YOU NEED 20G OF VRAM

YOU NO LONGER NEED 30.5 GB OF VRAM

TLDR:

- Works on 20GB VRAM now (3090 bros), without a need of massive RAM
- We need your help (any kind of help, but mostly devs and researchers)
- Code got forked and privated and the guy is planning to turn it SAAS
- Distributed training might be possible
- we need developers, developers, developers, developers, developers, developers, developers, developers, developers
- Haru spent 3K so far

Info on: https://discord.gg/VvuQX3a5h6 & https://discord.gg/AFHNGG8APj

Important Notice:

Its been a while since I have updated this guide. I will no longer update this guide, the docs will move to docs.conlaboro.xyz soon.
Yes, yes it is now possible to run the whole process on a single RTX TItan, RTX 3090, RTX 3090ti, and any card with at least 20GB of VRAM.
This is done with finally porting bf16, and reducing batch and worker size to 1. The documentation will be updated.
We are also looking into a way to do **distributed training** via **merging**, but more info on that is available on the server.

HOWEVER....

Someone has forked our code, made it private, implemented his own modifications, and is going to make it software as a service (SAAS), the model is going to be closed source from what I have heard. That's right, the repo doesn't even has a proper license yet. I have been talking with haru to try to move it to AGPL, but right now the top priority is **getting researchers and developers involved**. We have been talking with unstable diffusion (yes I hate them too) to try to colaborate, as we have seen that **MANY** trainers have already been "manipulated" by this guy to try and join him, although it is unknown wether their models will be public or not. Just as a fact, Haru is doing all of this on his own and its costing him 3,000$ so far.

Our progress so far...

Heres a list of what has changed since the last guide revision:

- Updated Dataloader to use Webdataset rather than directory based loading (30% boost on 8x config)
- bf16 working reducing necessary vram to 20gb (although training results have not been great, there should also be a 20% boost)
- Currently developing a way to train on any resolution (not necessarily 512px)
- Config file that reduces necessary vram to 23gb (batch and worker size = 1)

This was written on 26/09/2022

THE FOLLOWING IS OUTDATED!

~~**In the future, it might be possible to reduce it to around 20gb, for updates on fp16 join the touhou ai discord**~~

Training News:

11/09/2022:
~~TURNS OUT BRONIES HAD THIS TECHNOLOGY ALL THIS TIME! They didnt even shared the code~~
~~Supposedly its even possible on 12gb gpus... More information later today.~~
somewhat fake, the only brony we know didnt used 12gb gpus rather 8 A6000
could be another one though

12/09/2022:
~~Its a little bit hard to explain, but basically we have been trying multiple versions of attention.py and model.py located on the ldm folder.~~
~~Some of these modifications you might already have, such as in the REDACTED repo.~~

~~There are multiple versions, the latest and most updated version currently is:
https://github.com/lstein/stable-diffusion/tree/development/ldm/modules
but rather than decreasing VRAM usage it increases it to 35.5gb.
Waifu Diffusion repo had a pr for days with a similar solution:
https://github.com/harubaru/waifu-diffusion/pull/9

~~If you want to keep up to date and/or get more information make sure to subscribe/check out the following sources:

~~https://github.com/lstein/stable-diffusion/issues/431~~
~~https://discord.gg/VvuQX3a5h6 - research finetune channel~~
~~https://discord.gg/jDHZYHburh - everything really~~
~~https://github.com/neonsecret/stable-diffusion~~

~~On another note, I will not say names, but there are two similar users, do not confuse them. If you know you know.~~

15/09/2022:
Training on lower mem has been confirmed, not as efficient though.
Expect a colab notebook soon.
Also I will be releasing more self-trained models, datasets, and scripts on the discord.

informal training guide by dep#2171
look I will write this with no organization whats so ever beacuse thats what takes time
If you have any problems, you can always join "stable-diffusion-training-lounge" discord: https://discord.gg/xVsyrmhQWS
I am also writing this because currently im running a file_ext extractor thats going to take a while so yeah I dont want to waste my time
this is pseudo windows and linux.

