CONTAINERS FOR SCIENCE

What are they?

Why are they useful?

How to use them?

How to build them?

WHOAREWE

Tom Shaw
@thomcat992
t.shaw@uq.edu.au

Steffen Bollmann
@sbollmann_MRI
steffen.bollmann@cai.uq.edu.au

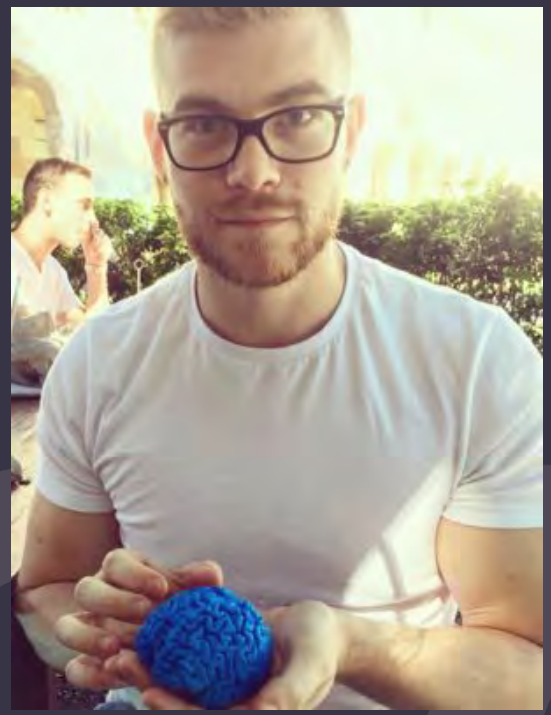

Centre for Advanced Imaging

The University of Queensland

Mattermost: https://tinyurl.com/helpdesk -docker
Jitsi room: https://tinyurl.com/jitsi -docker

The University Of Queensland
AUSTRALIA
Create change

CONTAINERS FOR SCIENCE

- **What are containers?**
- Why are containers useful?
- How can we use containers?
- How can we build containers?

POLL – WHAT ARE CONTAINERS?

a) A different name for a virtual machine

b) Plastic boxes to store food in and keep it separated

c) A collection of tools to keep programs from interfering with each other

WHAT ARE CONTAINERS?

- isolate software from its surroundings
- container image includes: code, runtime, system tools, system libraries, settings
- resource management provided by the Linux kernel (namespaces and cgroups)
- recipe = describes what should be in an image
- image = stores everything we need to run
- container = what we launch based on an image

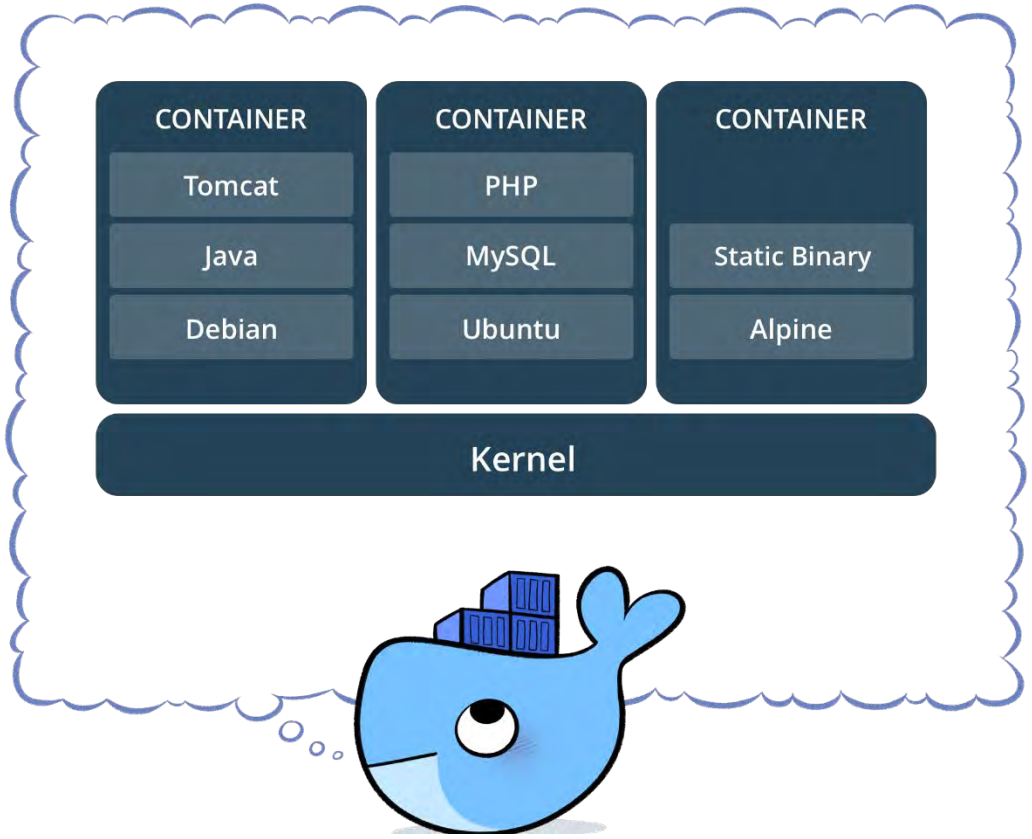

https://www.docker.com/what -container#/package_software

DOCKER

- started the container hype by providing easy to use packages for Linux, Windows, Mac
- widely adopted and supported by cloud providers, including orchestration of many containers (Kubernetes)

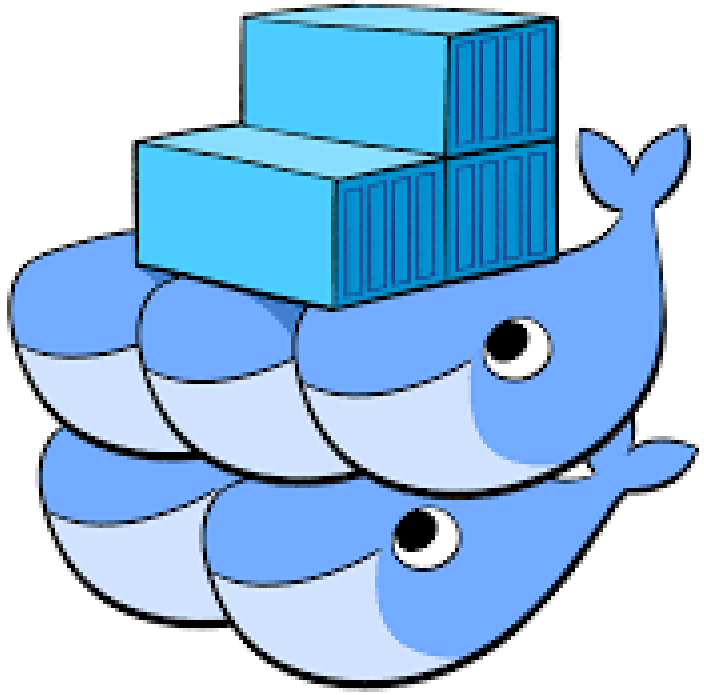

- not easy to run graphical applications ☹
- requires elevated privileges ☹

SINGULARITY

RESEARCH ARTICLE

Singularity: Scientific containers for mobility of compute

Gregory M. Kurtzer[1], Vanessa Sochat[2]*, Michael W. Bauer[1,3,4]

- Built to run on HPCs
  - simple to install
  - untrusted users running untrusted containers
  - same user inside image than outside
- built for data science
  - GPU devices can be accessed in container
  - graphical applications can run inside containers
- container format based on a single file
  - simple file transfer and archive

