WaDi: Weight Direction-aware Distillation for One-step Image Synthesis

Lei Wang [1] Yang Cheng [1] Senmao Li [1] Ge Wu [1] Yaxing Wang [1,3](✉) Jian Yang [1,2](✉)

[1]PCA Lab, VCIP, College of Computer Science, Nankai University [2]PCA Lab, School of Intelligence Science and Technology, Nanjing University [3]Shenzhen Futian, NKIARI

Motivation

Why do standard fine-tuning adapters struggle for one-step distillation?

1. Teacher and student weights are strongly coupled in norm and direction, but the actual useful change is highly anisotropic.
2. Empirical analysis on DMD2 and PixArt-alpha DMD shows directional changes are much larger than norm changes across layers.
3. Replacing the student direction with the teacher direction severely degrades generation, while replacing only the norm has little effect.

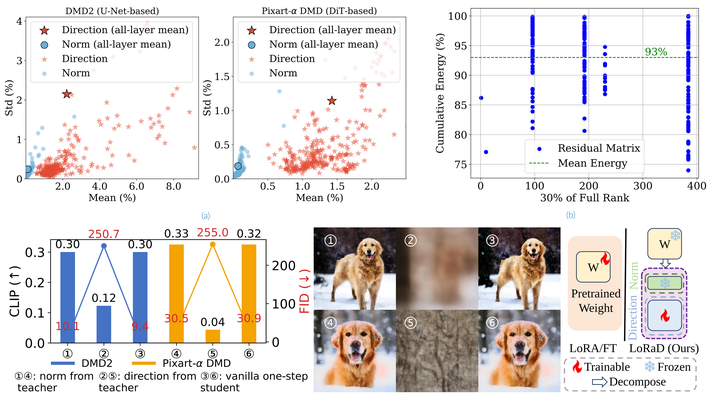

Figure 1. Motivational analysis of our method.

Highlights

- A weight-space analysis showing that direction reconstruction is the key mechanism behind one-step distillation.
- LoRaD: a parameter-efficient adapter that rotates pre-trained weight directions with learnable low-rank angles.
- WaDi: a VSD-based one-step distillation framework that uses LoRaD in both student and fake model branches.
- Strong generalization to controllable generation, relation inversion, high-resolution synthesis, and customization.

Why LoRaD Works

Key idea: keep the norm, rotate the direction.

Instead of additive updates, LoRaD applies orthogonal 2×2 rotations over odd-even weight pairs, then factorizes the rotation angles as $A \cdot B$ to reduce trainable parameters.

Result: more expressive direction updates, preserved norms, and a better quality- efficiency trade-off than LoRA / DoRA / FT.

Takeaways

- Direction-aware updates provide a more faithful target for one-step diffusion distillation than generic low-rank additive tuning.
- The method is backbone-agnostic and has been validated on SD 1.5, SD 2.1, and PixArt-α.
- WaDi achieves SOTA or near-SOTA image quality and text-image alignment while remaining training-efficient.

Wechat

WaDi Links

Code
https://github.com/gudaochangsheng/WaDi

Project
https://gudaochangsheng.github.io/WaDi-Page/

Arxiv
https://arxiv.org/abs/2603.08258

Method Overview

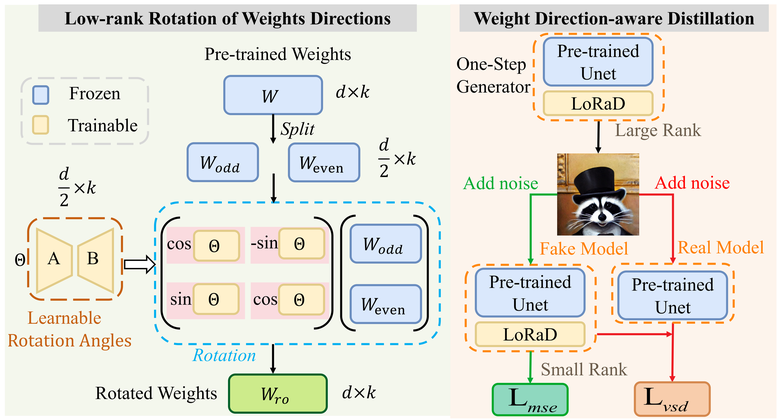

Figure 2. (Left) Detailed architecture of the Low-rank Rotation of weight Direction (LoRaD) module. (Right) Overview of the Weight Direction-aware Distillation (WaDi) framework.