requirements for personal instance, can be win or linux:
anything really, just a big drive with at least 30gb free
requirements for linux instance:
at least 30gb of disk storage, 30.5gb of vram, and 11gb of ram
you need both not just one

if you can't understand something in this guide again, go to the discord cuz im not going to edit this a million times like the prior guide I did
and a ton of people accused me for spreading misinformation
Im also writing this off my head beacuse paperspace just terminated me temporaly

my setup is as follows:
windows with cygterminal installed
linux with a6000 40gb of vram AND ram, 50gb of disk storage

lets start with the dataset
go down to see the dataset format

download cygwin if on windows, if on linux just install rsync
LETS SAY I want to download batches 0000-0009, do this:
rsync --progress rsync://176.9.41.242:873/danbooru2021/512px/000* ./512px/
this will download the folders 0000 to 0009 onn 512px
now the metadata:
rsync --progress rsync://176.9.41.242:873/danbooru2021/metadata/posts000000000000.json ./metadata/
same thing but on metadata folder

now lets extract the images along with its tags.
download the extractor
wget https://github.com/chavinlo/waifu-diffusion/blob/patch-3/danbooru_data/local/extractfromjson_danboo21.py

flags:
--jsonpath path to json
--extractpath path to where to extract the images and text files
--rating optional, select ratings from e (explicit), g (general), q (questionable) or s (safe), only put the letter in small caps
--convtohuman ignore its broken currently
--imagespath optional, put the path to images, if none it will default to 512px

so it should be:
py extractfromjson_danboo21.py --jsonpath metadata/posts000000000000.json --extractpath extracted --imagespath 512px --rating s
why rating s? beacuse everything in the 512px folder is only rated safe anyways. explicit and rest are on original folder but for more info on

that join the discord this will just explain the basics

it will output everything to the folder "extracted"
there will be txt files and jpg files both with same names
now you compress them and send it to your training (linux) instance

once there, run the following:

```
git clone https://github.com/harubaru/waifu-diffusion
cd waifu-diffusion
pip install -e .
pip install omegaconf einops pytorch-lightning==1.6.5 test-tube transformers kornia
pip install -e git+https://github.com/CompVis/taming-transformers.git@master#egg=taming-transformers
pip install -e git+https://github.com/openai/CLIP.git@main#egg=clip
pip install setuptools==59.5.0
pip install pillow==9.0.1
pip install torchmetrics==0.6.0
```

IF ANY OF THE DEPENDENCIES FAIL TO INSTALL, just install them again but without the specific versions and individually
IF YOU GET "module not found taming/clip" ERROR, just copy as follows:
download the zip from https://github.com/CompVis/taming-transformers (source code)
unzip, and copy the "taming" folder to the root directory of the waifu diffusion repo
same for CLIP, download zip from https://github.com/openai/CLIP.git , unzip, copy the folder "clip" and to root directory of repo

download your dataset.zip

```
unzip dataset.zip
cd dataset
mkdir txt
mkdir img
mv .txt txt
mv img
cd ..
mv dataset danbooru-aesthetic
```

Unless haru (waifu mantainer) confirms that the loading directory is no longer hardcoded, DO NOT CHANGE THE FOLDER NAME AT ALL.
this caused me so many problems.

nano my-train.sh
add the following:

```
python3 main.py -t -n "waifu-diffusion" --resume_from_checkpoint model.ckpt --base ./configs/stable-diffusion/v1-finetune-4gpu.yaml --no-test --seed 25 --scale_lr False --gpus 0, --data_root "./danbooru-aesthetic"
```

you could try with the 1gpu yaml config, but the 4gpu works exactly the same.
Remember to add a coma next to gpu count
--resume_from_checkpoint does passes the model now, you can try with a nonexistent file and you will see it fails loading it

ctrl-x, save, yes same name

```
chmod +x my-train.sh
./my-train.sh
```

to resume training, just change the --resume_from_checkpoint path to the latest model, remember that they will be stored on the logs folder.
Everytime you run the script, a new folder will be created, most of which will weight about 3gb each, so be careful with storage.