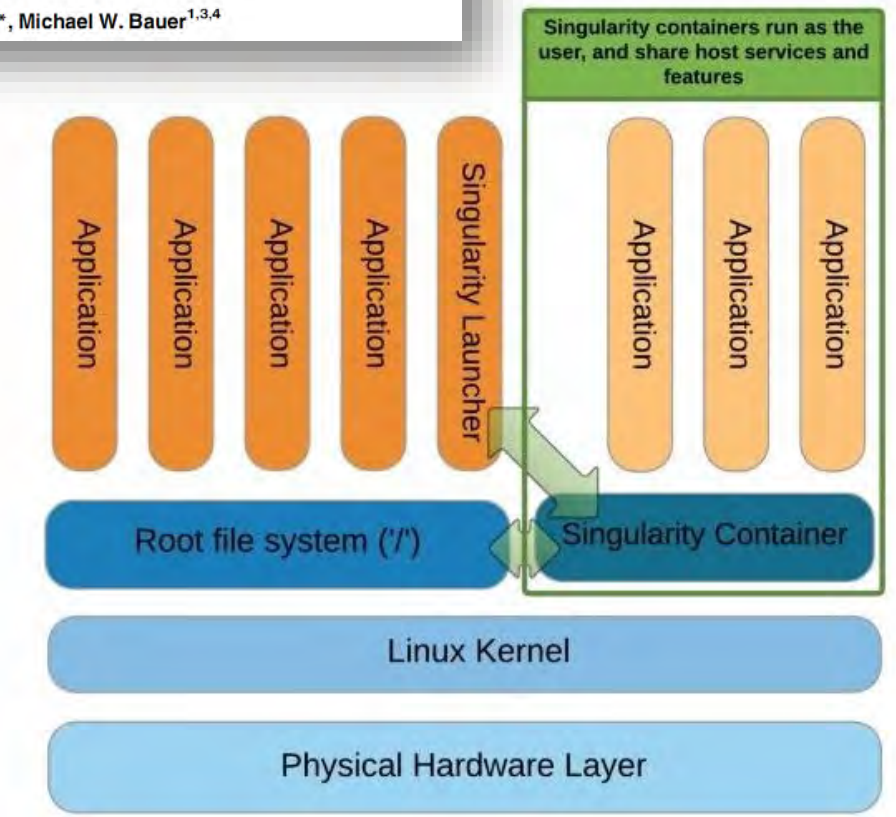

https://www.hpcwire.com/2017/05/04/singularity-hpc-container-technology-moves-lab/

National Imaging Facility

The University Of Queensland AUSTRALIA

Create change

VIRTUAL MACHINES VS CONTAINERS

Storage: 10 GB
Startup : 15s
RAM: 4GB

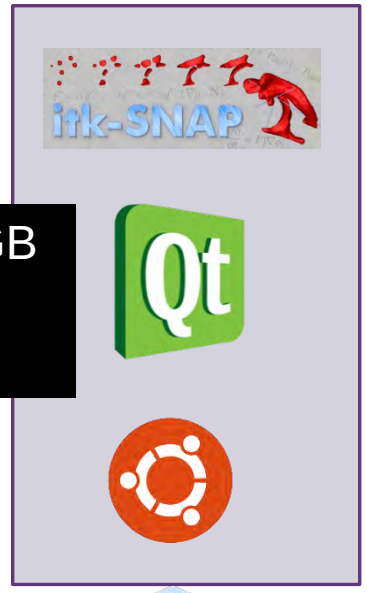

Application	• e.g. itksnap
Libraries	• e.g. QT4
Guest OS	• e.g. Ubuntu 16.04
Hypervisor	• e.g. Virtualbox
Host OS	• e.g. Centos 6
Hardware	• e.g. Dell Precision

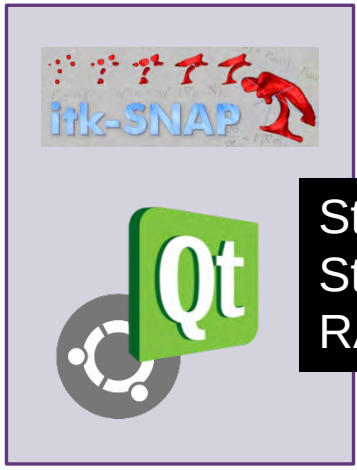

Storage: 0.1 GB
Startup : 0.2 s
RAM: 0.1 GB

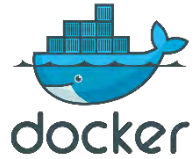

National Imaging Facility

The University Of Queensland AUSTRALIA
Create change

QUESTIONS?

CONTAINERS FOR SCIENCE

- What are containers?
- **Why are containers useful?**
- How can we use containers?
- How can we build containers?

WHY ARE CONTAINERS USEFUL?

Reproducibility of neuroimaging analyses across operating systems

Tristan Glatard[1,2], Lindsay B. Lewis[1], Rafael Ferreira da Silva[3], Reza Adalat[1], Natacha Beck[1], Claude Lepage[1], Pierre Rioux[1], Marc-Etienne Rousseau[1], Tarek Sherif[1], Ewa Deelman[3], Najmeh Khalili-Mahani[1] and Alan C. Evans[1]

```
expf(1.54051852226257324218750000000)
      =4.667009353637695312500000

expf(1.54051852226257324218750000000)
      =4.667009830474853515625000
```

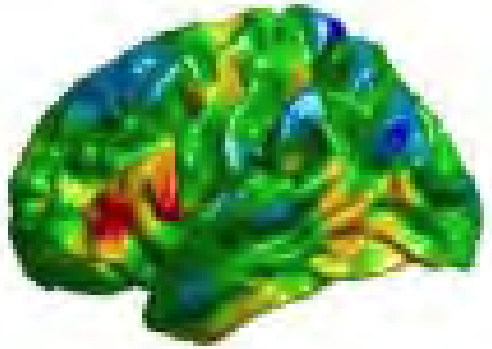

- glibc 2.5 vs 2.18 deliver different floating -point results
- leads to significant differences in long pipelines

National Imaging Facility

The University Of Queensland AUSTRALIA
Create change

USE CASE – RE-RUNNING FREESURFER ANALYSIS FROM 5 YEARS AGO

surfer.nmr.mgh.harvard.edu/fswiki/DownloadAndInstall

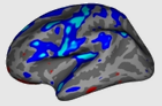

Login

DownloadAndInstall

FreeSurferWiki | RecentChanges | FindPage | HelpContents | DownloadAndInstall

Immutable Page Discussion Info Attachments More Actions:

FreeSurfer Download and Install

You can read and compare the version 7 and 6 release notes here.

Latest Version 7 Release (May 2020)

Freesurfer 7 release downloads and setup: here.

Previous Version 6 Release (Jan 2017)

Freesurfer 6 release downloads and setup: here.

Important Note: *When processing a group of subjects for your study, it is essential to process all your subjects w to ensure that results match across platforms, there are none-the-less system-level libraries that are OS depende. all stream) if you check with us first (for instance you may run the longitudinal processing with newer versions).*

Development Version: Daily builds of FreeSurfer can be downloaded from here.

Older Releases: Previous releases of FreeSurfer can be downloaded from here.

Freeview: For instructions on how to update Freeview - FreeSurfer's mri viewing application - go to the following

License

A license key must be obtained to make the FreeSurfer tools operational. Obtaining a license is free and comes in by the FREESURFER_HOME environment variable.

Follow this link to obtain a license key.

surfer.nmr.mgh.harvard.edu/pub/dist/freesurfer/

Name	Last modified	Size	Description
Parent Directory		-	
0.7.0/	11-Dec-2018 15:48	-	
0.8.0/	11-Dec-2018 17:20	-	
2.2.0/	21-Feb-2020 06:00	-	
3.0.3/	11-Dec-2018 17:20	-	
3.0.4/	27-Jan-2009 16:35	-	
3.0.5/	21-Feb-2020 06:03	-	
4.0.0/	21-Feb-2020 06:05	-	
4.0.1/	27-Jan-2009 15:41	-	
4.0.2/	27-Jan-2009 14:15	-	
4.0.3/	27-Jan-2009 14:05	-	
4.0.4/	27-Jan-2009 13:39	-	
4.0.5/	27-Jan-2009 13:22	-	
4.1.0/	21-Feb-2020 06:08	-	
4.2.0/	21-Feb-2020 06:11	-	
4.3.0/	21-Feb-2020 06:13	-	
4.3.1/	21-Feb-2020 06:15	-	
4.4.0/	21-Feb-2020 06:17	-	
4.5.0/	21-Feb-2020 06:44	-	
5.0.0/	11-Dec-2018 17:21	-	
5.1.0/	11-Dec-2018 17:23	-	
5.2.0-deprecated/	14-Jun-2013 12:04	-	
5.3.0-HCP/	21-Feb-2020 06:49	-	
5.3.0-patch/	26-Aug-2014 14:42	-	
5.3.0-tracula-addons/	25-Aug-2014 05:20	-	
5.3.0/	11-Dec-2018 17:33	-	
6.0.0-patch/	05-Dec-2019 16:52	-	
6.0.0/	04-Sep-2018 18:08	-	
6.0.1/	11-Dec-2018 15:50	-	
7.0.0-beta/	04-Jun-2020 03:31	-	
7.0.0/	03-May-2020 20:14	-	
7.1-tracula-beta/	06-Jun-2020 07:01	-	
7.1.0/	11-May-2020 20:30	-	

