

Docker tutorial

Step 1: Understand docker

- Read <https://docs.docker.com/get-started/docker-overview/> and understand (docker) image, container, daemon, client, hub (registry), pull, push.
- Install docker (Linux Mint as example. Information on other OSes are welcome too.)
 - `$ sudo apt install docker.io`
 - Add your username to the "docker" group (or by editing `/etc/group` directly).
 - `$ sudo usermod -aG docker your_username`
 - Logout & login to make the change of group into effect, or simply reboot your Linux.
- Run a ubuntu container (docker client will automatically pull a ubuntu image and create a local container.). If you get a console with command prompt "`root@xxxxxxxxxxxx:/`," it worked.
 - `$ docker run -i -t ubuntu /bin/bash`
- Exit the container by `exit` (notice the prompt changed to "`#`", showing you are the root)
 - `# exit`

Step 2 (if you have an NVidia GPU): Install Nvidia GPU support

- Follow <https://docs.nvidia.com/datacenter/cloud-native/container-toolkit/latest/install-guide.html>

Step 3: Install PyTorch (for development)

- `$ docker pull pytorch/pytorch:latest`
- `$ docker run --rm --gpus all pytorch/pytorch:latest python -c "import torch; print(torch.cuda.is_available())"`
- If you see "True", it worked.

References

- <https://github.com/saikhu/Docker-Guide-for-AI-Model-Development-and-Deployment> (Chapter 6 is obsolete. See the above Step 2 for the latest information)

From:
<https://aw.gsais.kyoto-u.ac.jp/wiki/> - Future Wisdom @ GSAIS (Shishu-Kan) , Kyoto U.

Permanent link:
<https://aw.gsais.kyoto-u.ac.jp/wiki/doku.php?id=public:docker&rev=1740319267>

Last update: 2025/02/23 14:01