TROUBLESHOOTING HERE:
there might be a million errors more after this guide, so again, JOIN THE SERVER IF YOU HAVE QUESTIONS

ERROR 1: KeyError: 'Trying to restore optimizer state but checkpoint contains only the model. This is probably due to ModelCheckpoint.save_weights_only being set to True.'
FIX 1:
patch pytorch lighting as follows, credits to haru and astralite for the fix:
your terminal will most likely print a path that points to a file named "checkpoint_connector.py" during the logtrace, copy that path

nano PATH TO checkpoint_connector.py

search for "state_dict" and remove everything after it and replace it for "and False:"
should be something like this, remember that it varies through versions:
https://media.discordapp.net/attachments/979237546028597288/1016238769562734672/unknown.png
https://i.imgur.com/32YLzl1.png

search for the call of "self.restore_optimizers_and_schedulers()" or the definition of "restore_training_state"
comment out the call of "self.restore_optimizers_and_schedulers()"
should be like this:
https://media.discordapp.net/attachments/979237546028597288/1016239598202015825/unknown.png
https://i.imgur.com/iAGud4J.png

ERROR 2: `FileNotFoundError: [Errno 2] No such file or directory: ''`
full traceback:

```
Traceback (most recent call last):
  File "/notebooks/waifu-diffusion/main.py", line 719, in <module>
    trainer.fit(model, data)
  File "/usr/local/lib/python3.9/dist-packages/pytorch_lightning/trainer/trainer.py", line 770, in fit
    self._call_and_handle_interrupt(
  File "/usr/local/lib/python3.9/dist-packages/pytorch_lightning/trainer/trainer.py", line 721, in
_call_and_handle_interrupt
    return self.strategy.launcher.launch(trainer_fn, args, trainer=self, **kwargs)
  File "/usr/local/lib/python3.9/dist-packages/pytorch_lightning/strategies/launchers/subprocess_script.py", line 93, in
launch
    return function(args, kwargs)
  File "/usr/local/lib/python3.9/dist-packages/pytorch_lightning/trainer/trainer.py", line 811, in _fit_impl
    results = self._run(model, ckpt_path=self.ckpt_path)
  File "/usr/local/lib/python3.9/dist-packages/pytorch_lightning/trainer/trainer.py", line 1236, in _run
    results = self._run_stage()
  File "/usr/local/lib/python3.9/dist-packages/pytorch_lightning/trainer/trainer.py", line 1323, in _run_stage
    return self._run_train()
  File "/usr/local/lib/python3.9/dist-packages/pytorch_lightning/trainer/trainer.py", line 1342, in _run_train
    self._pre_training_routine()
  File "/usr/local/lib/python3.9/dist-packages/pytorch_lightning/trainer/trainer.py", line 1335, in _pre_training_routine
    self._call_callback_hooks("on_pretrain_routine_start")
  File "/usr/local/lib/python3.9/dist-packages/pytorch_lightning/trainer/trainer.py", line 1636, in _call_callback_hooks
    fn(self, self.lightning_module, *args, kwargs)
  File "/notebooks/waifu-diffusion/main.py", line 260, in on_pretrain_routine_start
    os.makedirs(self.logdir, exist_ok=True)
  File "/usr/lib/python3.9/os.py", line 225, in makedirs
    mkdir(name, mode)
FileNotFoundError: [Errno 2] No such file or directory: ''
```

FIX 2: put the model on waifu-diffusion/logs/original/checkpoints/last.ckpt
waifu-diffusion being the root of the repo
CHANGE THE my-train.sh FILE TOO:

```
python3 main.py -t -n "waifu-diffusion" --resume_from_checkpoint logs/original/checkpoints/last.ckpt --base
./configs/stable-diffusion/v1-finetune-4gpu.yaml --no-test --seed 25 --scale_lr False --gpus 0, --data_root "./danbooru-
aesthetic"
```

