Linux @ FNNDSC

What I Wish I Knew, v1.0.0

Jennings Zhang | 2023-08-16

Outline

1. About this presentation
2. About us
3. How to get help
4. Beginner's Guide to Linux
5. The Command Line Interface
6. Best Practices
7. Beyond Linux: E2 and *ChRIS*

About this Presentation

- Orange text: “fun fact”, “good to know”
  - E.g. UNIX is a historical predecessor to MacOS and Linux
- Blue text: FNNDSC-specific information
  - E.g. you can use the same account to login to any workstation
- $ starts_with_dollarsign ← this is a command

Meet the team

And who to get help from

$ whoami

Rudolph Pienaar

Technical Director

Sandip Samal

Software Developer

Jennings Zhang (me)

Research Assistant

Gideon Pinto

Software Developer

We are the Advanced Computing Group

(a.k.a. the *ChRIS* core developer group)

Newborn_FNNDSCdev-dl@childrens.harvard.edu

Email us for computer help and technical questions!

(or you could email dev@babyMRI.org)

#dev-help on Matrix/Slack

[matrix]

or

https://matrix.to/#/#chris-devhelp:fedora.im
(preferred)

https://fnndsc.slack.com

Who are we *not*?

What we do

- Create FNNDSC Linux accounts (Firstname.LastName)
- Manage Linux computers at the Landmark Center

What we *don't* do

- Create BCH accounts (chXXXXXX)
- Manage E2 cluster and access
- Wifi, VPN, ...

Hospital IT Contact

- FNNDSC Advanced Computing Group (ACG)
  - Newborn_FNNDSCdev-dl@childrens.harvard.edu
    - Landmark Center Linux computers and software
    - *ChRIS*
    - PACS Query/Retrieve
- BCH Research Computing (RC)
  - Research.Computing@childrens.harvard.edu
    - E2 cluster
- BCH Information Services Department (ISD)
  - help.desk@childrens.harvard.edu
    - chXXXXXX account
    - Email (Outlook)
    - WiFi
    - Remote work and VPN
    - Kronos, PeopleSoft, ...

Beginner's Guide to Linux

Why Linux?

- ~~Because everyone else is using it~~
- Free, open-source software
- Because everyone else is using it

- Why not:
  - Apple MacOS: costs money, restrictive policies
  - Microsoft Windows: costs money, bad security, defiant of standards

App Alternatives

- Microsoft Word --> LibreOffice Write
- Microsoft PowerPoint --> LibreOffice Impress
- Microsoft Excel --> LibreOffice Calc

