



Flux Training Concept POC
RunDiffusion.com

contact@rundiffusion.com

<https://huggingface.co/RunDiffusion/Wonderman-Flux-POC>

Darin Holbrook
Chief Technology Officer
RunDiffusion.com



Train the Data

All tasks were performed on a local workstation equipped with an RTX 4090, i7 processor, and 64GB RAM. Note that 32GB RAM will not suffice, as you may encounter out-of-memory (OOM) errors when caching latents. We did use [RunDiffusion.com](https://huggingface.co/RunDiffusion.com) for testing the LoRAs created, enabling us to launch five servers with five checkpoints to determine the best one that converged

We're not going to dive into the rank and learning rate and stuff because this really depends on your goals and what you're trying to accomplish. But the rules below are good ones to follow.

- We used Ostris's ai-toolkit available here: <https://github.com/ostris/ai-toolkit/tree/main>
- Default config with LR: 4e-4 at Rank 16
- 2200 - 2600 steps saw good convergence. Even some checkpoints into the 4k step range turned out pretty good.

If targeting finer details, you may want to adjust the rank up to 32 and lower the learning rate. You will also need to run more steps if you do this.

Training a style: Using simple captions with clear examples to maintain a coherent style is crucial. Although caption-less LoRAs can sometimes work for styles, this was not within the scope of our goals, so we cannot provide specific insights.

Training a concept: You can choose either descriptive captions to avoid interfering with existing tokens or general captions that might interfere, depending on your intention. This choice should be intentional.

Captioning has never been more critical. Flux "gives you what you ask for" - and that's a good thing. You can train a LoRA on a single cartoon concept and still generate photo realistic people. You can even caption a cartoon in the foreground and a realistic scene in the background! This capability is BY DESIGN - so do not resist it - embrace it! (Spoiler alert to the right)

You'll see in the next page of examples where the captioning really helps or hurts you. Depending on your goals again you will need to choose the path that fits what you're trying to accomplish.

Total time for the LoRA was about 2 to 2.5 hours. \$1 to \$2 on RunPod, Vast, or local electricity will be even cheaper.

Now for the results! (This next file is big to preserve the quality)