surfer.nmr.mgh.harvard.edu/pub/dist/freesurfer/5.3.0/

Index of /pub/dist/freesurfer/5.3.0

Name	Last modified	Size	Description
Parent Directory		-	
Freeview_v53_lion.dmg	18-Jul-2013 10:38	41M	
Freeview_v53_snowleopard.dmg	18-Jul-2013 10:38	40M	
freesurfer-Darwin-lion-stable-pub-v5.3.0-HCP.dmg	11-Dec-2018 17:28	3.4G	
freesurfer-Darwin-lion-stable-pub-v5.3.0.dmg	15-May-2013 00:02	3.4G	
freesurfer-Darwin-lion-stable-pub-v5.3.0.tar.gz	23-Oct-2013 16:11	3.4G	
freesurfer-Darwin-snowleopard-i686-stable-pub-v5.3.0-HCP.dmg	11-Dec-2018 17:29	3.2G	
freesurfer-Darwin-snowleopard-i686-stable-pub-v5.3.0.dmg	14-May-2013 14:36	3.2G	
freesurfer-Darwin-snowleopard-i686-stable-pub-v5.3.0.tar.gz	23-Oct-2013 16:13	3.2G	
freesurfer-Linux-centos4-stable-pub-v5.3.0.tar.gz	14-May-2013 13:22	3.8G	
freesurfer-Linux-centos4_x86_64-stable-pub-v5.3.0.tar.gz	14-May-2013 13:24	4.0G	
freesurfer-Linux-centos6_x86_64-stable-pub-v5.3.0.tar.gz	14-May-2013 13:25	4.1G	
freesurfer-Virtualbox-linux-x86-stable-pub-v5.3-full.vdi.gz	15-Aug-2013 12:49	9.0G	

Apache/2.2.15 (CentOS) Server at surfer.nmr.mgh.harvard.edu Port 443

National Imaging Facility

The University Of Queensland AUSTRALIA

Create change

USE CASES: SCALABILITY

- **Example: Running fMRIPrep on your pilot data**
- Your pilot data turns into 100 participants.
- OR you have to include a larger dataset in your publication
- Use the same software on different platforms (your notebook/PC, Linux workstations, different clusters/ high performance computers)
- Run the same pipeline with no changes. Hooray!

https://i.guim.co.uk/img/media/

https://rcc.uq.edu.au/article/2017/10/new -hpc -qld -research -goodbye -euramoo -hello -awoonga

USE CASES: RUNNING THINGS ON DIFFERENT OS

- **Example: Upgrade your computer! Now what?**
- Different OS?
- Run Linux software on Mac and Windows
- Containers make your work reproducible and robust to OS changes.

USE CASES: SHARING

- **Example: Sharing a reproducible pipeline including the software and the data!**

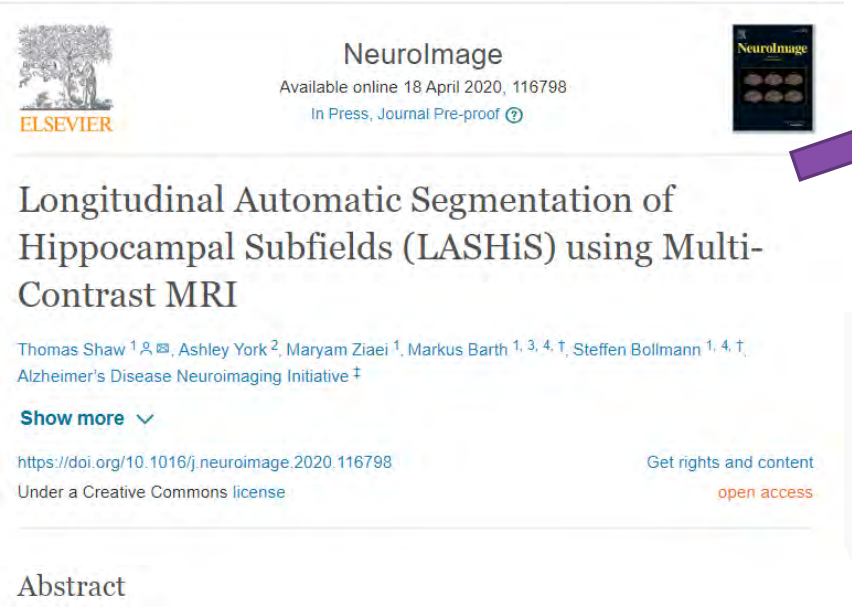

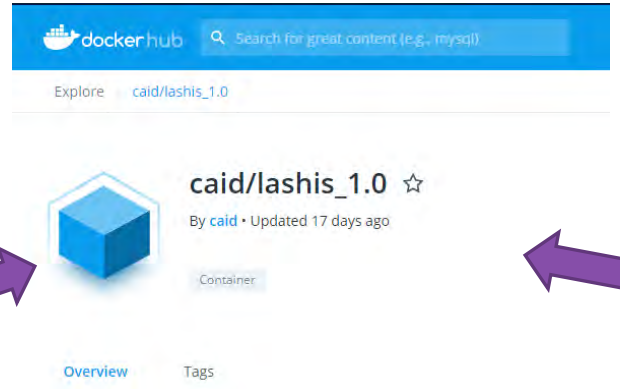

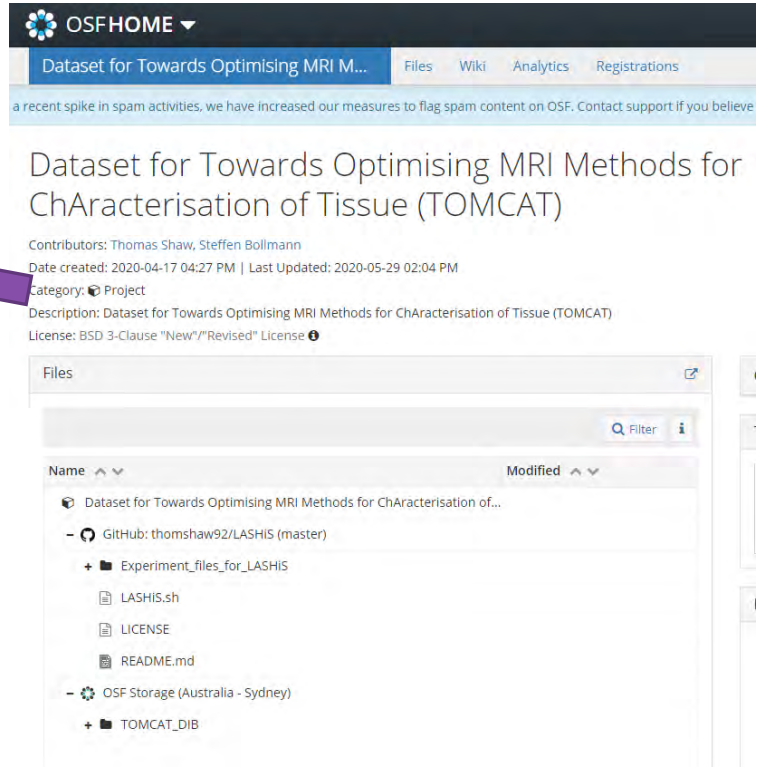

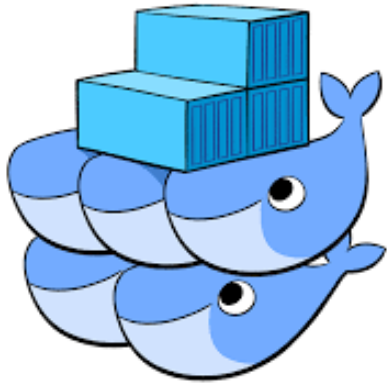

National Imaging Facility

The University Of Queensland AUSTRALIA
Create change

BENEFITS

- Longitudinal stability of software pipeline (e.g., upgrade of Ubuntu 16.04 breaks fslview -> libpng12 gone)
- Reproducible? Yes and no: version of container can change. Is the container build reproducible?
- Portable
- Isolated – are not exposed to the host system
- Ease of use (e.g., compiling ANTs on a HPC without sudo access vs using compiled ANTs container)
- Some software is easier to use in containers (e.g., fmriprep , mriqc , BIDSApps / integration.)