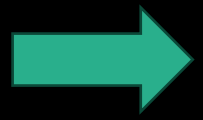

LibreOffice
The Document Foundation

Using Office Online

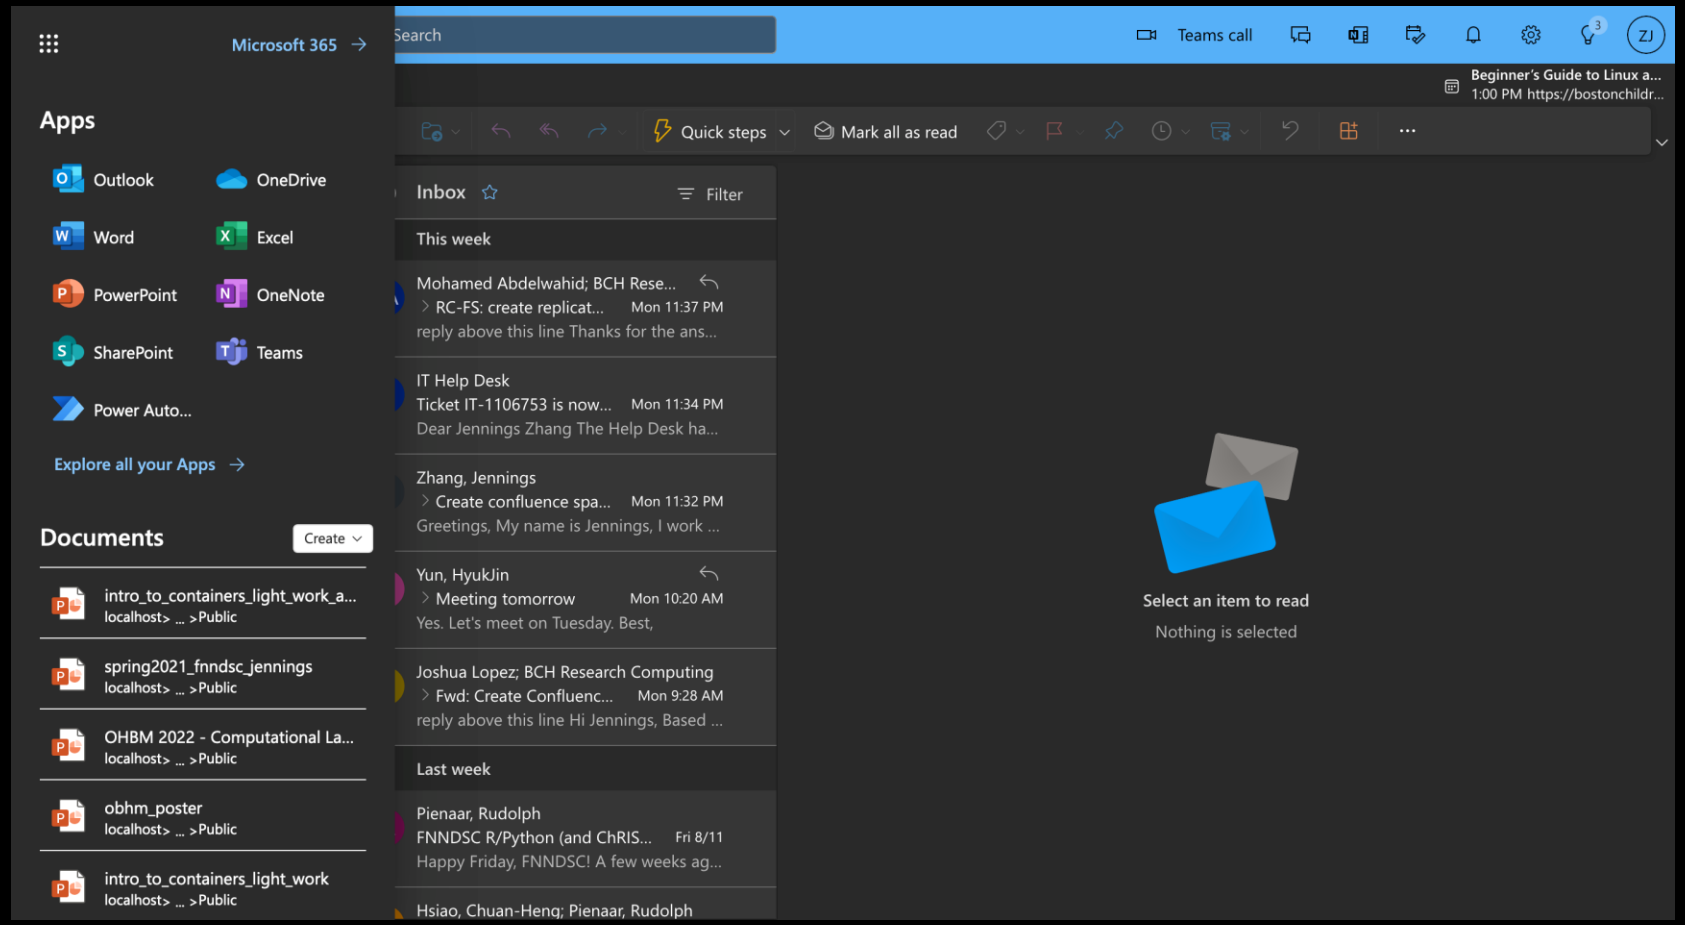

Other Online Software

- Note: can't download from BCH OneDrive
- Using Google docs? Be careful about PHI!
  - BCH Google Drive → http://drive.bchresearch.org/ (you need to request access)
- Use Dropbox for sharing files
  - http://chbwiki.tch.harvard.edu/display/BCL/Dropbox

- Research Computing Wiki: http://rcwiki
- BCH Intranet: http://web2.tch.harvard.edu

How to Install Apps?

- Download from the internet, create .desktop file in ~/.local/share/applications

- $ sudo apt install <package>
  - Ask your PI or Rudolph for the "toor" password
- $ apptainer build image.sif docker://repo/app

- Avoid Snap (it's bad)
- Avoid Flatpak (it's broken)
- AppImage is great!

