

Open Infrastructure for open science

How Binder Powers an Open Stack in the Cloud

Chris Holdgraf, UC Berkeley and Project Jupyter






andrewosh



betatim




captainsafia



Carreau



choldgraf



consideratio



ellisonbg




fperez



freeman-lab



henhc




JamiesHQ



jhamrick



jzf2101



lheagy




mbmilligan



minrk



mpacer




parente




rgbkrk



Ruv7



ryanlovelt



sgibson91



takluyver




tgeorgeux



willingc



yuvipanda




Zsailer



you???





THE LEONA M. AND HARRY B.
HELMSLEY
CHARITABLE TRUST



ALFRED P. SLOAN
FOUNDATION

GORDON AND BETTY
MOORE
FOUNDATION

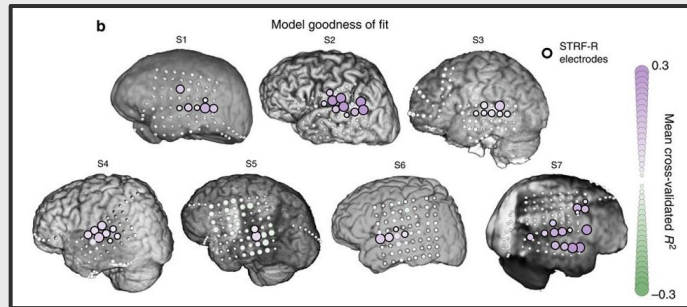


A bit about me then...

Berkeley Neuroscience
HELEN WILLS NEUROSCIENCE INSTITUTE



Cognitive Neuroscience



Open Source




A bit about me now...

Research and Open Source



Education and Open Source



Jupyter @ Berkeley





a *community* of people and an *ecosystem* of open tools and standards for interactive computing





create things that are **language-agnostic** and **modular**.
Empower people to use **other open tools**.



Aside: Jupyter and the last mile problem



You



San Jose
Coffee!



You



San Jose
Coffee!



One option: walk there by myself

Another option: pay somebody to drive me

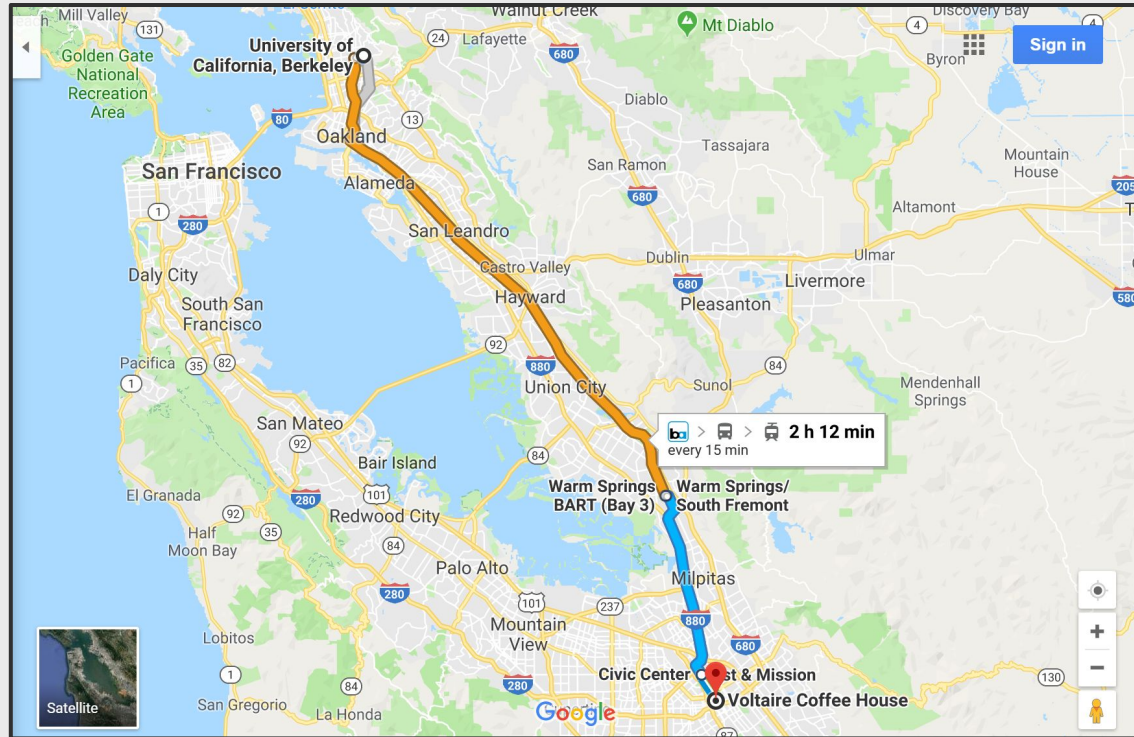
My favorite option: use public infrastructure



You



San Jose
Coffee!