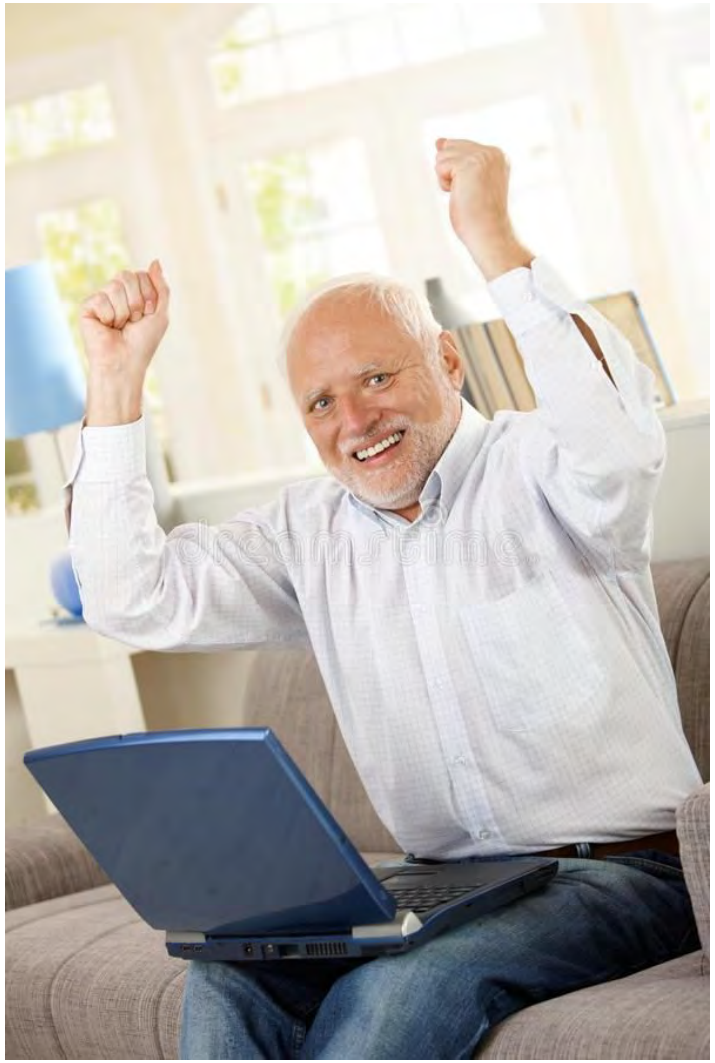

CHALLENGES

- Hard to use – bind mounts/syntax of commands
- Docker doesn't run with HPCs
- Versioning of containers comes with no guarantees – images may not exist tomorrow, hubs disappear
- Licensing issues (Freesurfer/ FSL/ MATLAB)
- Black box – how was it made?
- Version of Singularity or Docker changes

QUESTIONS?

CONTAINERS FOR SCIENCE

- What are containers?
- Why are containers useful?
- **How can we use containers?**
- How can we build containers?