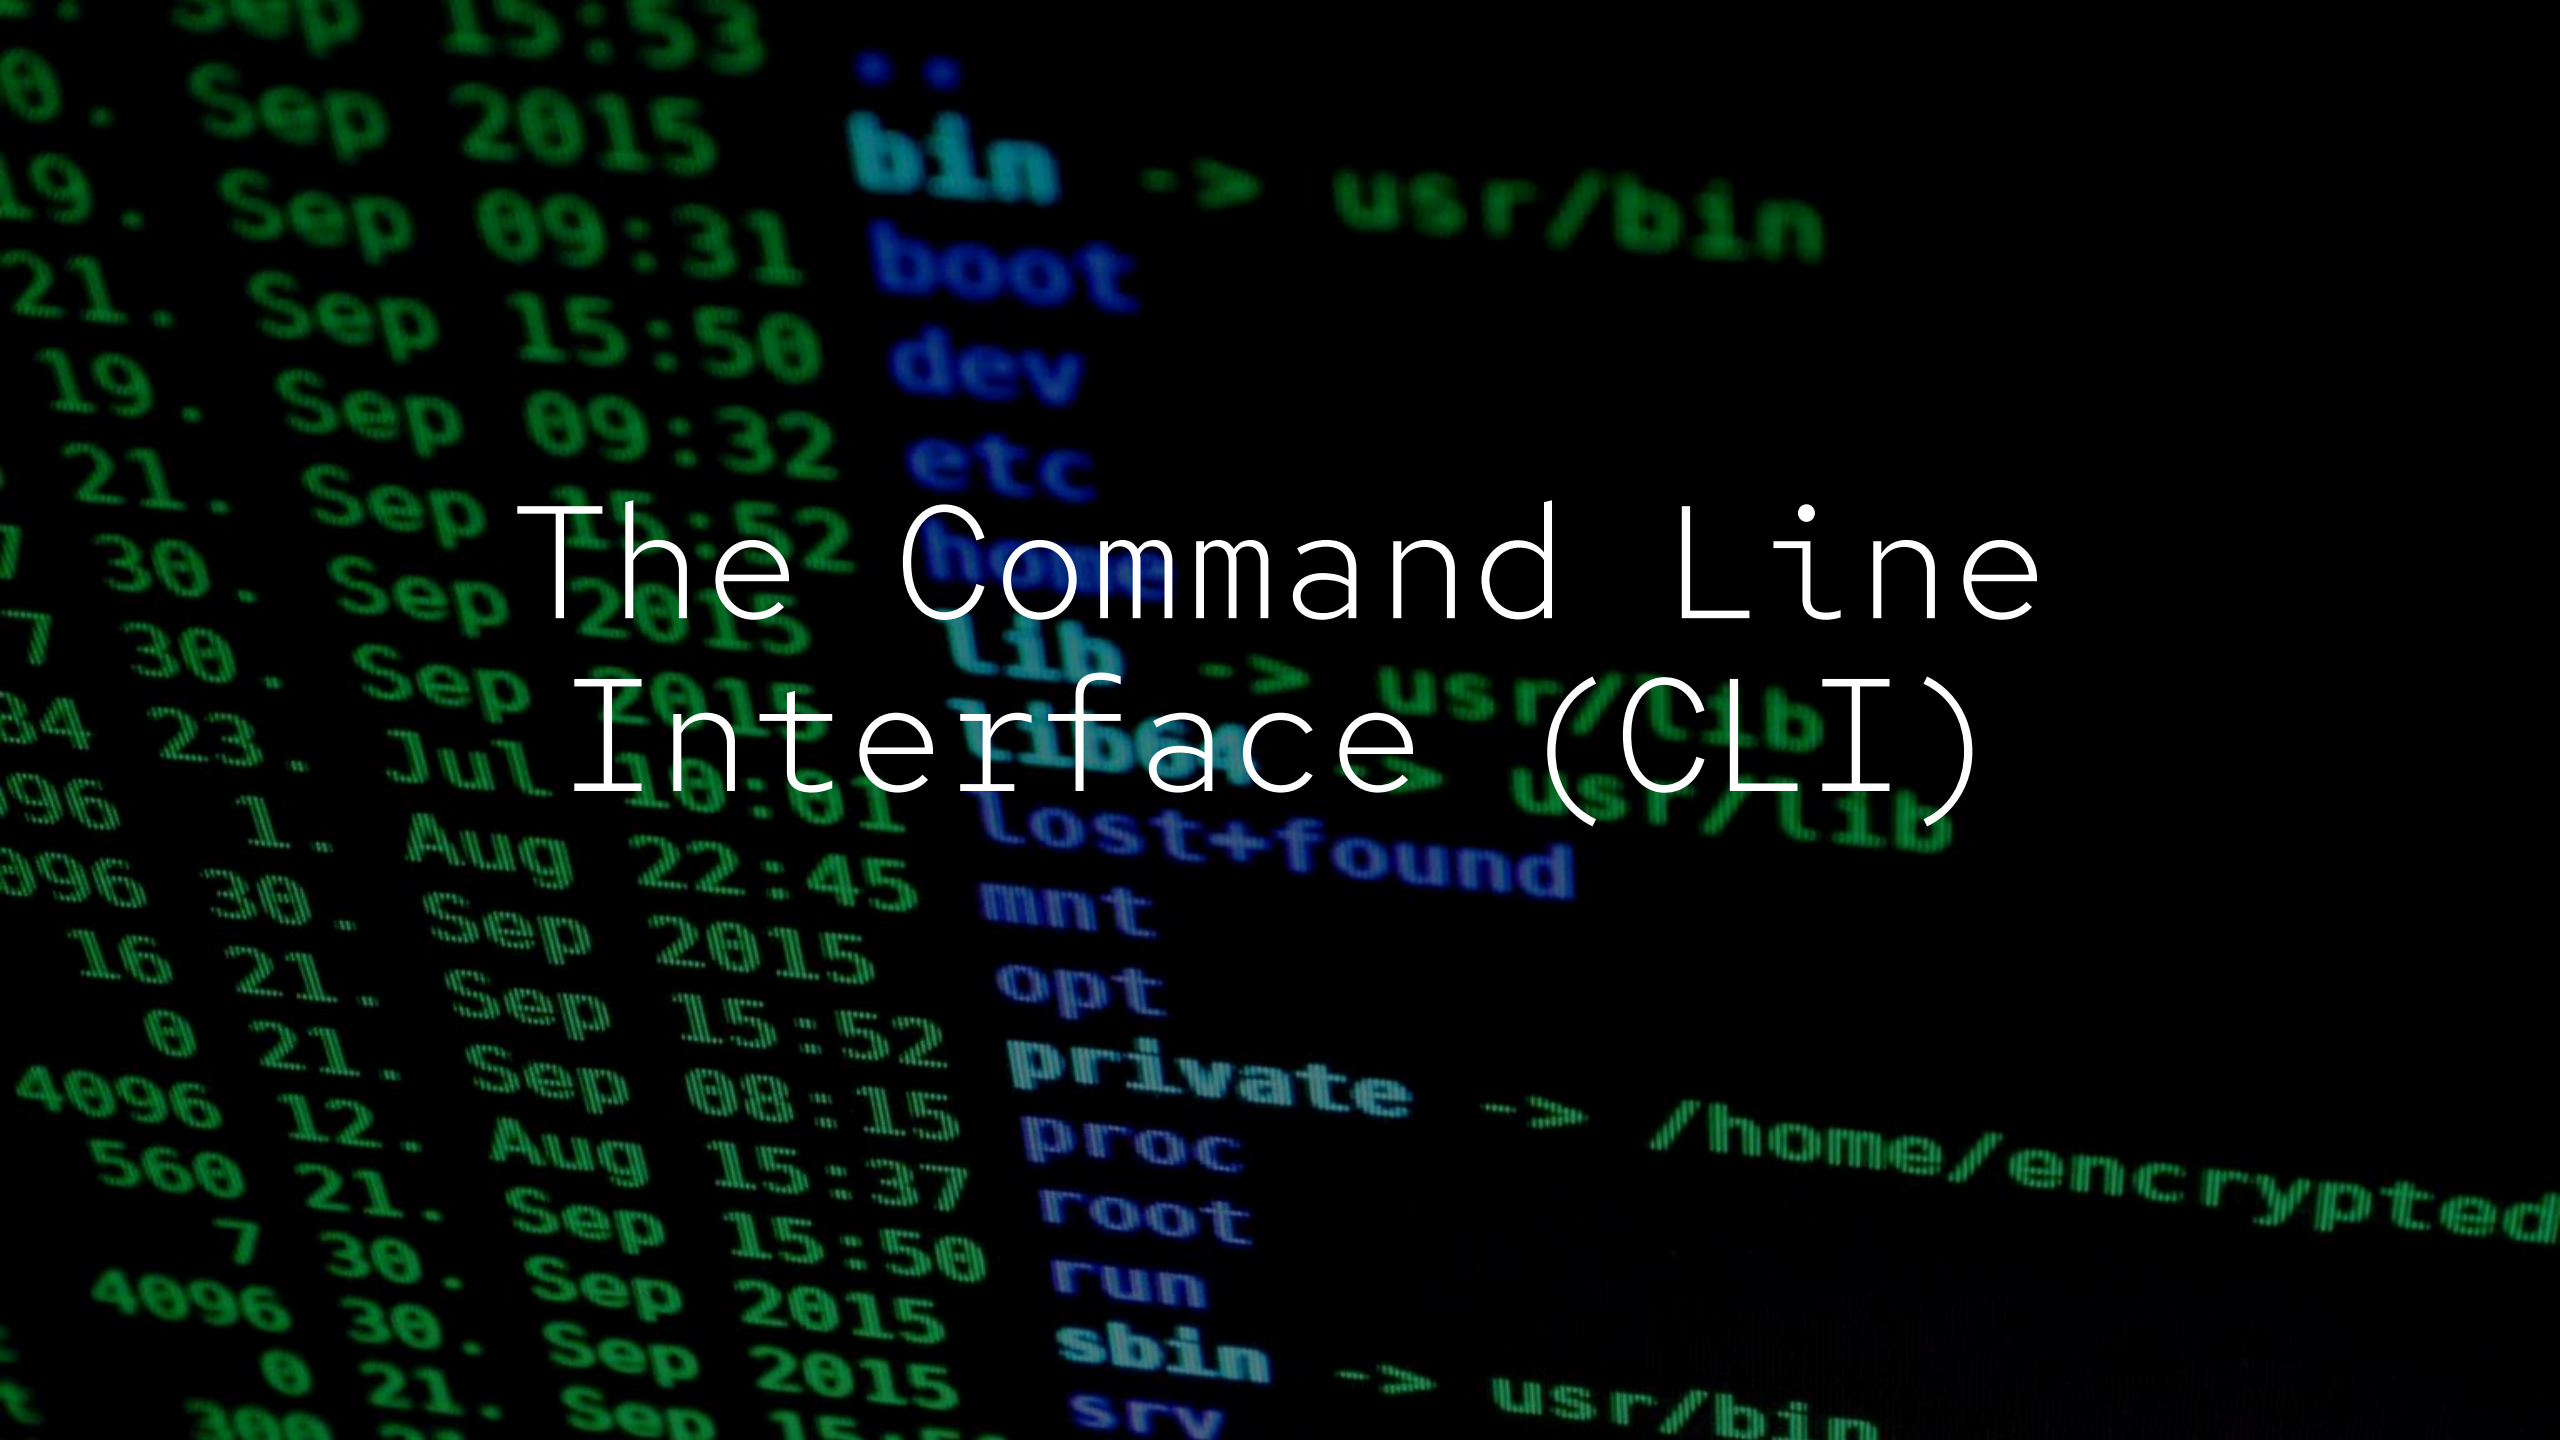

The Command Line Interface (CLI)

