

April

An automatic update to cuda has broken my virtual environments. Creating new virtual environments doesn't seem to help. I get an error saying that all available cuda devices are busy when there are actually no devices in use. (e.g. nvidia-smi shows 0% gpu utilization, no vram occupied and no processes claiming the gpu).

In addition, snapper has non-sane defaults on ubuntu and thus my rollback was broken. After rolling back, the .snapshots folder was lost so I was unable to rollback again. I found a solution which involves writing something to your fstab for .snapshots. I copied the actual config from my opensuse tumbleweed install and it worked. However, no matter how far back I go with my snapshots, I still get the cuda device busy error. :(.

I booted another ubuntu os from my portable ssd with cuda installed on it and sure enough it works, I can train the model (and see both the gpus in nvidia-smi).

I am going to try and see if my m40 gpu is compatible with the cuda drivers on opensuse leap (Hopefully this is the case as I've purchased another m40 for cheap to put in my opensuse leap system). If, indeed it does work, I will be switching to opensuse leap for mwanafunzi. I know that snapper has good defaults and works well on tumbleweed and leap so hopefully this rollback issue won't be a problem in the future and if nvidia breaks my setup, I can just rewind. This was the original reason why I was using snapper but apparently I had it misconfigured.

Revision #2

Created Fri, Apr 16, 2021 2:02 PM by [kenneth](#)

Updated Fri, Apr 16, 2021 2:10 PM by [kenneth](#)