You



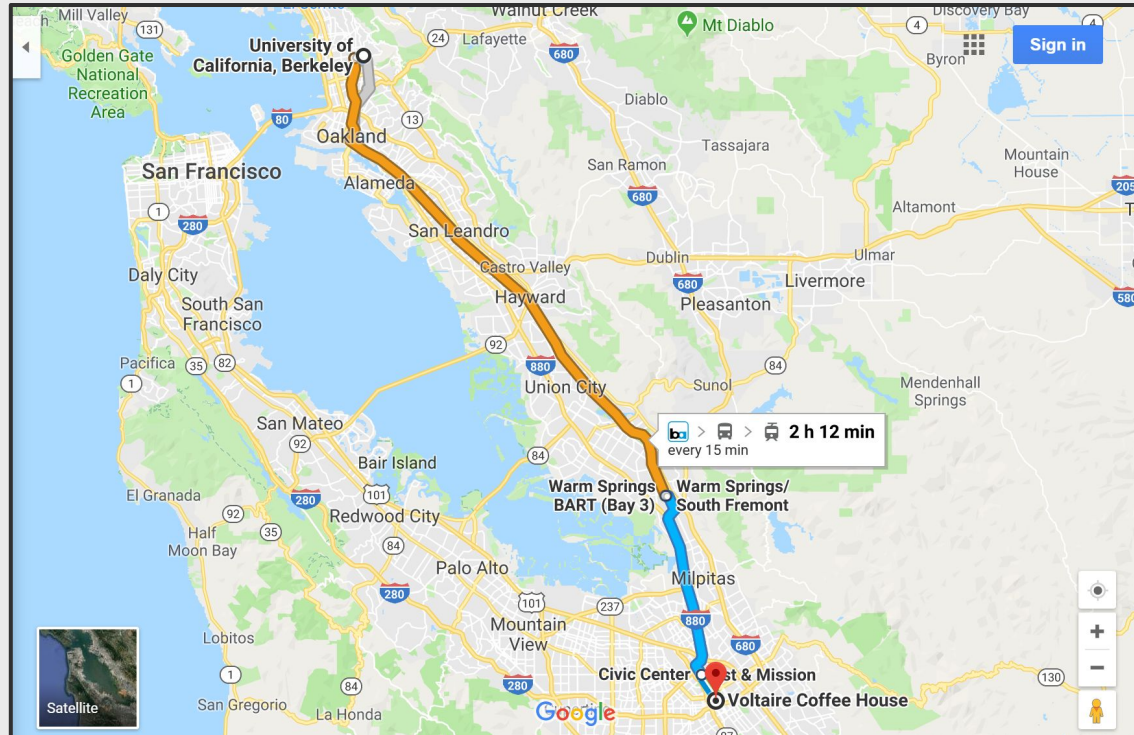
BART



Bus



San Jose
Coffee!



You



BART



Bus



San Jose
Coffee!



Public infrastructure gets us
closer to our goal.

It makes the last mile shorter.



How does **Jupyter**
fit in to this?