Shell Concepts: Commands

```
$ command --parameter value ...
    |          |
    |          └── program options
    |                 may or may not be optional
    |                 may or may not be order dependent
    └─────  program name
               Found in /bin, /usr/bin, or ~/.local/bin
```

How to get help for a command

```
$ whatis command

$ man command

$ command --help

$ command -h

$ command -help
```

Environment Variables and $PATH

- Commands ARE programs (most of them)
- Programs ARE commands
- Environment variables are shared with the programs you run
- $PATH is a variable where programs are found
  - Default PATH=/bin:/usr/bin:/usr/local/bin

Environment Variables and $PATH

```
$ printenv

$ printenv | grep PATH

$ export VARIABLE_NAME=value

$ export PATH="$HOME/bin:$PATH"
```

Shell Basics

```
$ cd /neuro/labs/grantlab/research

$ ls

$ ls -lh

$ cd ~

$ ls -lh
```

ssc && . neuro-fs stable

"Legacy" software definition: outdated, but we're stuck using it

```
$ ssc: system setup color
$ . neuro-fs stable
```

- Modifies $PATH to load *FreeSurfer*
- WARNING: this version of FreeSurfer is outdated!

The FNNDSC Network Filesystem (NFS)

- /neuro/users/<firstname.lastname>
  - Private-ish home directory
- /neuro/labs/grantlab/research
  - Lab shared data and programs

Shell Scripts

- Commands *are* a programming language!
  (but not a good one)

scriptname.sh

```
#!/bin/bash -e
count=$(find "$1" -type f -name '*.py' | wc -l)
echo "There are $count Python files in $1"
```

Best Practices

- Log out of your computer, don't lock screen (sorry, it's a bug)
- Don't delete files from /neuro/labs/grantlab/research
- Write documentation, preferably on http://rcwiki
- Share code on Github with MIT license https://github.com/FNNDSC

Beyond Linux

E2 SLURM Cluster

```
$ srun -A fnndsc -p bch-compute.     \
       --cpus-per-task=2 --mem=2g    \
       --time=30:00                  \
       --pty /bin/bash
```

https://ondemand-computing.tch.harvard.edu

http://rcwiki/display/RCK/HPC+E2+Cluster

ChRIS Research Integration System

- PACS query and retrieve
- Access Galena & E2 via website
- Reproducibility = data lineage + containerized software

Log in with your FNNDSC account at http://chris-next.tch.harvard.edu

Overview

Data

Library

PACS Query/Retrieve

Analysis

New and Existing Analyses

Plugins

Compute

Pipelines

New and Existing Analyses (5)

Create New Analysis

Select all

1 - 5 of 5

1 of 1

Search Analysis By Name

Name

Id	Analysis	Created	Creator	Run Time	Size	Sta...
59	Pull Marisol's subplate segmentations again, accounting for backwards left&right labels	15 Aug 2023, 15:06	jennings.zhang	01:09:16	11 GB	
20	Registration of medial cut masks to fetal WM surfaces	26 Jul 2023, 17:17	jennings.zhang	02:40:13	3 GB	
18	Inner SP fitting experiment	22 Jul 2023, 02:08	jennings.zhang	198:14:23	49 GB	
14	Extract inner SP masks	20 Jul 2023, 10:34	jennings.zhang	00:09:05	3 GB	

