

AyyMD Stable Diffuse v1.4 for Wangblows 10 (by anon)

When you want it to run on GNU/Linux, follow the ROCm guide. (<https://reentry.org/sdamd>)

Prerequisites

You have installed Python 3.x (min. 3.7)

<https://www.python.org/downloads/windows/>

A huggingface.co account (to generate and download the stable-diffusion v1.4 model)

<https://huggingface.co/>

Accepted terms of use for stable diffusion with your account (download will fail without it)

<https://huggingface.co/CompVis/stable-diffusion-v1-4>

Installation

- Make a folder where you will download all stuff into
- Open Command Line (as Administrator, or it will fail later on because of symlinks)
- Clone special hack version by harishanand95

```
git clone https://github.com/harishanand95/diffusers.git
cd diffusers && git checkout dml && pip install -e .
pip install transformers ftfy scipy
```

- Check for the newest nightly of onnx and download it (at time of writing it was "1.13.0.dev20220901005")

https://aiinfra.visualstudio.com/PublicPackages/_artifacts/feed/ORT-Nightly/PyPI/ort-nightly-directml/overview/

If you are on Python 3.7 download the file that ends with `** -cp37-cp37m-win_amd64.whl`

If you are on Python 3.8 download the file that ends with `** -cp38-cp38-win_amd64.whl`

If you are on Python 3.9 download the file that ends with `** -cp39-cp39-win_amd64.whl`

If you are on Python 3.10 download the file that ends with `** -cp310-cp310-win_amd64.whl`

- I still use Python 3.7 so for me it is:

```
ort_nightly_directml-1.13.0.dev20220901005-cp37-cp37m-win_amd64.whl
```

- Download it and install it

```
pip install ort_nightly_directml-1.13.0.dev20220901005-cp37-cp37m-win_amd64.whl
```

An error message like this "ERROR: ort_nightly_directml-1.13.0.dev20220830001-cp38-cp38-win_amd64.whl is not a supported wheel on this platform." means that there is mismatch in python version and the downloaded package supported python version.

Generating a onnx-based stable-diffusion v1.4 model

- Go to folder "diffusers\examples\inference"
- Login via cli and generate a token (huggingface-cli was installed by transformers package)

```
huggingface-cli login
```

<https://huggingface.co/settings/tokens>

- Start downloading and generating of stable-diffusion v1.4 model

```
python ./save_onnx.py
```

Generate a image

- Open dml_onnx.py with notepad and edit what you want to generate at the bottom of the file in "prompt" variable
- Default for testing is: `prompt = "a photo of an astronaut riding a horse on mars"`

- Save and execute with "python ./dml_onnx.py"
- After a bit of time you will get a png in the "diffusers\examples\inference" folder.

I have about ~3.30/it on my good ol AMD RX 480 and a image takes about "2:30 min". You should have a higher number.