You



Your
awesome
report



You



Your
Awesome
report



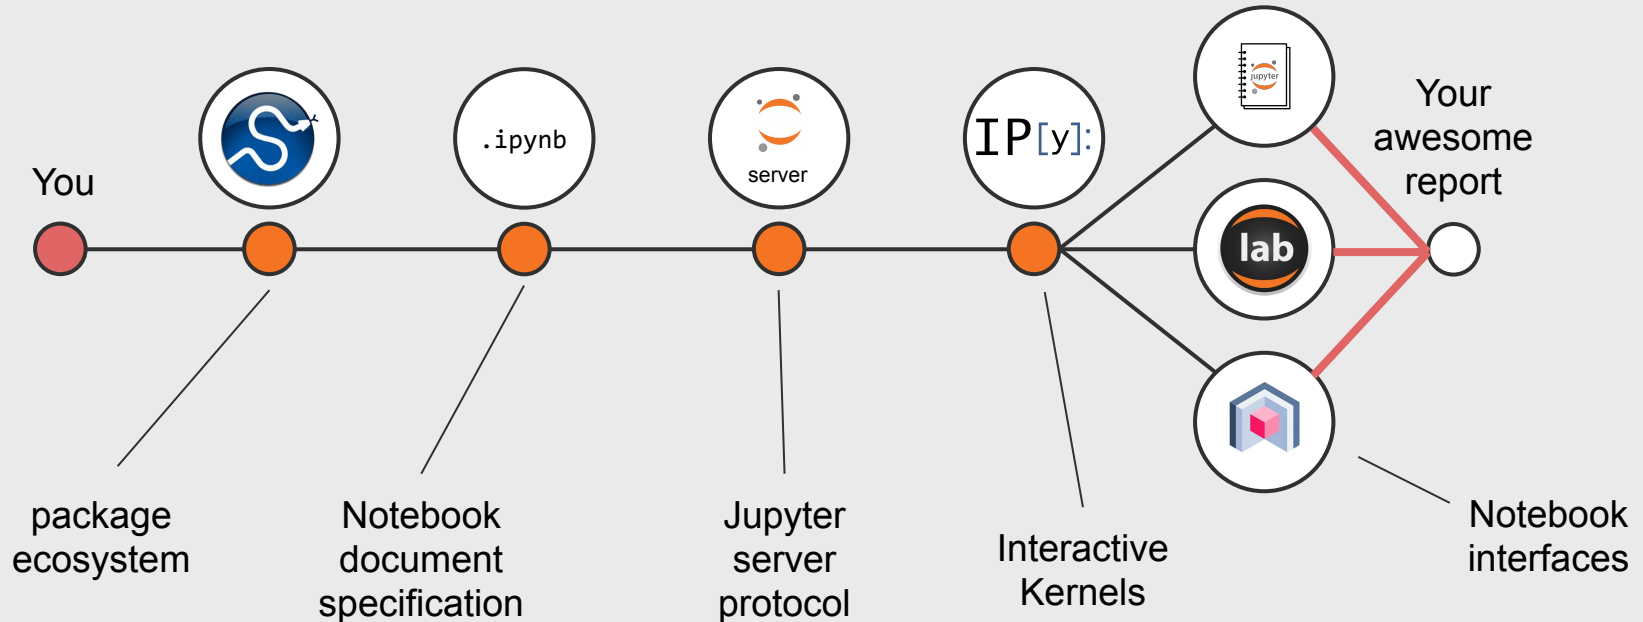
One option: build it from scratch

Another option: pay somebody for a product

My favorite option: build on open, modular tools



Jupyter shortens the last mile by creating and leveraging public infrastructure



Back to our talk...



The science is the code

*An article about computational science in a scientific publication is not the scholarship itself, it is merely advertising of the scholarship. The **actual scholarship** is the complete software development environment and the complete set of instructions which generated the figures.*

Buckheit and Donoho
(paraphrasing John Claerbout)
WaveLab and Reproducible Research, 1995



Our mission for this talk



Part 1: from your laptop
to the cloud with
JupyterHub



(some) data science should
be taught to everyone
(no, really)



Here's what this means at Berkeley...



How can we connect people
with computation?



What is JupyterHub?