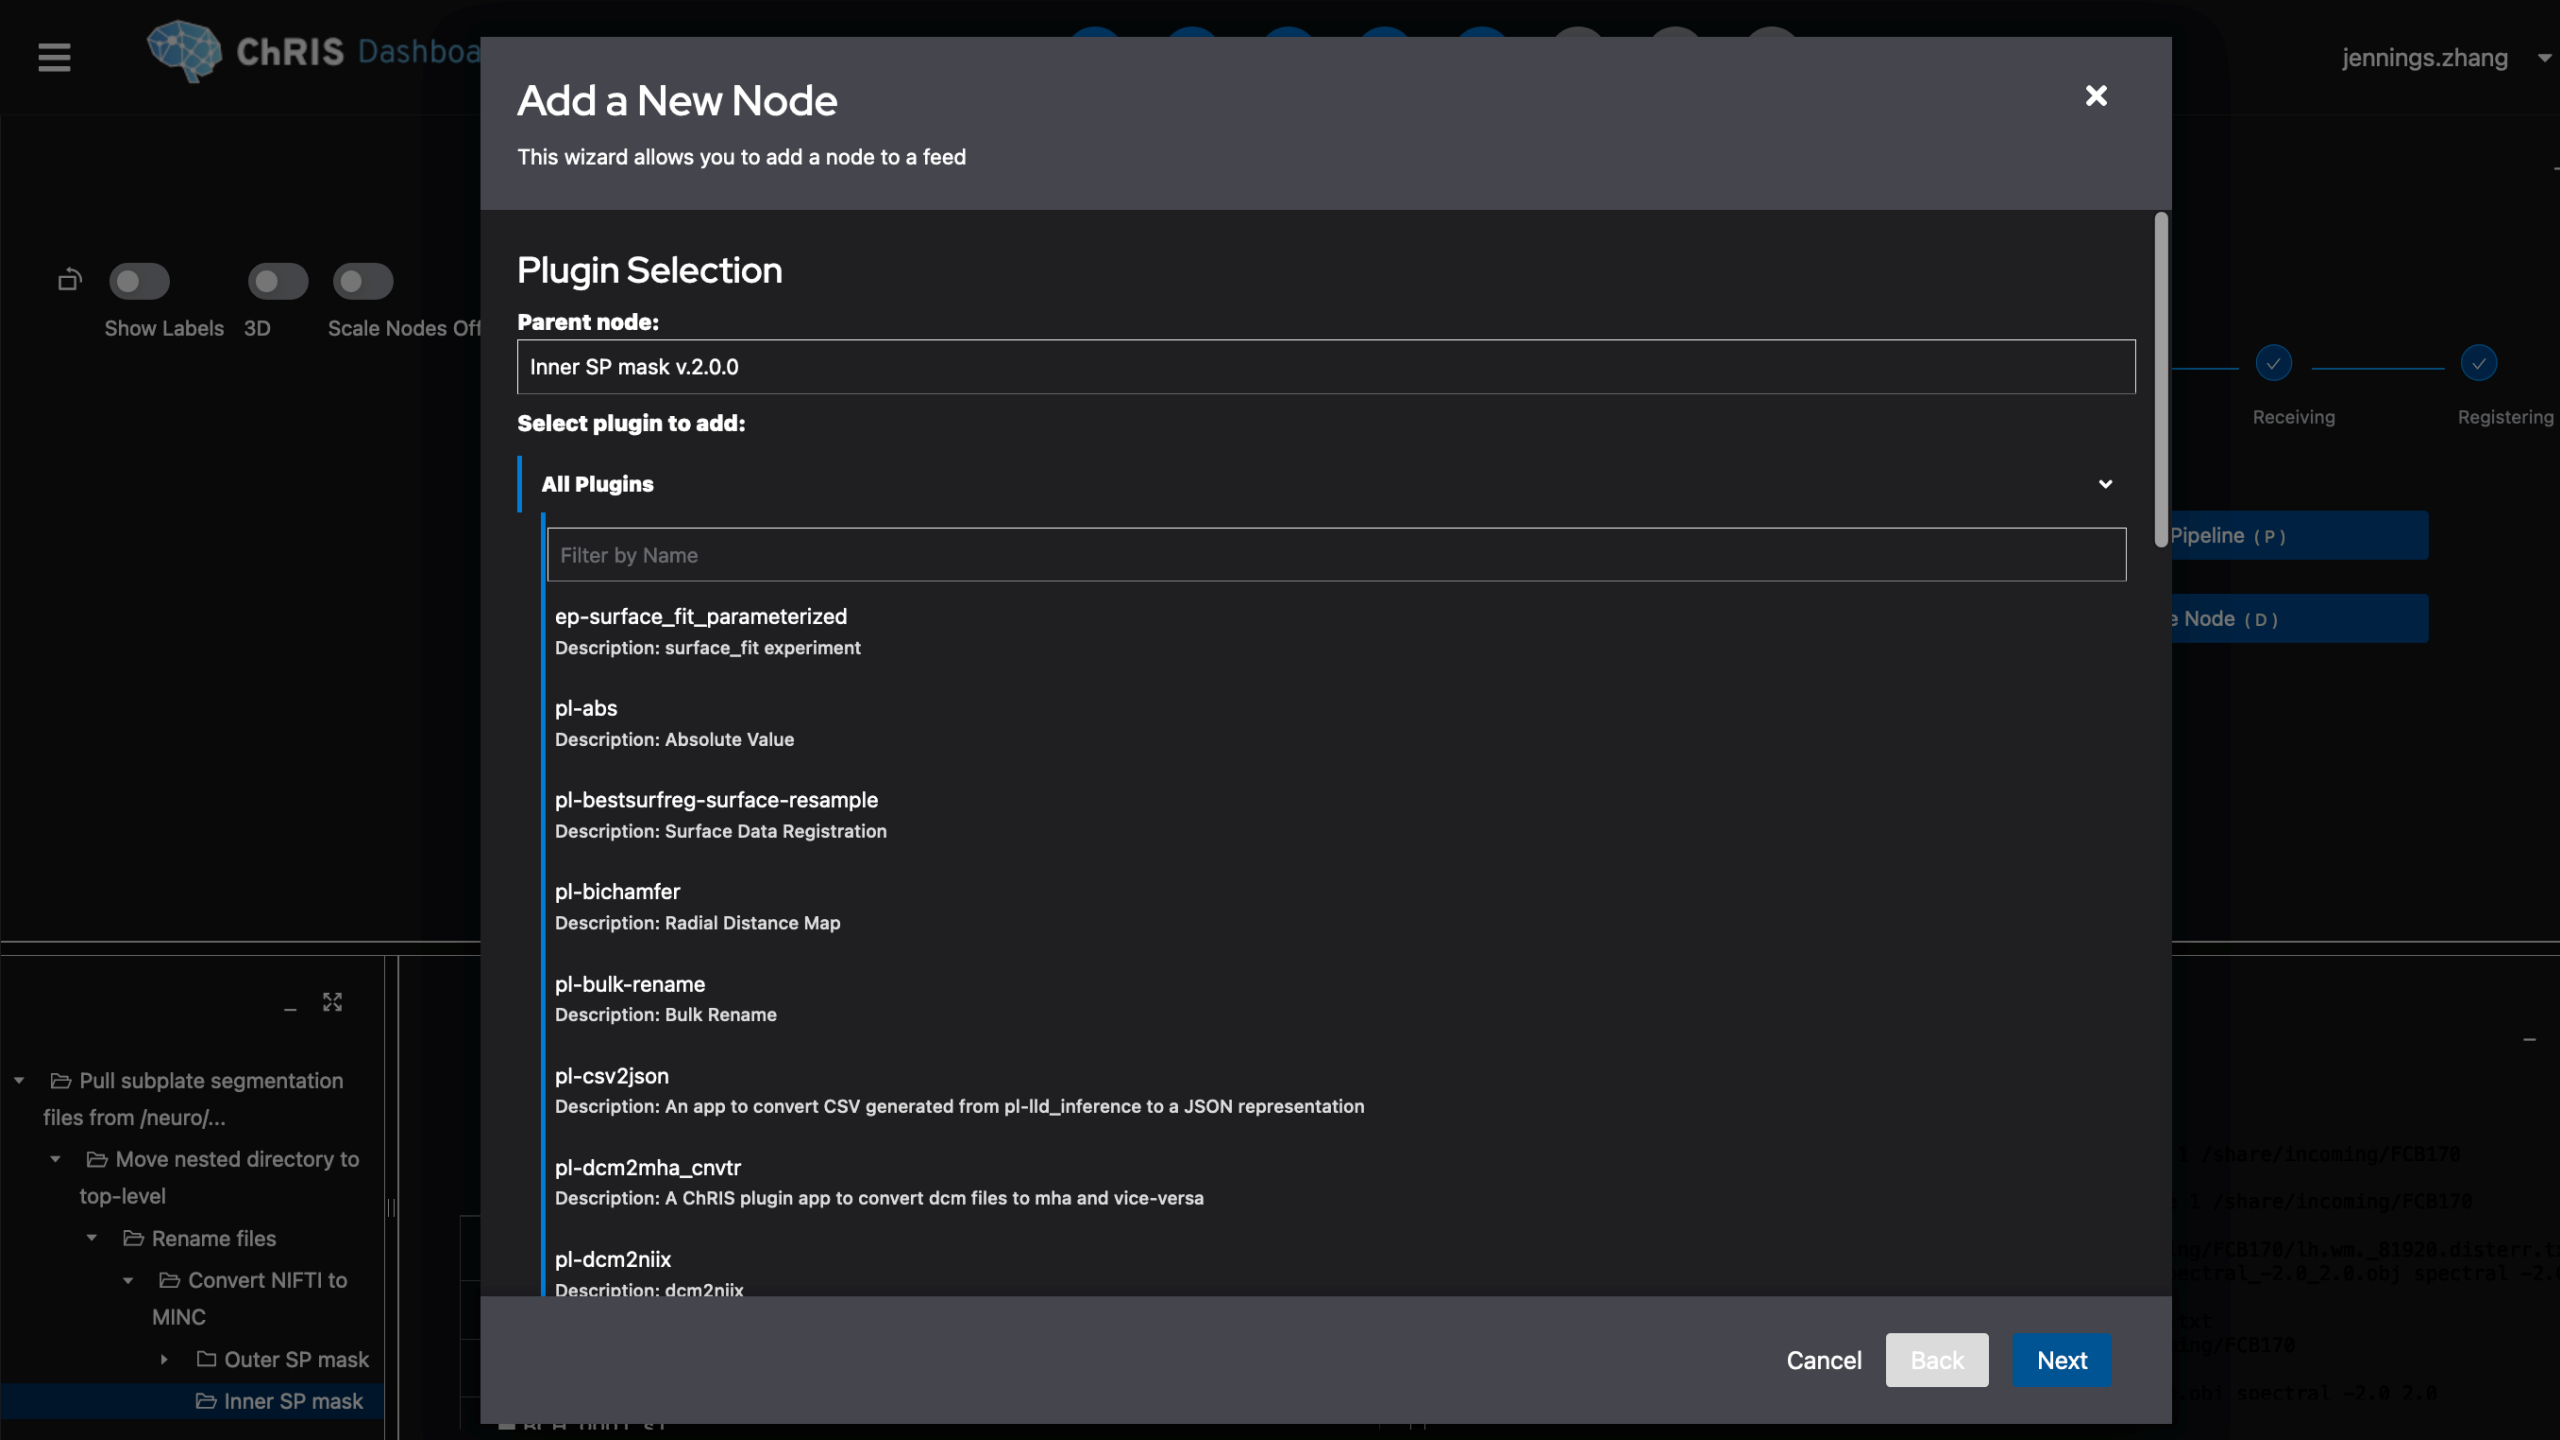
ChRIS Dashboa
jennings.zhang
Show Labels
3D
Scale Nodes Off
Receiving
Registering
Pipeline (P)
e Node (D)
Pull subplate segmentation files from /neuro/...
Move nested directory to top-level
Rename files
Convert NIFTI to MINC
Outer SP mask
Inner SP mask
Add a New Node
This wizard allows you to add a node to a feed
Plugin Selection
Parent node:
Inner SP mask v.2.0.0
Select plugin to add:
All Plugins
Filter by Name
ep-surface_fit_parameterized
Description: surface_fit experiment
pl-abs
Description: Absolute Value
pl-bestsurfreg-surface-resample
Description: Surface Data Registration
pl-bichamfer
Description: Radial Distance Map
pl-bulk-rename
Description: Bulk Rename
pl-csv2json
Description: An app to convert CSV generated from pl-lld_inference to a JSON representation
pl-dcm2mha_cnvtr
Description: A ChRIS plugin app to convert dcm files to mha and vice-versa
pl-dcm2niix
Description: dcm2niix
Cancel
Back
Next