BRAIN IMAGING DATA STRUCTURE

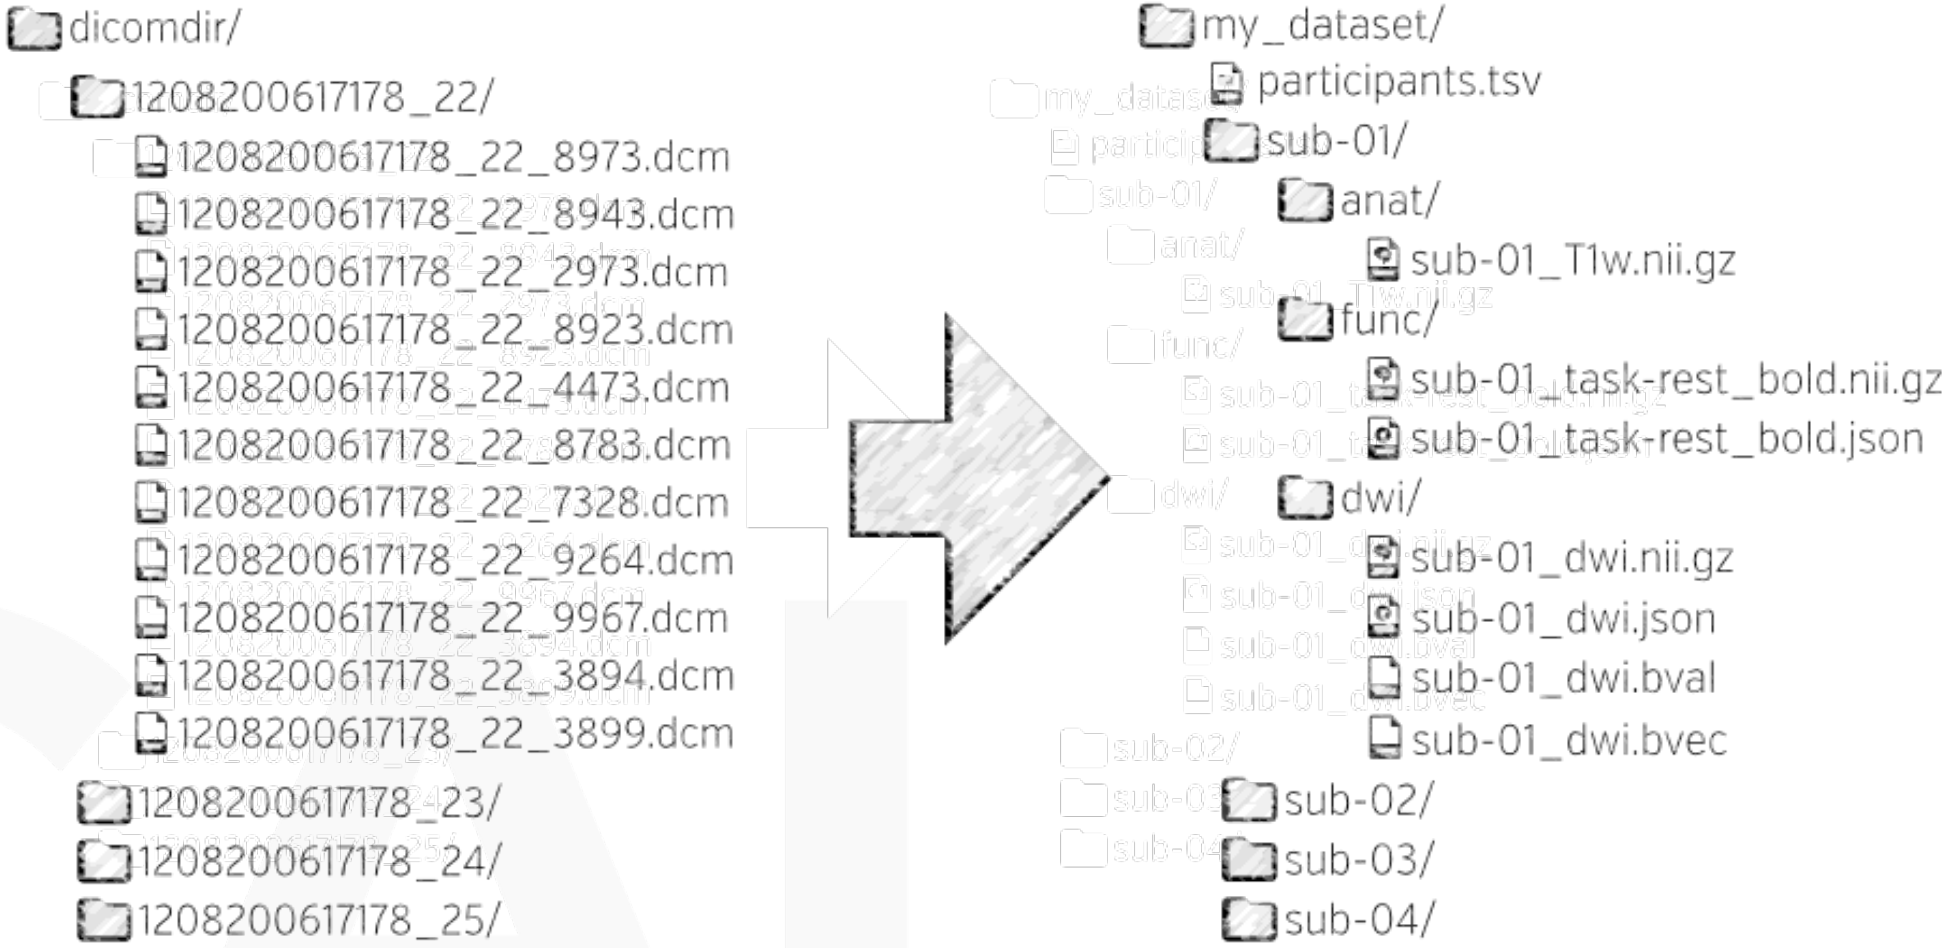

http://bids.neuroimaging.io/

BIDS APPS

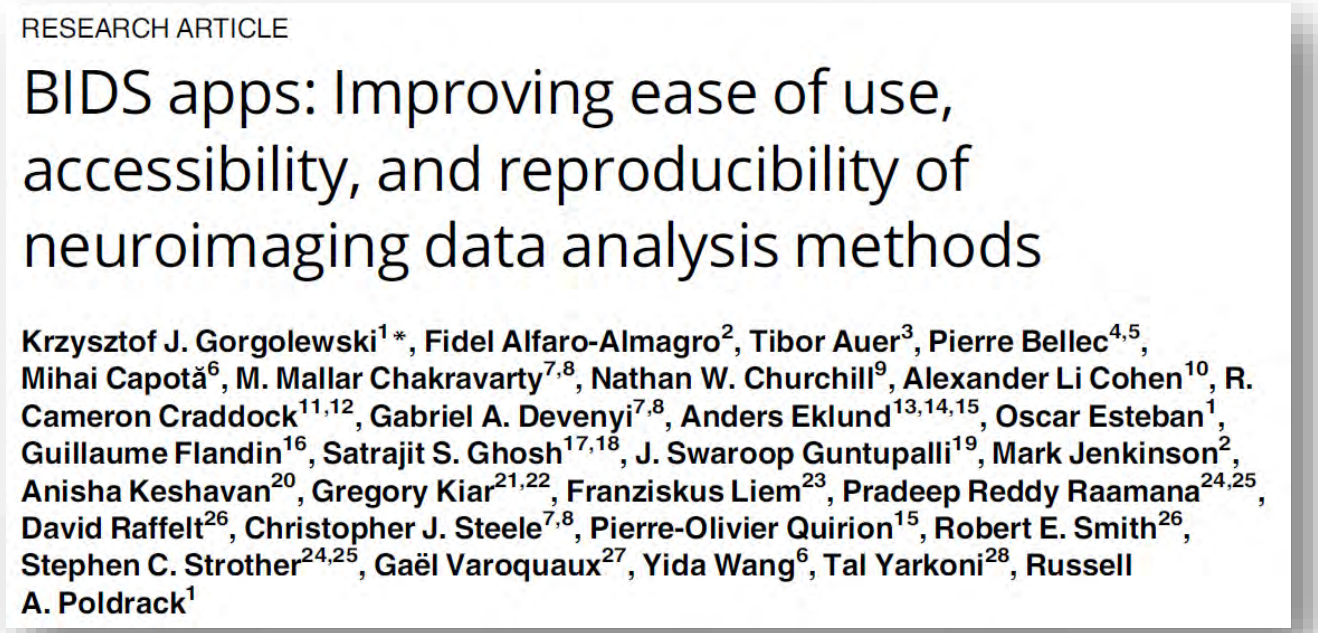

RESEARCH ARTICLE

BIDS apps: Improving ease of use, accessibility, and reproducibility of neuroimaging data analysis methods

Krzysztof J. Gorgolewski[1]*, Fidel Alfaro-Almagro[2], Tibor Auer[3], Pierre Bellec[4,5], Mihai Capotă[6], M. Mallar Chakravarty[7,8], Nathan W. Churchill[9], Alexander Li Cohen[10], R. Cameron Craddock[11,12], Gabriel A. Devenyi[7,8], Anders Eklund[13,14,15], Oscar Esteban[1], Guillaume Flandin[16], Satrajit S. Ghosh[17,18], J. Swaroop Guntupalli[19], Mark Jenkinson[2], Anisha Keshavan[20], Gregory Kiar[21,22], Franziskus Liem[23], Pradeep Reddy Raamana[24,25], David Raffelt[26], Christopher J. Steele[7,8], Pierre-Olivier Quirion[15], Robert E. Smith[26], Stephen C. Strother[24,25], Gaël Varoquaux[27], Yida Wang[6], Tal Yarkoni[28], Russell A. Poldrack[1]

- neuroimaging tools and workflows in docker containers
- e.g. FreeSurfer, SPM, MRtrix3, AFNI, ANTS, HCPPipelines ....

https://github.com/BIDS -Apps

National Imaging Facility

The University Of Queensland AUSTRALIA

Create change

CONTAINERS FOR SCIENCE

- What are containers?
- Why are containers useful?
- How can we use containers?
- **How can we build containers?**

- https://tinyurl.com/ohbm2020 -containers

HELLO WORLD OF DOCKER

```
docker run hello-world
```

```
Unable to find image 'hello-world:latest' locally
latest: Pulling from library/hello-world
0e03bdcc26d7: Already exists
Digest: sha256:d58e752213a51785838f9eed2b7a498ffa1cb3aa7f946dda11af39286c3db9a9
Status: Downloaded newer image for hello-world:latest

Hello from Docker!
This message shows that your installation appears to be working correctly.
```

DOCKER CACHES IMAGES

- Does not download image again!

```
PS C:\Users\uqsbollm> docker run hello-world

Hello from Docker!
This message shows that your installation appears to be working correctly.
```

SHOW DOWNLOADED IMAGES

- This can fill up your hard drive ...

```
PS C:\Users\uqsbollm> docker images
REPOSITORY              TAG                 IMAGE ID            CREATED             SIZE
<none>                  <none>              9fdba2ebb1a4        19 hours ago        125MB
ubuntu                  16.04               005d2078bdfa        7 weeks ago         125MB
gigantum/labmanager     fa7d5e79            ec37c9898625        4 months ago        962MB
hello-world             latest              bf756fb1ae65        5 months ago        13.3kB
```

in windows all docker images are stored in a single hyper-v virtual machine disk at:
C:\ProgramData\DockerDesktop\vm-data\DockerDesktop.vhdx

CLEAN UP DOCKER IMAGES

```
PS C:\Users\uqsbollm> docker rmi -f hello-world
Untagged: hello-world:latest
Untagged: hello-world@sha256:d58e752213a51785838f9eed2b7a498ffa1cb3aa7f946dda11af39286c3db9a9
Deleted: sha256:bf756fb1ae65adf866bd8c456593cd24beb6a0a061dedf42b26a993176745f6b
```

```
PS C:\Users\uqsbollm> docker images
REPOSITORY              TAG                 IMAGE ID            CREATED             SIZE
<none>                  <none>              9fdba2ebb1a4        19 hours ago        125MB
ubuntu                  16.04               005d2078bdfa        7 weeks ago         125MB
gigantum/labmanager     fa7d5e79            ec37c9898625        4 months ago        962MB
```

