

Using GPU in A219

For the more computationally demanding exercises, you can use GPU on the local machines in A219 with a docker image containing necessary libraries.

Important Things to Keep In Mind

- As with all provided resources they are shared resources, so be mindful of others who might be using them
- Machines have a limited `/var/tmp` space (place where the docker images are stored), have one docker image on one machine, do not needlessly create multiple docker images
- If you are no longer using the image, remove it, for example by calling:
`podman system prune -a`
- As the docker image is stored in a temporary directory, data in `/var/tmp` might be removed during the semester, so be careful

How to Run

On the machine `nymfe<number>`:

```
> git clone https://github.com/MIR-MU/pv211-utils.git
> cd pv211-utils
> git checkout spring2025
> docker build -t pv211:latest .
> podman run -p 8888:8888 --device nvidia.com/gpu=0 -it pv211:latest
```

In the created container:

```
> jupyter-lab --ip="0.0.0.0"
```

You will be provided with the server url, where you can run and modify your code in the jupyter notebook.

Remote Access

You can connect to the machine remotely, with some caveats.

- 1) Check the availability of the machine before using it [here](#).
- 2) You have to connect to the machine through another MU server (for example [aisa](#)).

Further information available [here](#).

To connect from your local machine to an available machine in A219:

```
local: > ssh <your_username>@aisa.fi.muni.cz
```

```
aisa: > ssh <your_username>@nymfe<number>
```