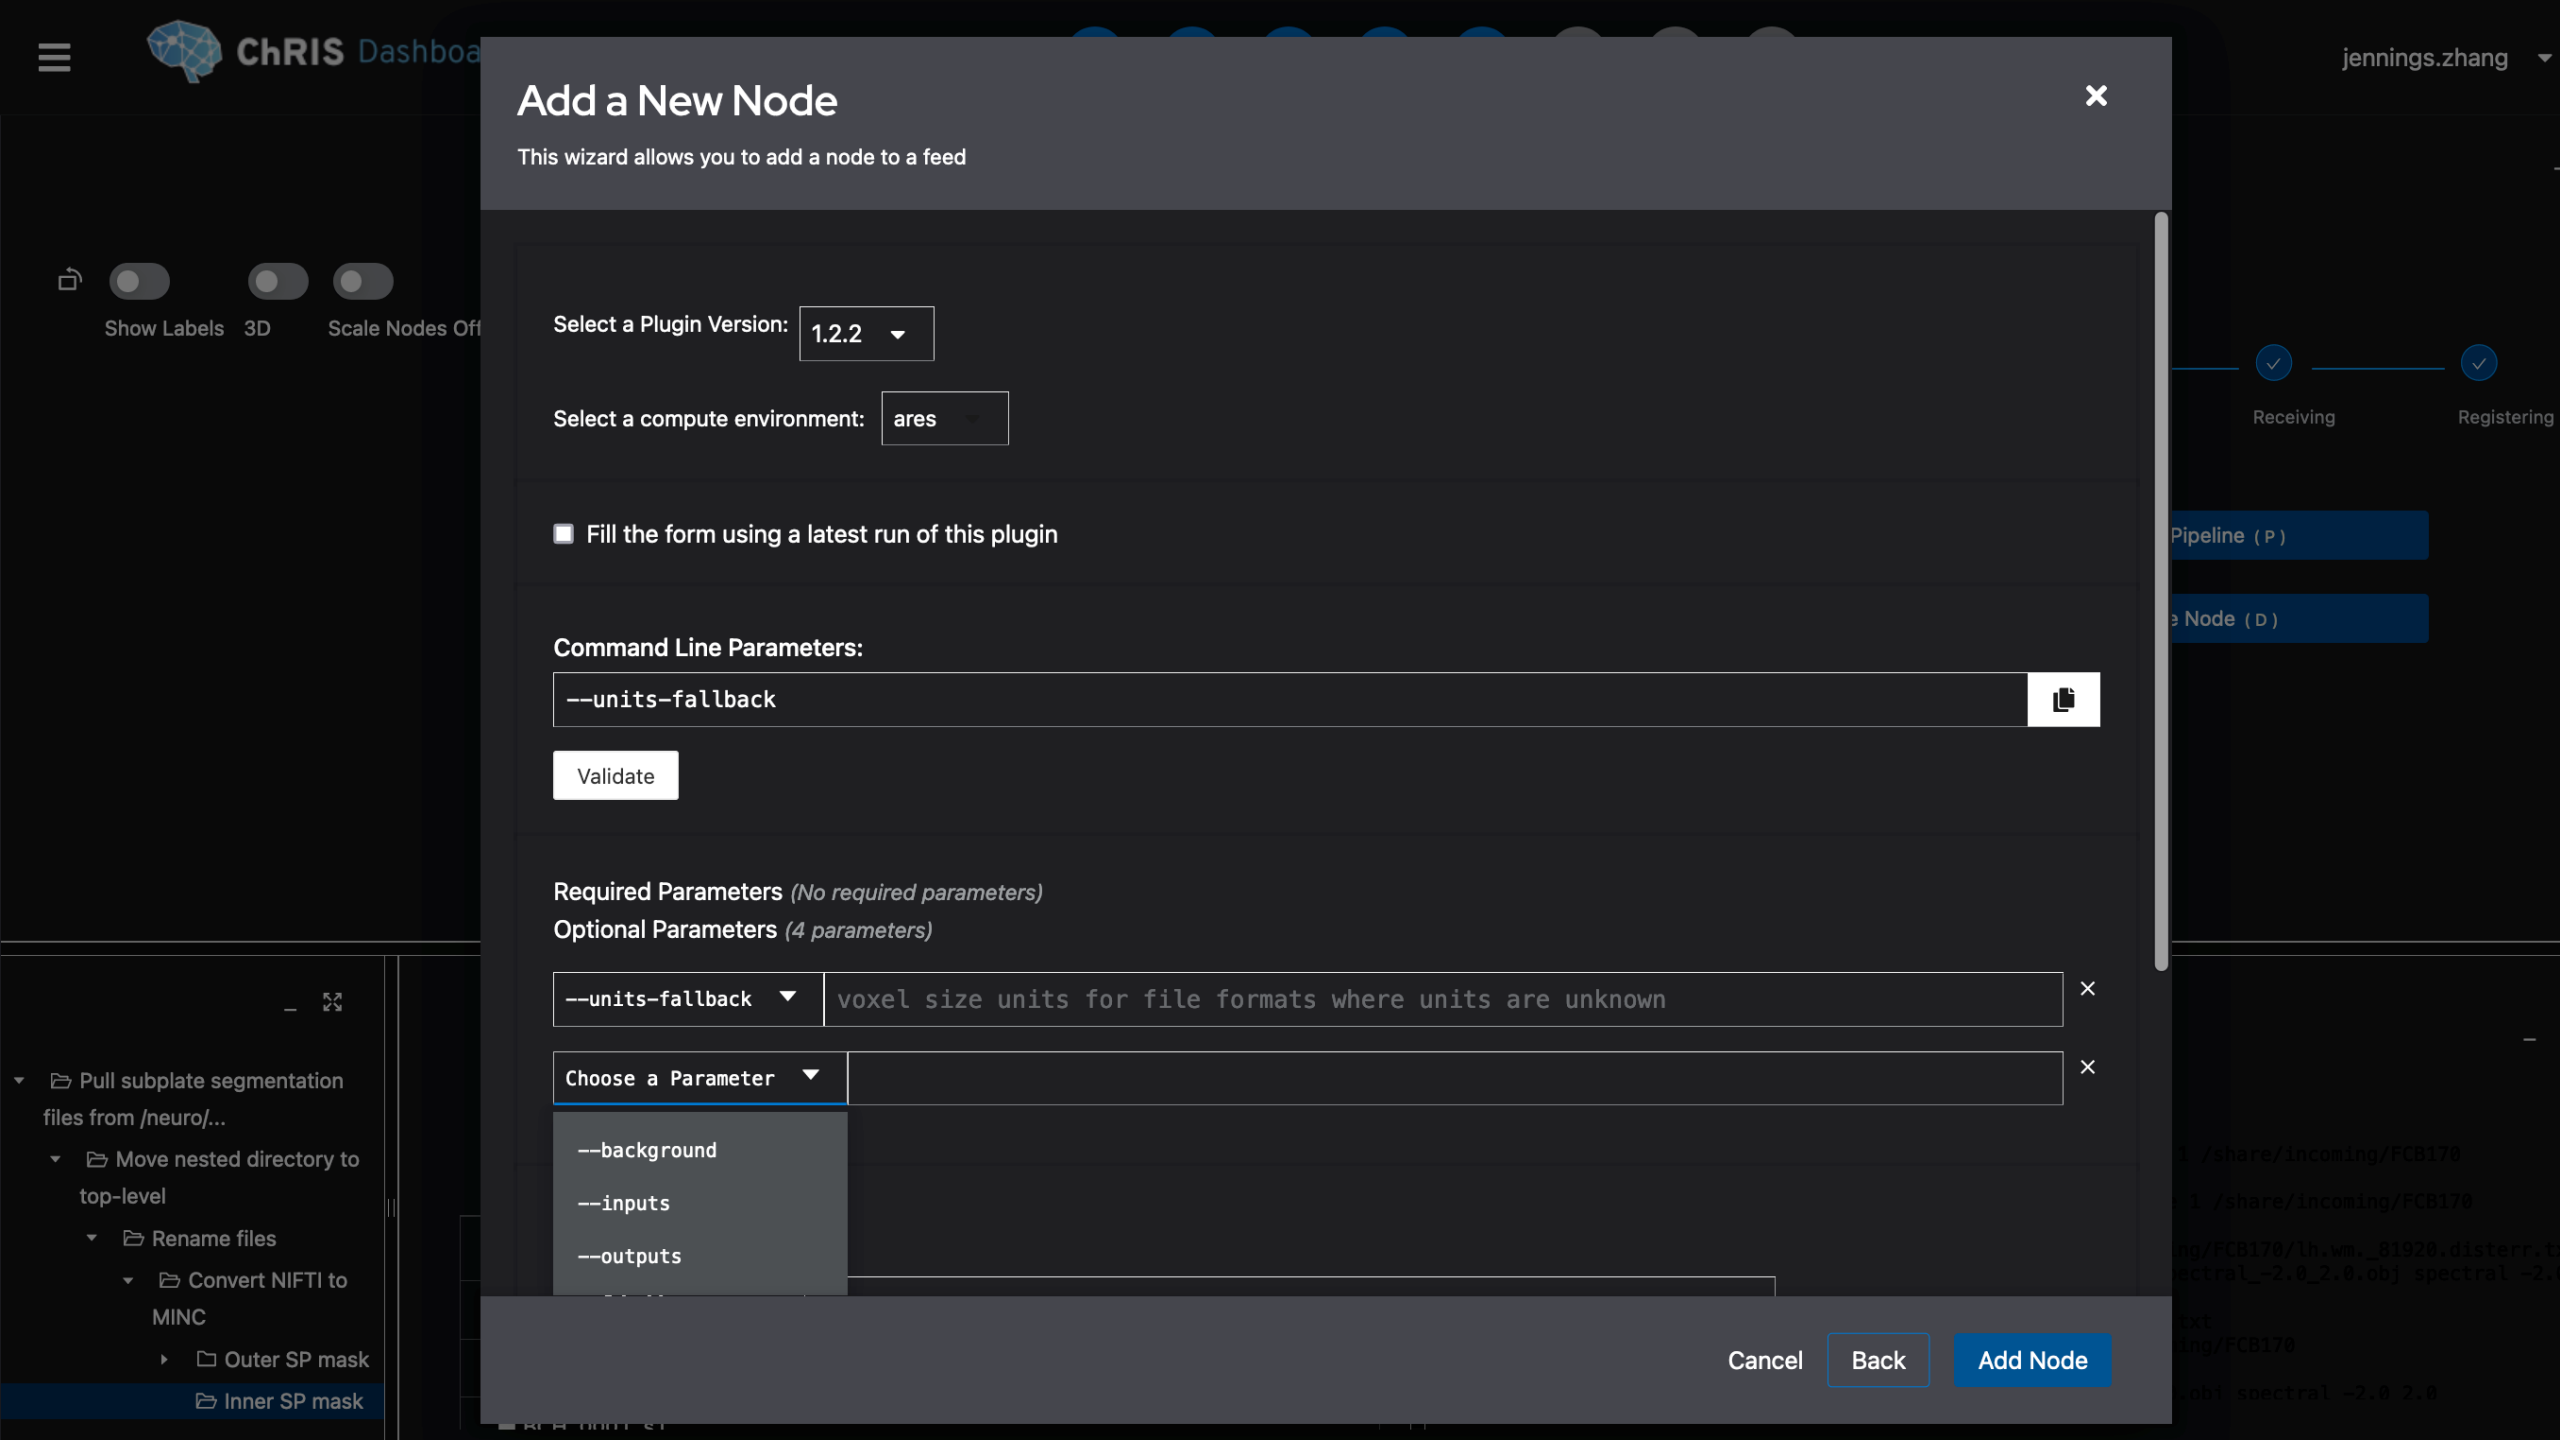
ChRIS Dashboa
jennings.zhang
Show Labels
3D
Scale Nodes Off
Receiving
Registering
Pipeline (P)
e Node (D)
Pull subplate segmentation files from /neuro/...
Move nested directory to top-level
Rename files
Convert NIFTI to MINC
Outer SP mask
Inner SP mask
Add a New Node
This wizard allows you to add a node to a feed
Select a Plugin Version:
1.2.2
Select a compute environment:
ares
Fill the form using a latest run of this plugin
Command Line Parameters:
--units-fallback
Validate
Required Parameters (No required parameters)
Optional Parameters (4 parameters)
--units-fallback
voxel size units for file formats where units are unknown
Choose a Parameter
--background
--inputs
--outputs
Cancel
Back
Add Node