Host pre-configured data science environments on shared infrastructure




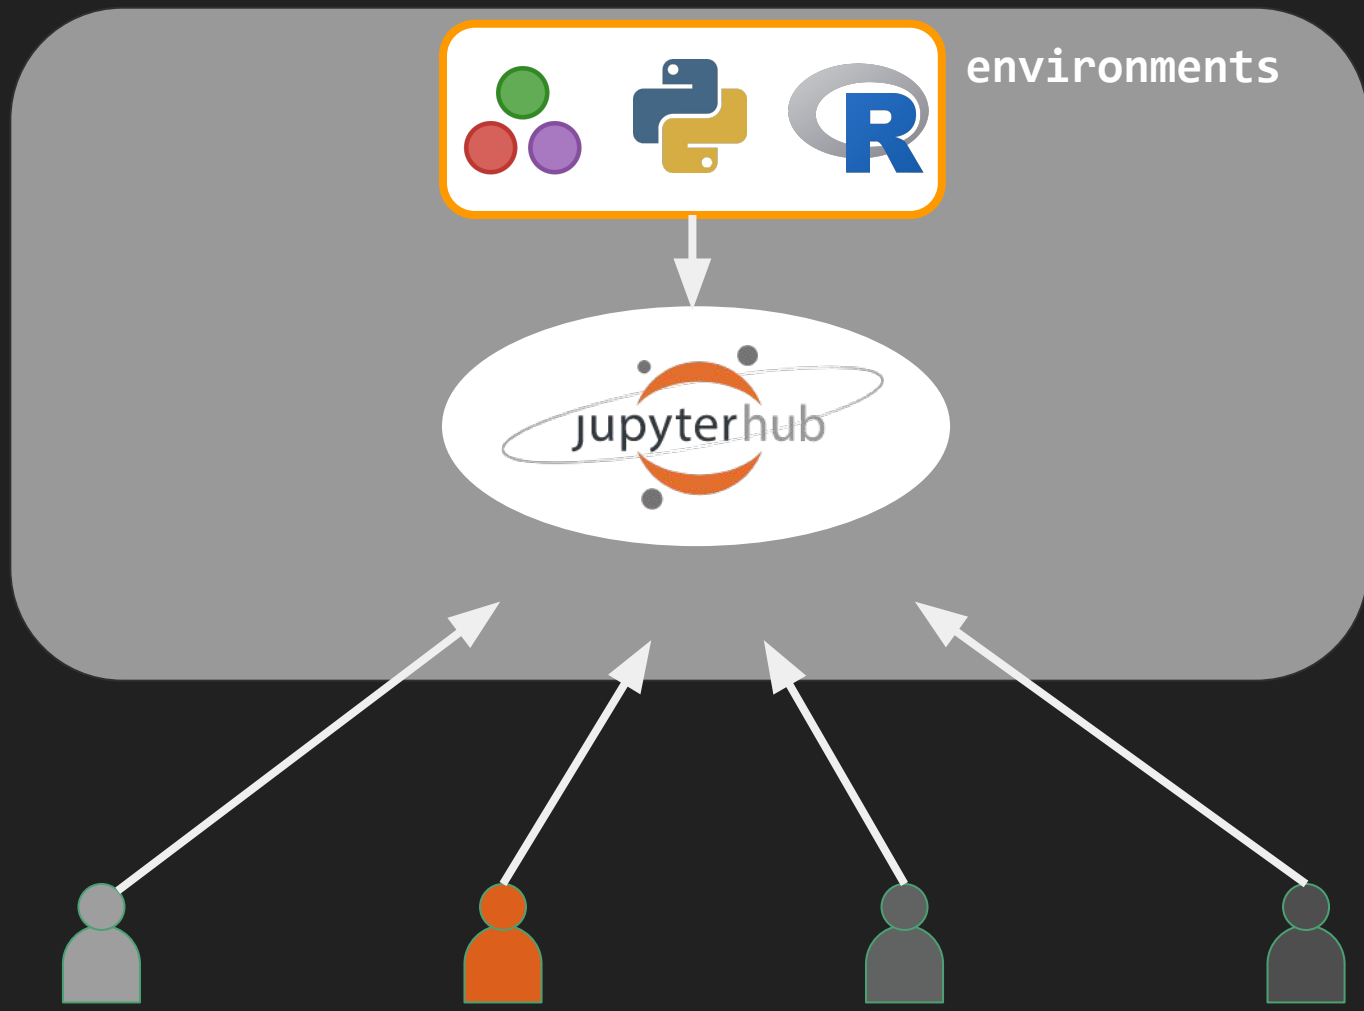
jupyter.org/hub




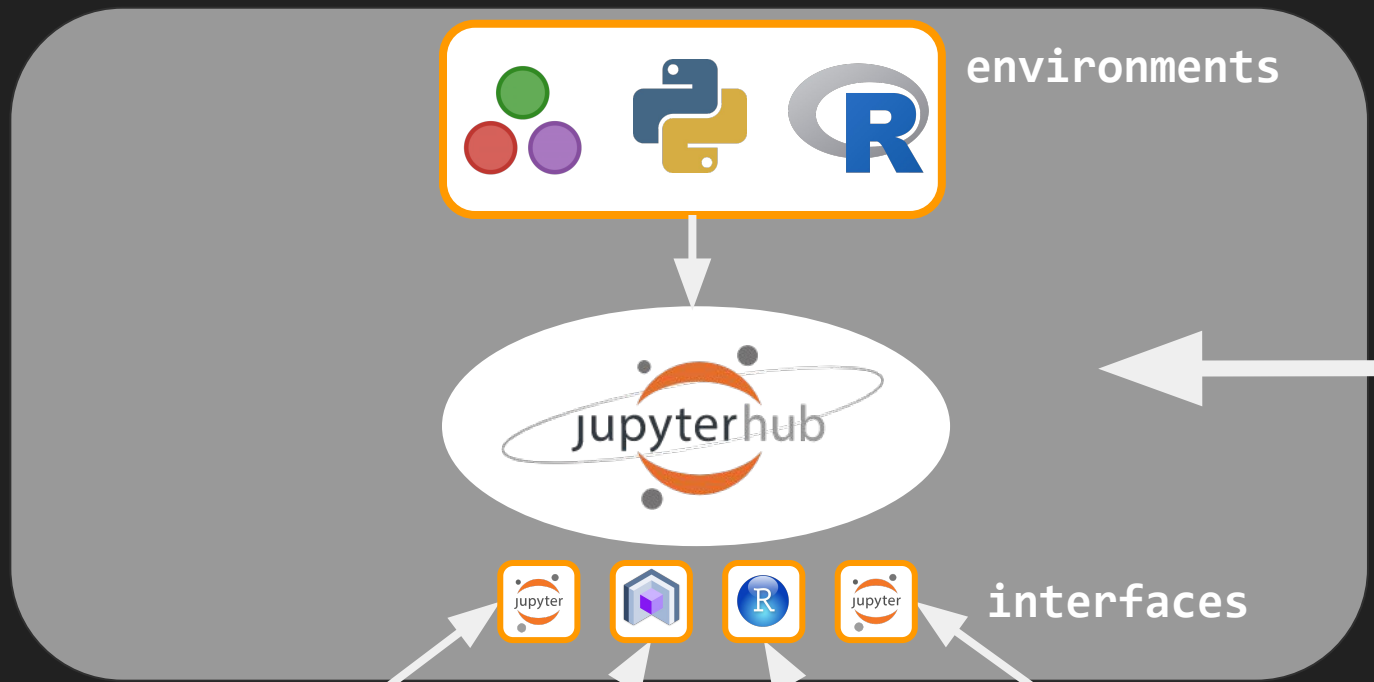
My fancy machine in the cloud









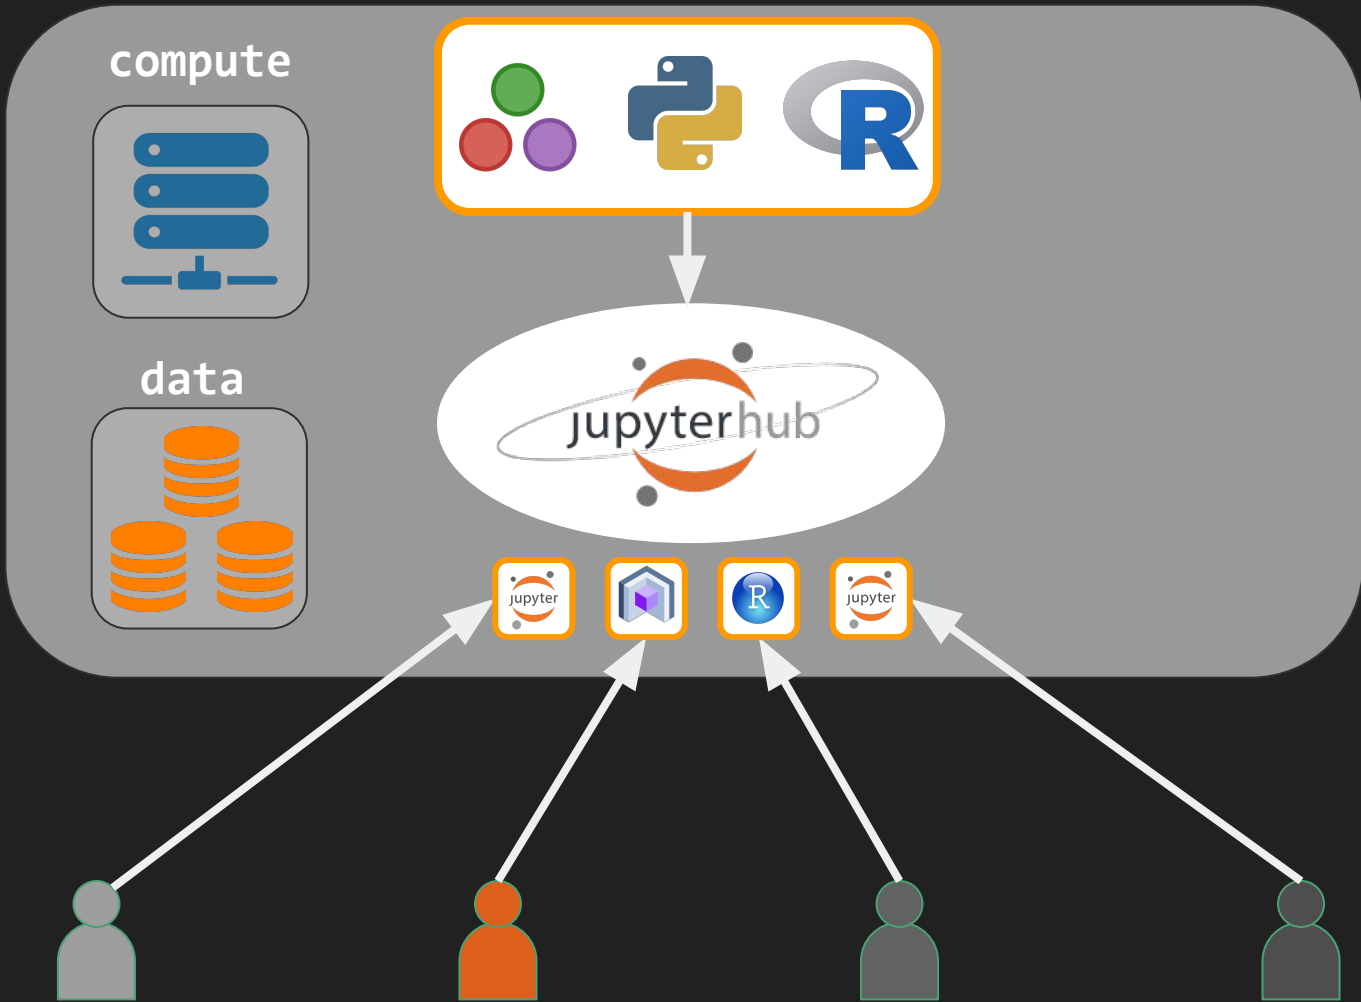


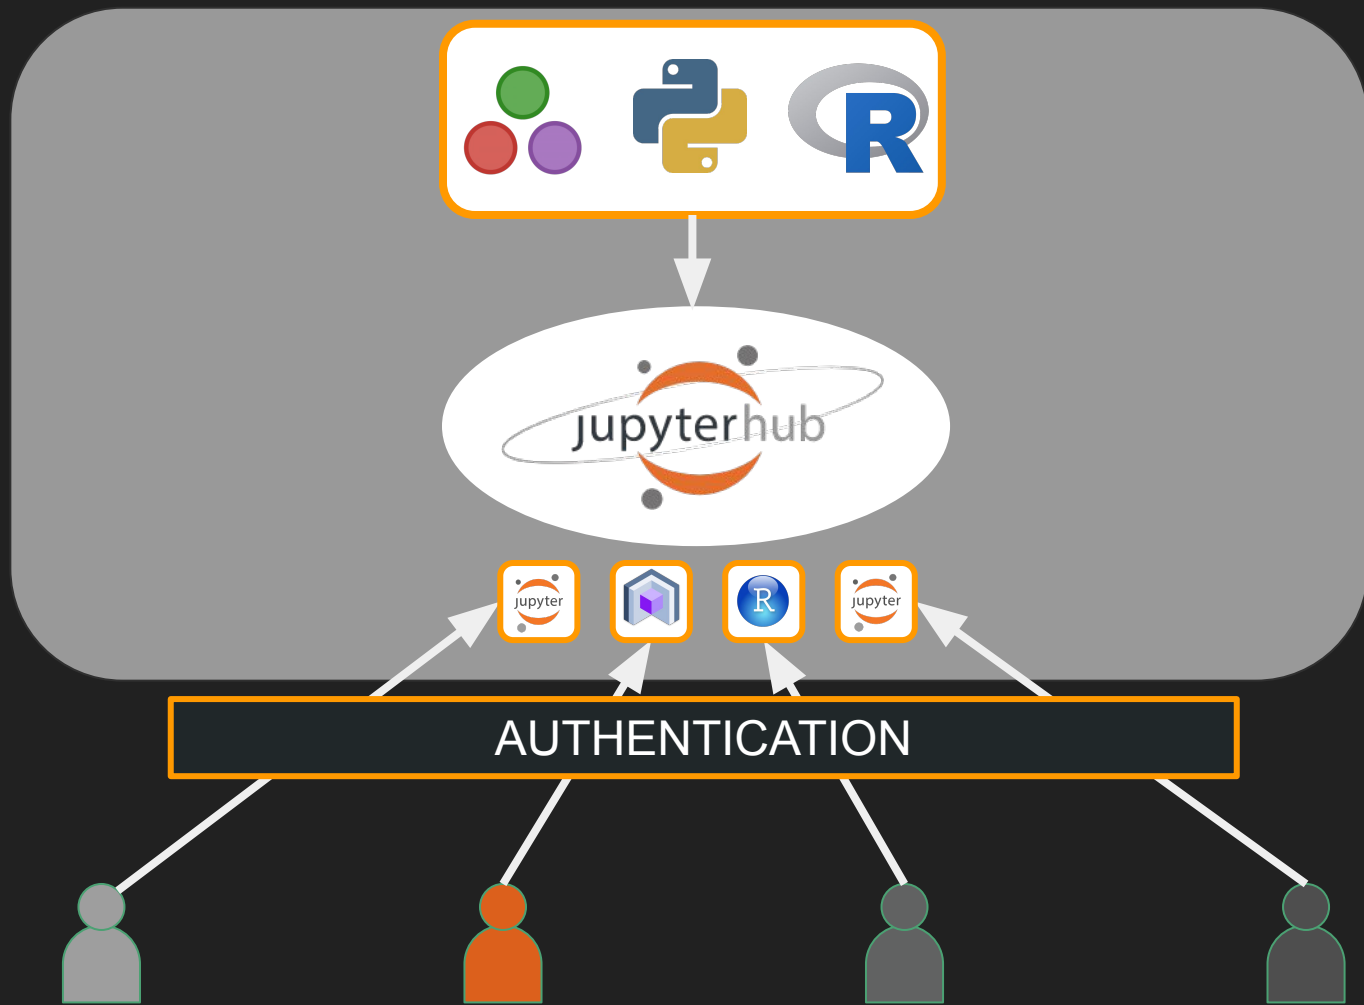
myhub.org



documents







Chris Is Trying A Live Demo

Hopefully he doesn't embarrass himself too badly.

- * Press ESC to cancel and return to Windows.
- * Press ENTER to close this application that is not responding. You will lose any unsaved information in this application.