Visualization

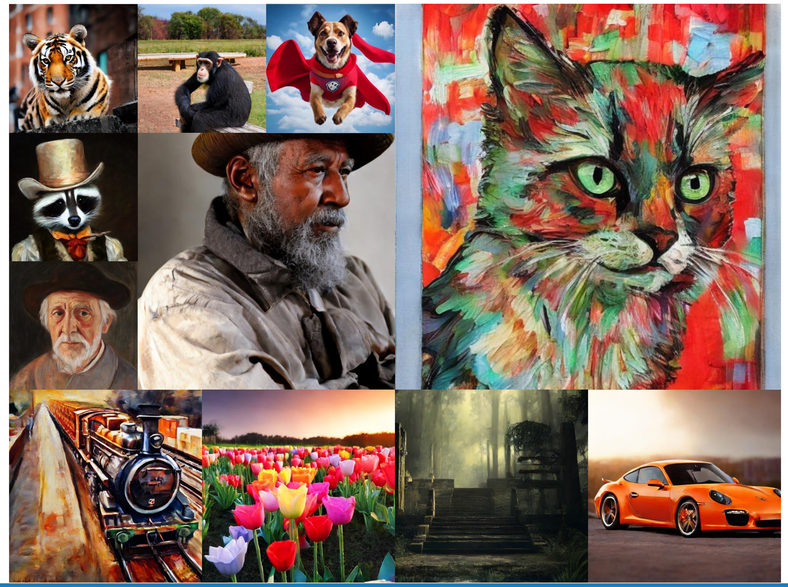

Qualitative Comparison

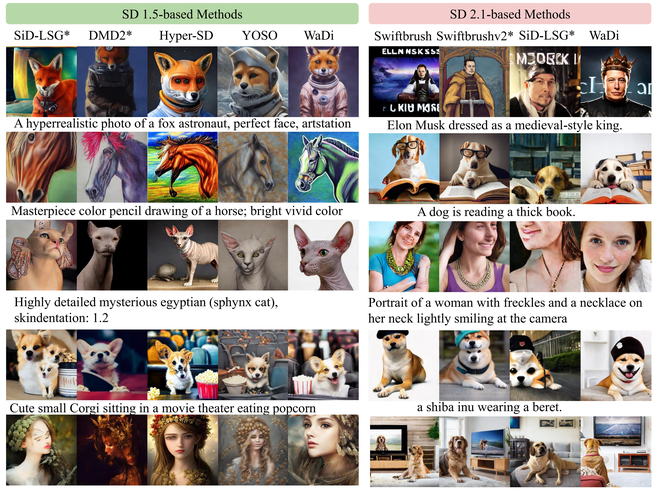

Figure 4. Qualitative comparison with other methods, where * indicates our reproduced results.

Downstream Tasks

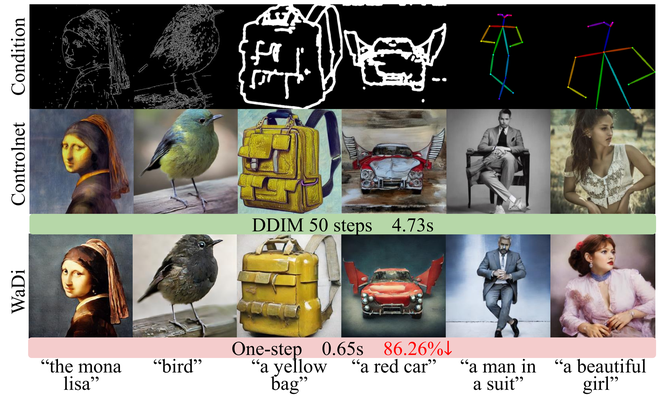

Figure 5. Quality results by Controlnet with or without WaDi.

Acknowledgements

This work was supported by the National Science Fund of China under Grant Nos, 62361166670 and U24A20330, the "Science and Technology Yongjiang 2035" key technology breakthrough plan project (2024Z120), the Shenzhen Science and Technology Program (JCYJ20240813114237048), the Chinese government-guided local science and technology development fund projects (scientific and technological achievement transfer and transformation projects) (254Z0102G), and the Supercomputing Center of Nankai University (NKSC).