ChRIS Dashboard

jennings.zhang

- Overview
- Data
- Library
- PACS Query/Retrieve
- Analysis
- New and Existing Analyses
- Plugins
- Compute
- Pipelines

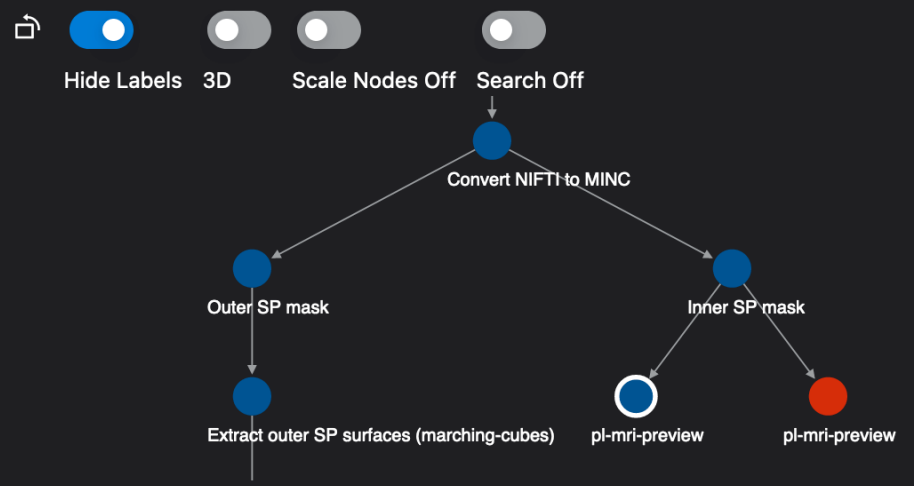

Search 0 matches

```
main:123 - input files: ['.nii', '.nii.gz', '.mnc', '.mgz']
main:124 - output formats: ['.png', '.txt']
main:125 - background threshold: 0.0
main:142 - /share/incoming/BCH_0024_s1/lh.spinner.mnc: 73395 voxels, volume=46581.6
main:146 -         -> /share/outgoing/BCH_0024_s1/lh.spinner.png
main:146 -         -> /share/outgoing/BCH_0024_s1/lh.spinner.txt
main:142 - /share/incoming/BCH_0024_s1/rh.spinner.mnc: 69048 voxels, volume=43822.7
main:146 -         -> /share/outgoing/BCH_0024_s1/rh.spinner.png
main:146 -         -> /share/outgoing/BCH_0024_s1/rh.spinner.txt
main:142 - /share/incoming/FCB035/lh.spinner.mnc: 63581 voxels, volume=40352.969646
main:146 -         -> /share/outgoing/FCB035/lh.spinner.png
main:146 -         -> /share/outgoing/FCB035/lh.spinner.txt
main:142 - /share/incoming/FCB035/rh.spinner.mnc: 61721 voxels, volume=39172.482967
main:146 -         -> /share/outgoing/FCB035/rh.spinner.png
main:146 -         -> /share/outgoing/FCB035/rh.spinner.txt
main:142 - /share/incoming/BCH_0055_s1/lh.spinner.mnc: 65821 voxels, volume=41774.6
```

- Pull subplate segmentation files from /neuro/...
  - Move nested directory to top-level
    - Rename files
      - Convert NIFTI to MINC
        - Outer SP mask
        - Inner SP mask

data › BCH_0005_s1

(4 items)

Name	Size
rh.spinner.txt	35 B
rh.spinner.png	19 KB
lh.spinner.txt	35 B
lh.spinner.png	19 KB

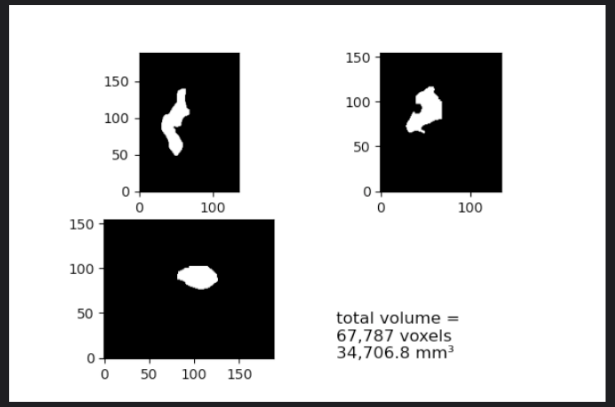

Thank you, happy computing!

- Questions?

Link to slides →