National Imaging Facility

The University Of Queensland AUSTRALIA

Create change

FMRIPREP EXAMPLE

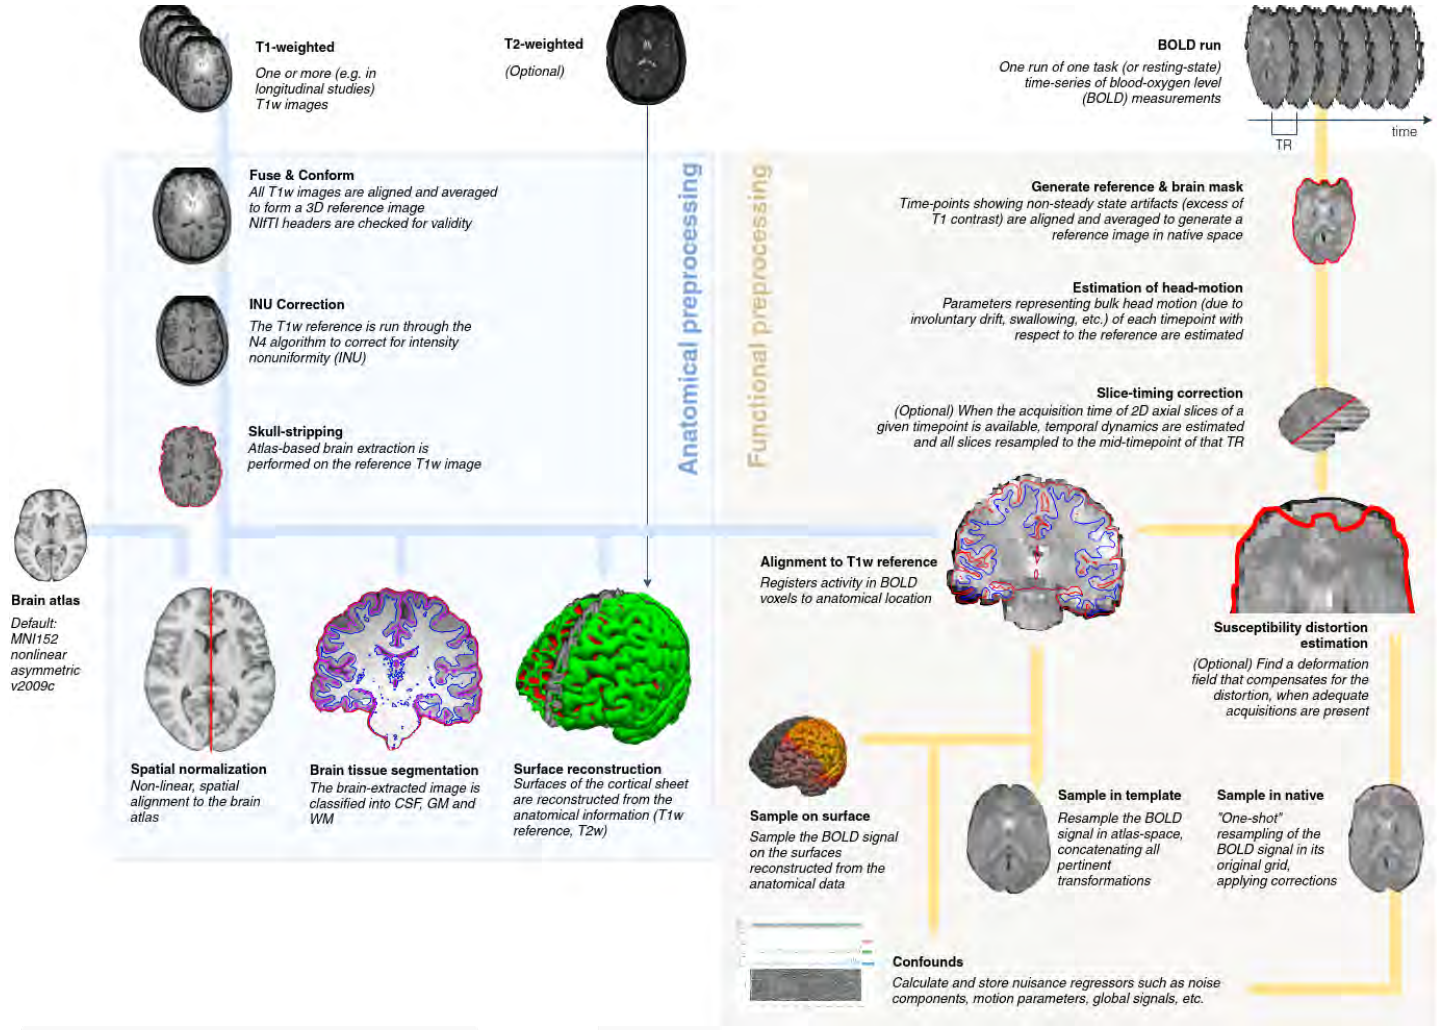

FMRIPREP EXAMPLE

```
PS C:\repos\TrainTrack_containers_2020\1_use_docker> docker run -ti --rm `
>> -v $HOME/data/BrainHackOHBM2020:/data:ro `
>> -v $HOME/data/derivatives:/out `
>> -v $HOME/data/work:/work `
>> poldracklab/fmriprep:latest /data /out/fmriprep-latest `
>> participant -w /work
```

FMRIPREP EXAMPLE

```
bids-validator@1.4.0

        1: [WARN] You should define 'SliceTiming' for this file. If you don't provide this information slic
e time correction will not be possible. (code: 13 - SLICE_TIMING_NOT_DEFINED)
                ./sub-01/func/sub-01_task-functionallocalizer_run-01_bold.nii.gz
                ./sub-01/func/sub-01_task-view_run-01_bold.nii.gz
                ./sub-01/func/sub-01_task-view_run-02_bold.nii.gz
                ./sub-02/func/sub-02_task-functionallocalizer_run-01_bold.nii.gz
                ./sub-02/func/sub-02_task-view_run-01_bold.nii.gz
                ./sub-02/func/sub-02_task-view_run-02_bold.nii.gz
                ./sub-03/func/sub-03_task-functionallocalizer_run-01_bold.nii.gz
                ./sub-03/func/sub-03_task-view_run-01_bold.nii.gz
                ./sub-03/func/sub-03_task-view_run-02_bold.nii.gz
                ./sub-04/func/sub-04_task-functionallocalizer_run-01_bold.nii.gz
                ... and 23 more files having this issue (Use --verbose to see them all).

        Please visit https://neurostars.org/search?q=SLICE_TIMING_NOT_DEFINED for existing conversations ab
out this issue.


        Summary:                Available Tasks:        Available Modalities:
        53 Files, 2.77GB        Functional localizer    T1w
        11 - Subjects           Passive Viewing         bold
        1 - Session


        If you have any questions, please post on https://neurostars.org/tags/bids.

200616-00:20:28,20 nipype.workflow IMPORTANT:

    Running fMRIPREP version 20.1.1:
      * BIDS dataset path: /data.
      * Participant list: ['01', '02', '03', '04', '05', '06', '07', '08', '09', '10', '11'].
      * Run identifier: 20200616-002015_dd0a3499-2db2-466a-a999-b3b075e206ae.
      * Output spaces: MNI152NLin2009cAsym:res-native.
      * Pre-run FreeSurfer's SUBJECTS_DIR: /out/fmriprep-latest/freesurfer.
```

PULL FMRIPREP TO HPC

- **I want to run fMRIPrep on the HPC to impress my boss**
- Docker doesn't run – need Singularity
- Push to docker hub and pull on HPC (so you can run the software **without root access)**
- Or pull an image from docker/singularity hub

```
uqtshaw@awoonga1:/90days/uqtshaw/TrainTrack_containers_2020> mkdir ./images
uqtshaw@awoonga1:/90days/uqtshaw/TrainTrack_containers_2020> singularity build ./images/fmriprep-latest.simg docker://poldracklab/fmriprep:latest
INFO:    Starting build...
Getting image source signatures
Copying blob 0a01a72a686c done
Copying blob cc899a5544da done
Copying blob 19197c550755 done
Copying blob 716d454e56b6 done
Copying blob b5bf898e214a done
Copying blob 42da0942cc0e [=========>----------------------------] 38.8MiB / 147.8MiB
```

OUTPUT TERMINAL DEBUG CONSOLE

```
2020/06/16 12:46:12  info unpack layer: sha256:a6b9c2f8b7ca1854bf38a928d60db0c8defea8407c3204fd01b4b6b8734ef4f6
2020/06/16 12:46:12  info unpack layer: sha256:be6e63d21e8caca0617df2e171de1934ff0aa15cb8dd5076a2cb65207a7e1e4e
2020/06/16 12:46:29  info unpack layer: sha256:02f9e6349e696228cea0a1caa4b41736b43fecc29a5111d69ba4ae9a78056790
2020/06/16 12:46:30  info unpack layer: sha256:f4b8d1e2864cc5bf3fc5b546eacc2c4b68ccad25f1c5438d31c779ffc09ae667
2020/06/16 12:47:07  info unpack layer: sha256:2baf8abf2b02dbc4babff3efbce47af04e43f328f11af8da693263268e68ac78
2020/06/16 12:47:07  info unpack layer: sha256:37d4bcbd72aba20c9e6a78602f0c060e573bd227aacf240add0854ef4883171a
2020/06/16 12:47:10  info unpack layer: sha256:0da53f606a4ee7574e8149a1d87ca08f520ba6e56b3c87638653536d113cb357
INFO:    Creating SIF file...
INFO:    Build complete: ./images/fmriprep-latest.simg
```

