parameter-
1087 efficient default setting.

1088 **Training Strategies.** We compared several training strategies, including a non-adversarial base-
1089 line. We also tested *regular distillation*, where the model learns to mimic the output of a 50-step
1090 DDIM inversion on real images, as well as a step-wise adversarial tuning similar to diffusion pre-
1091 training (*regular fine-tuning*) (Ho et al., 2020). For *regular distillation*, the performance ceiling
1092 effect is obvious if we were to directly learn from DDIM’s inversion of a distorted image. There-
1093 fore, our actual learning objective is defined as the L2 distance between the noise produced by the
1094 student model on an augmented image and the noise produced by the baseline DDIM on the corre-
1095 sponding clean image. For *regular fine-tuning*, the evaluation is conducted using a 50-step inversion.
1096 The results in Table 9 confirm our claims: adversarial training is vital for FARI’s robustness, the *reg-
1097 ular distillation* approach is limited by its teacher’s performance ceiling, and the step-wise objective
1098 of *regular fine-tuning* fails to learn global robustness.

1100 Table 8: The Performance of FARI on different
1101 trained modules.

	TR TPR	GS Bit Acc.
Attention LoRA FT	0.997	0.9841
Extended LoRA FT	0.996	0.9849
Full FT	0.998	0.9860

1100 Table 9: The Performance of FARI on different
1101 training strategies.

	TR TPR	GS Bit Acc.
FARI w/ adversarial training	0.997	0.9841
FARI w/o adversarial training	0.919	0.9788
Regular distillation	0.995	0.9823
Regular fine-tuning	0.963	0.9774

1109 **Choice of t .** Regarding the choice of t value in Equation 9, we conducted an ablation study by
1110 training with different t values. The results are shown in Table 10. The results demonstrate that
1111 $t \approx 0$ serves as a better initialization for optimization than larger t values.

1114 Table 10: Ablation study on different t values in Equation 9.

	$t = 0$	$t = 100$	$t = 500$	$t = 999$
GS Bit Acc.	0.9841	0.9835	0.9809	0.9772
TR TPR	0.997	0.996	0.995	0.988

1120 E.4 GENERAL EVALUATION AND VISUAL RESULTS
1121

1122 To demonstrate FARI’s broader compatibility with inversion-based watermarking methods beyond
1123 Gaussian Shading (Yang et al., 2024) and Tree-Ring (Wen et al., 2023), we directly measure and
1124 report the MSE of FARI’s reconstructed latent noise in Table 11, with visual results shown in Fig-
1125 ure 11.