ERROR 3: cannot find module "clip" or "taming"
FIX 3:
IF ANY OF THE DEPENDENCIES FAIL TO INSTALL, just install them again but without the specific versions and individually
IF YOU GET "module not found taming/clip" ERROR, just copy as follows:
download the zip from https://github.com/CompVis/taming-transformers (source code)
unzip, and copy the "taming" folder to the root directory of the waifu diffusion repo
same for CLIP, download zip from https://github.com/openai/CLIP.git , unzip, copy the folder "clip" and to root directory of repo

ERROR 4: `AttributeError: 'int' object has no attribute 'strip'`
traceback:

```
Traceback (most recent call last):
  File "C:\Users\T\waifu-diffusion\main.py", line 676, in <module>
    ngpu = len(lightning_config.trainer.gpus.strip(",").split(','))
AttributeError: 'int' object has no attribute 'strip'
```

FIX 4:
add a coma to gpu count
if one gpu:
--gpus 0,

ERROR 5: script just closes out of nowhere without output (well it prints one blank line and then it exits)
this one made me almost kill myself
FIX 5:
you are probably running train.sh which is a little bit broken, its hard to explain but heres my interaction with haru

now a similar error that I think I encountered before happens, it just exits
give me a sec I think its on the touhou disc
nope apparently not
some dead link here
haru — hoy a la 1:08
seems like it's working!
or not
dep — hoy a la 1:08
it just exited
haru — hoy a la 1:08
hm. i noticed before sometimes it executed a signal
dep — hoy a la 1:09
so something else is terminating it?
haru — hoy a la 1:09
not too sure honestly, what are the specs of the machine running it?
dep — hoy a la 1:10
A6000 45gb of ram
haru — hoy a la 1:13
you should see if it's maxing out the ram
dep — hoy a la 1:14
nop
max 10gb
haru — hoy a la 1:14
thats pretty strange
dep — hoy a la 1:15
is your instance working?
like, your version
haru — hoy a la 1:16
yeah
dep — hoy a la 1:16
do you think astraline would give me a copy of his version? like a zip file not a docker img
oh
haru — hoy a la 1:16
i just used whatever was on main
dep — hoy a la 1:16
could u send me yours
a
wait I will try with just the normal sh
do i add the -t flag?
dep — hoy a la 1:49
im going to shoot myself
nvm
haru — hoy a la 1:49
huh
haru — hoy a la 1:50
yeah
dep — hoy a la 1:50
it worked for a sec when the training folder renamed to danbooru-aesthetic
but it still exits
its working??
like im not even sure at this point
beacuse I remember doing this exact same thing
but the problem is that I don't think its training from the checkpoint, because if the argument is set to a non existent path it keeps working asusual
nvm it does
thank you so much
took a while but it works
haru — hoy a la 1:57
hm
must be an issue with the dataloader
dep — hoy a las 2:01
First, reinstalled all the dependencies from your message, then added --data_root "./danbooru-aesthetic" and renamed my data folder to that name (it was a step that I forgot) but even so it exited, then I tried with -n "waifu-diffusion" but got a conflict with --resume parameter so changed it to --resume_from_checkpoint with the same path, and then boom it worked

but what I am worried about if its actually using the checkpoint/weights because I remember you said not too long ago that it didnt even passed the --resume_from_checkpoint flag, although I just tried with a non-existent path and it does seem to know that it doesnt exists, but maybe its just that and it doesnt uses it
haru — hoy a las 2:09
hmm, it'll say if it's resuming from a checkpoint and will give a list of trainable parameters
haru — hoy a las 2:11
yep
looks like it resumed from it

alright that should be it
my parsing got completed, over 3million images which im defo not going to use
see you all next time and updates will be either on my discord (https://discord.gg/xVsyrmhQWS), touhou ai discord or the official wd repo
remember join the discord for support or any questions