National Imaging Facility

The University of Queensland AUSTRALIA
Create change

QUESTIONS

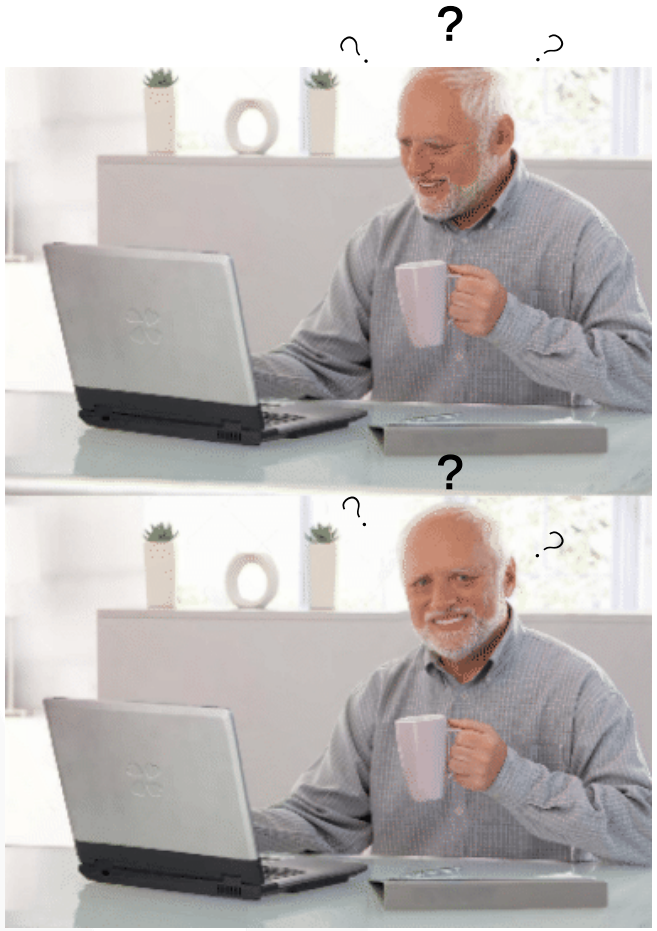

A SIMPLE DOCKER CONTAINER

A SIMPLE DOCKER CONTAINER

```
PS C:\repos\TrainTrack_containers_2020\2_build_itksnap_docker> docker build .
Sending build context to Docker daemon  5.632kB
Step 1/9 : FROM ubuntu:16.04
 ---> 005d2078bdfa
Step 2/9 : LABEL maintainer="Thom Shaw"
 ---> Using cache
 ---> c3f25d7ddfdd
Step 3/9 : LABEL org.label-schema.maintainer="Thom Shaw"
 ---> Using cache
 ---> 159ed9b7f638
Step 4/9 : ENV PATH="/opt/itksnap/bin/:${PATH}"
 ---> Using cache
 ---> 1dae10a5d039
Step 5/9 : ENV LD_LIBRARY_PATH=/opt/itksnap/lib/:${LD_LIBRARY_PATH}
 ---> Using cache
 ---> 82af4c8fae03
Step 6/9 : RUN apt-get update -y && apt-get install -y wget libglu1 libcurl4-openssl-dev libsm6 libxt6 libf
reetype6 libxrender1 libfontconfig1 libglib2.0-0 libqt4-dev libgtk2.0-dev && wget -O itksnap.tar.gz 'https:
//sourceforge.net/projects/itk-snap/files/itk-snap/Nightly/itksnap-nightly-master-Linux-gcc64-qt4.tar.gz/do
wnload' && tar -zxf itksnap.tar.gz -C /opt/ && mv /opt/itksnap-*/ /opt/itksnap/ && rm itksnap.tar.gz && use
radd -m -s /bin/bash itksnap
 ---> Using cache
 ---> 99c17fa41e18
Step 7/9 : USER itksnap
 ---> Using cache
 ---> bcc94c35bd2b
Step 8/9 : COPY --chown=itksnap:itksnap UserPreferences.xml /home/itksnap/.itksnap.org/ITK-SNAP/
 ---> Using cache
 ---> aa87bda36826
Step 9/9 : CMD ["itksnap"]
 ---> Using cache
 ---> 440f73a1ee8d
Successfully built 440f73a1ee8d
```

RUN THE CONTAINER INTERACTIVELY

- Running graphical applications requires some wizardry
- We will keep things simple here by going into the container

```
PS C:\Users\thoma> docker run -it thomshaw92/itksnap bash
itksnap@e240848aeced:/$ export DISPLAY=192.168.1.108:1.0
itksnap@e240848aeced:/$ itksnap
Return code : 0
itksnap@e240848aeced:/$
```

National Imaging Facility

VERSITY
NSLAND
A
Create change

RUN THE CONTAINER INTERACTIVELY

- Running graphical applications requires some wizardry
- We will keep things simple here by going into the container

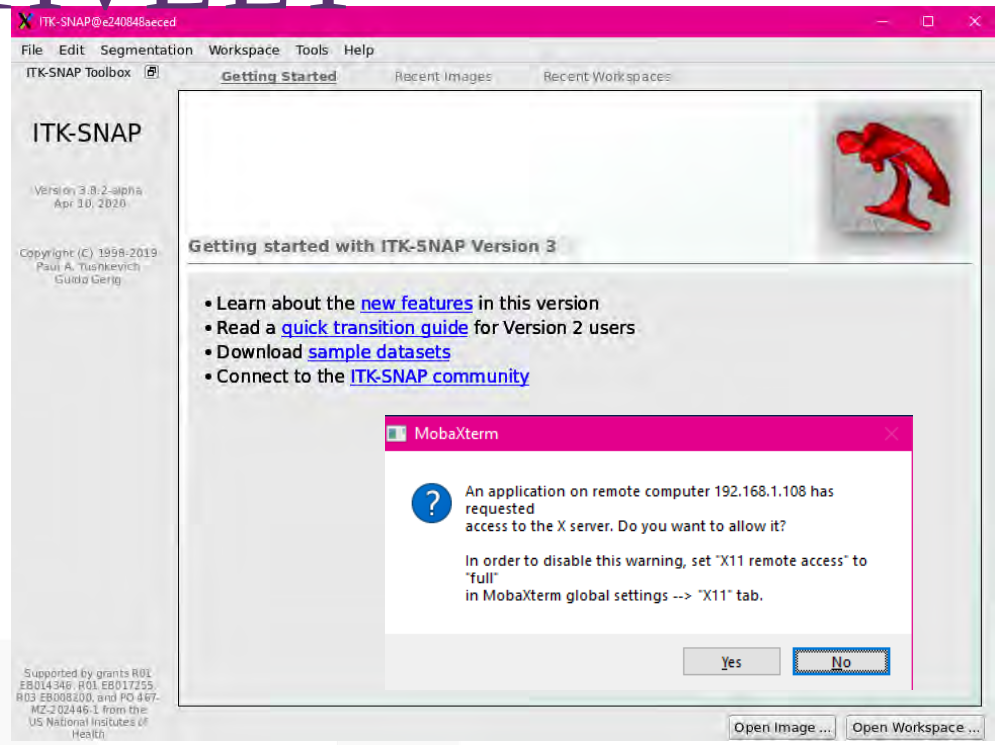

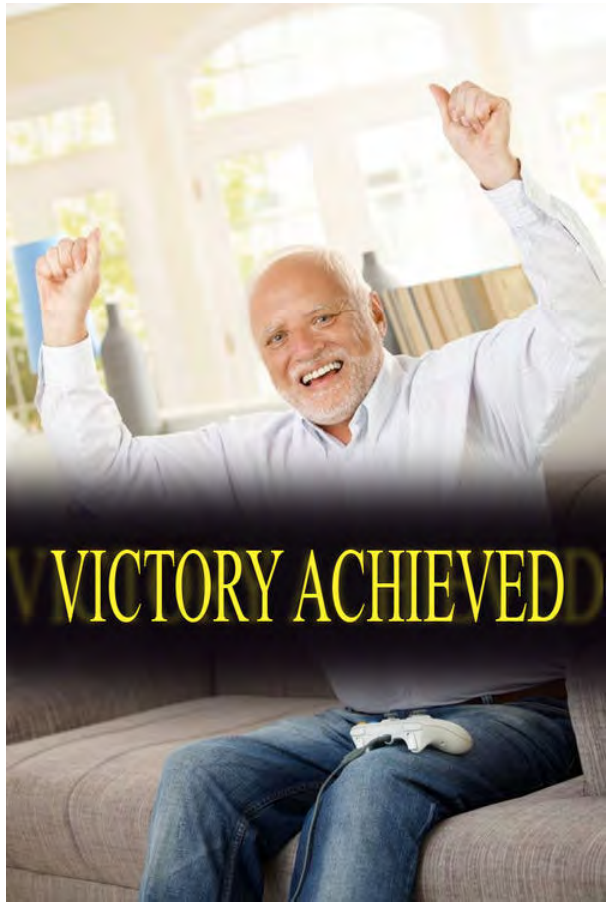


```
PS C:\Users\thoma> docker run -it thomshaw92/itksnap bash
itksnap@e240848aeced:/$ export DISPLAY=192.168.1.108:1.0
itksnap@e240848aeced:/$ itksnap
Return code : 0
itksnap@e240848aeced:/$
```

QUESTIONS?