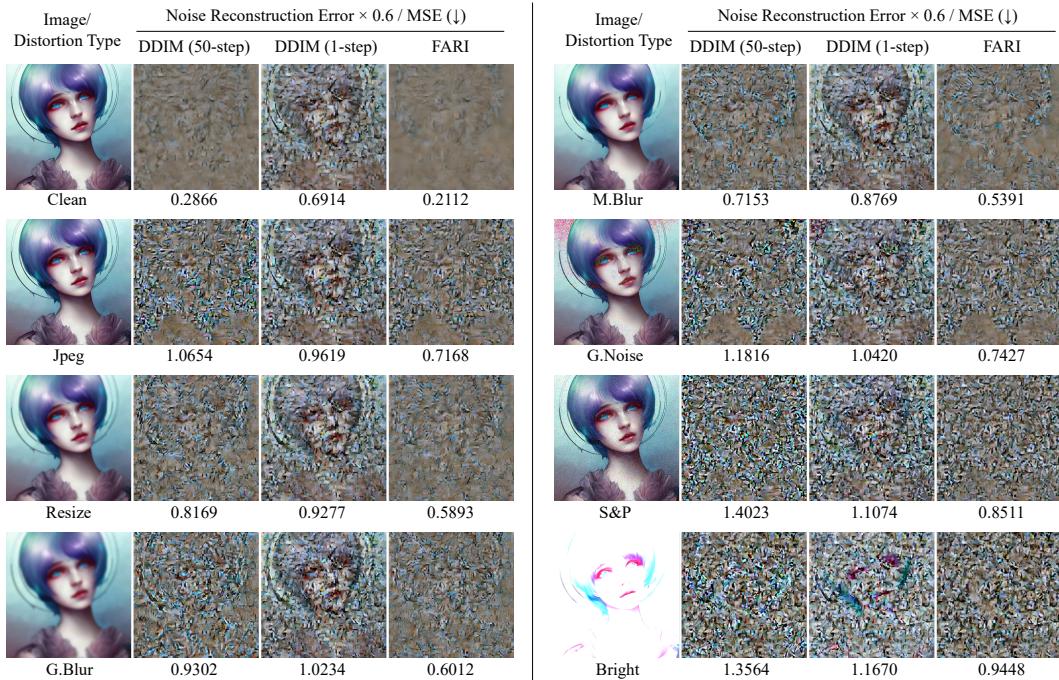
1127 F LLM USAGE STATEMENT
1128

1129 A large language model (LLM) was used as an assistive tool for the writing and editing of this
1130 paper. Its primary function was to polish and rephrase sentences for improved clarity, readability,
1131 and academic style. The authors have reviewed and take full responsibility for all content presented.

1134
1135
1136
1137
1138
1139

Table 11: The MSE of inversion methods under various image distortions.

DM	Methods	NFE	Clean	Adv.	Jpeg	R.Crop	R.Drop	Resize	G.Blur	M.Blur	G.Noise	S&P	Bright
SD v1.5	DDIM	50	0.2003	0.9285	0.8740	0.8734	0.8834	0.7671	0.9706	0.6920	1.1860	1.2963	0.8139
		1	0.7108	0.9016	0.8981	0.9646	0.9621	0.8294	0.8987	0.8157	0.9067	0.9316	0.9074
	EDICT	50	0.2646	1.0477	1.0400	1.0407	1.0547	0.8569	1.0756	0.7654	1.2852	1.3735	0.9370
	BELM	50	0.6593	1.1723	1.0230	1.4621	1.4384	0.9181	1.1744	0.8789	1.2193	1.2065	1.2305
	AMED [†]	2	0.3577	0.8375	0.8022	0.8756	0.8800	0.7116	0.8733	0.6439	0.9562	1.0033	0.7917
	LCM-LoRA	2	0.4488	0.8471	0.7773	0.8264	0.8331	0.7296	0.8614	0.7126	1.0413	1.0904	0.7517
	DMD2	1	0.7934	0.9201	0.9478	0.9472	0.9392	0.8673	0.8999	0.8578	0.9367	0.9587	0.9266
SD v2.1	FARI(Ours)	1	0.2033	0.6699	0.6138	0.7385	0.7396	0.5650	0.7037	0.5486	0.7146	0.7920	0.6135
	DDIM	50	0.2051	0.9657	0.9070	0.9395	0.9614	0.8012	1.0167	0.7080	1.1861	1.2946	0.8769
		1	0.6314	0.9573	0.8634	0.9934	1.0049	0.8861	1.0091	0.8471	1.0543	1.0456	0.9116
	EDICT	50	0.2659	1.0458	1.0419	1.0383	1.0476	0.8560	1.0745	0.7646	1.2700	1.3691	0.9503
	BELM	50	0.6706	1.1705	1.0213	1.4556	1.4313	0.9263	1.1767	0.8872	1.2099	1.1954	1.2190
	AMED [†]	2	0.3538	0.8357	0.8025	0.8729	0.8791	0.7095	0.8726	0.6424	0.9522	0.9990	0.7915
	ExactDPM	> 150	0.2495	1.1306	1.0733	1.2565	1.2645	0.9105	1.1577	0.8165	1.2495	1.4110	1.0359
FARI(Ours)	1	0.2214	0.6862	0.6335	0.7461	0.7530	0.5905	0.7216	0.5638	0.7294	0.7962	0.6418	

1151
1152
1153
1154
1155
1156
1157
1158
1159Figure 11: **FARI** reduces reconstruction error, especially under distortion. Naive DDIM single-step inversion produces visible artifacts. Error scaled by 0.6 to avoid oversaturation.1182
1183
1184
1185
1186
1187