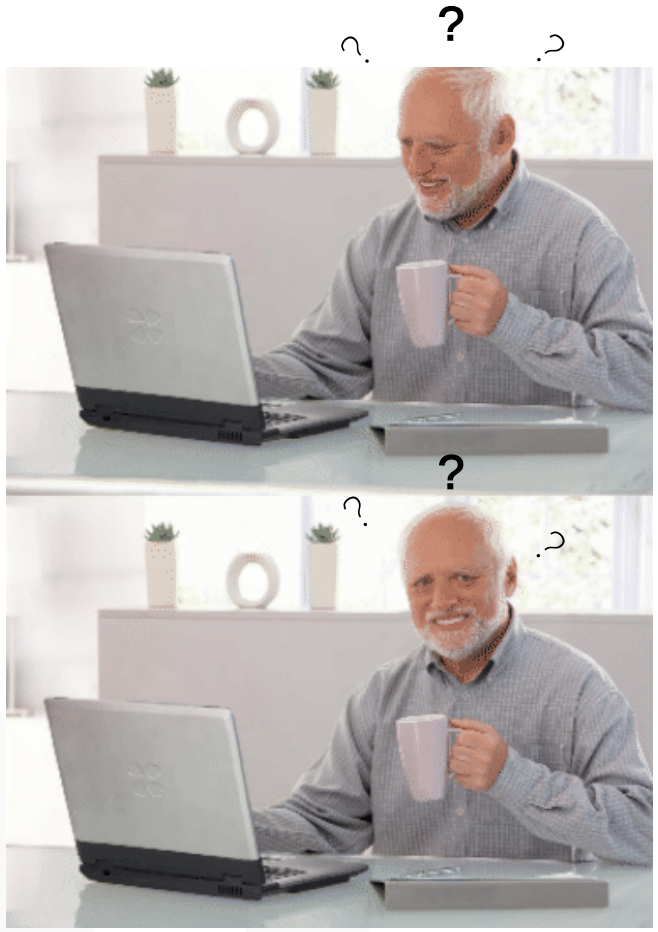

NEURODOCKER

- Neurodocker generates docker+singularity recipes for almost all the tools we need ☺

software	argument	description
AFNI	version*	Either 17.2.02 or latest.
	install_r	If true, install R and AFNI R packages. False by default.
	install_python2	If true, install Python 2.
	install_python3	If true, install Python 3.
ANTs	version*	2.2.0, 2.1.0, 2.0.3, or 2.0.0
	use_binaries	If true (default), use pre-compiled binaries. If false, build from source.
	git_hash	Git hash to checkout to before building from source (only used if use_binaries is false).
Convert3D	version*	"1.0.0" or "nightly".
dcm2niix	version*	"latest", "master", git commit hash, or git tag.
FreeSurfer	version*	Any version for which binaries are provided.
	license_path	Relative path to license file. If provided, this file will be copied into the Docker image. Must be within the build context.
	min	If true, install a version of FreeSurfer minimized for recon-all. See freesurfer/freesurfer#70. False by default.
FSL**	version*	Any version for which binaries are provided.
	eddy_5011	If true, use pre-release version of FSL eddy v5.0.11
	eddy_5011_cuda	6.5, 7.0, 7.5, 8.0; only valid if using eddy pre-release
	use_binaries	If true (default), use pre-compiled binaries. Building from source is not available now but might be added in the future.
	use_installer	If true, use FSL's Python installer. Only valid on CentOS images.
MINC	version*	1.9.15
Miniconda	env_name*	Name of this conda environment.
	yaml_file	Environment specification file. Can be path on host or URL.
	conda_install	Packages to install with conda. e.g., `conda_install="python=3.6 numpy traits"`
	pip_install	Packages to install with pip.
	conda_opts	Command-line options to pass to `conda create`. e.g., `conda_opts="-c vida-nyu"`
	pip_opts	Command-line options to pass to `pip install`.
	activate	If true (default), activate this environment in container entrypoint.
	miniconda_version	Version of Miniconda. Latest by default.
MRtrix3	use_binaries	If true (default), use pre-compiled binaries. If false, build from source.
	git_hash	Git hash to checkout to before building from source (only used if use_binaries is false).
NeuroDebian	os_codename*	Codename of the operating system (e.g., stretch, zesty).
	download_server*	Server to download NeuroDebian packages from. Choose the one closest to you. See `neurodocker generate --help` for the full list of servers.
	pkgs	Packages to download from NeuroDebian.
	full	If true (default), use non-free sources. If false, use libre sources.
PETPVC	version*	1.2.0-b, 1.2.0-a, 1.1.0, 1.0.0
SPM	version*	12 (earlier versions will be supported in the future).
	matlab_version*	R2017a (other MCR versions will be supported once earlier SPM versions are supported).

https://github.com/repronim/neurodocker

National Imaging Facility

The University Of Queensland AUSTRALIA

Create change

NEURODOCKER RECIPE

- Let’s make it easier

```
# Generate singularity recipe
neurodocker generate singularity \
--base ubuntu:16.04 \
--pkg-manager apt \
--itksnap version=3.8.0 \
--env DEPLOY_PATH=/opt/itksnap-3.5.0/bin/ \
--entrypoint /opt/itksnap-3.5.0/bin/itksnap \
--user=neuro \
> Singularity.itksnap
```

SINGULARITY BUILD

- Build singularity container locally
- Or if you don't have root access: https://cloud.sylabs.io/builder
- Build online
- Run!

```
uqtshaw@awoonga1:.../uqtshaw/TrainTrack_containers_2020/3_build_neurodocker> singularity remote login
INFO:    Authenticating with default remote.
Generate an API Key at https://cloud.sylabs.io/auth/tokens, and paste here:
API Key:
INFO:    API Key Verified!
uqtshaw@awoonga1:.../uqtshaw/TrainTrack_containers_2020/3_build_neurodocker> singularity build --remote ${imageName}_${buildDate}.sif ./Singularity.itksnap
Copying blob sha256:e92ed755c008afc1863a616a5ba743b670c09c1698f7328f05591932452a425f
Copying blob sha256:b9fd7cb1ff8f489cf082781b0e1fe0c13b840e20147e8fc8204b4592da7c2f70
Copying blob sha256:ee690f2d57a128744cf4c5b52646ad0ba7a5af113d9d7e0e02b62c06d35fd14c
Copying blob sha256:53e3366ec435596bed2563cc882ba47ec25df6be2b1027e3243e83589c667c1e
Copying config sha256:9387a5fd608d7a23de506446be6fcc9b8c13eed5f40b72106957ebee499dc1ce
Writing manifest to image destination
Storing signatures
```

QUESTIONS

TRANSPARENT SINGULARITY

- **Do I always have to type**

singularity exec itksnap_3.8.0_20200505.sif itksnap

to run itksnap ?

No – we can automatically build wrapper scripts and make our life easier ☺

```
https://github.com/CAIsr/transparent-singularity
```

TRANSPARENT SINGULARITY

```
git clone https://github.com/CAIsr/transparent-singularity.git afni_20.1.06
```

```
cd afni_20.1.06/
```

```
./run_transparent_singularity.sh afni_20.1.06_20200522.sif
```

```
-----------------------------------
installing container afni_20.1.06_20200522.sif
-----------------------------------
-----------------------------------
IMPORTANT: you need to set your system specific mount points in your .bashrc!: e.g. export SINGULARITY_BINDPATH="/opt,/data"
-----------------------------------
checking for singularity ...
deploying in /gpfs1/scratch/30days/uqsbollm/test/ohbm-container-talk/afni_20.1.06
checking if container needs to be downloaded
pulling image now ... this will take some time!
making container executable
checking which executables exist inside container
create singularity executable for each regular executable in commands.txt
creating activate script that runs deactivate first in case it is already there
deactivate script
create module files one directory up
```

TRANSPARENT SINGULARITY

This created a wrapper for every executable inside the container ☺

```
(base) uqsbollm@awoonga1:.../test/ohbm-container-talk/afni_20.1.06> cat suma
#!/usr/bin/env bash
export PWD=`pwd -P`
singularity exec --pwd $PWD /gpfs1/scratch/30days/uqsbollm/test/ohbm-container-talk/afni_20.1.06/afni_20.1.06_20200522.sif suma $@
```

And it made these wrapper scripts known to the HPC module system ☺

```
(base) uqsbollm@awoonga1:.../test/ohbm-container-talk/afni_20.1.06> module avail

------------------------------------------------------------------------------ /
   afni/20.1.02    afni/20.1.06 (D)    ants/2.3.0    ants/2.3.1 (D)    freesurfer/6.0.1
```

TRANSPARENT SINGULARITY

- Now we can just use all tools inside the container (and combine tools from different containers)

```
module load afni/20.1.06
```

```
(base) uqsbollm@awoonga1:.../test/ohbm-container-talk/afni_20.1.06> suma

suma:
     No input specified, loading some toy surfaces...
     Use '.' and ',' to cycle between them.
     See suma -help for assistance.

oo     Warning suma (SUMA_suma.c:1179):

 No sumarc file found. You should create one by running the following:

              suma -update_env

 I also recommend you run 'suma -update_env' whenever you update AFNI.
```

Thank you

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

Centre for Advanced Imaging

@thomcat992

@sbollmann_MRI

National Imaging Facility

Australian Government
Australian Research Council

AUSTRALIAN CANCER RESEARCH FOUNDATION

NCRIS
National Research Infrastructure for Australia
An Australian Government Initiative

Queensland Government