- * Press CTRL+ALT+DEL again to restart your computer. You will lose any unsaved information in all applications.

[Textbook link](#)
[Interact link](#)





jupyter compare_fitbit_healthkit (autosaved) Logout Control Panel

File Edit View Insert Cell Kernel Widgets Help Not Trusted | Python 3

+

An example of using Open Humans notebooks for activity data analysis

Getting started

In this notebook we want to compare our personal `fitbit` and `Apple HealthKit` data. If you want to run this notebook on your personal data in the [Open Humans Notebooks](#) you need to have both [Fitbit](#) as well as the [HealthKit Integration](#) data in your Open Humans account.


Initialize packages

We'll need some Python packages to get our personal data out of Open Humans and to play around with it.

```
In [1]: %matplotlib inline
import os, requests, json
import datetime
import pandas as pd
import matplotlib
import numpy as np
```

notebooks.openhumans.org





How can users
package+share their work?



Part 2: packaging and sharing your environment with **repo2docker**



What is repo2docker?

*Convert a repository into a Docker image
that runs the code inside.*



repo2docker.readthedocs.io



github.com/minrk/ligo-binder





repo2docker
*what does it
do?*

```
$ jupyter repo2docker \  
>      https://github.com/minrk/ligo-binder
```



repo2docker
*what does it
do?*

```
$ jupyter repo2docker \  
>      https://github.com/minrk/ligo-binder
```

```
Cloning into  
'/var/folders/.../T/repo2dockermu6z66sd' ...  
Using CondaBuildPack builder  
Step 1/31 : FROM buildpack-deps:bionic  
----> 29f4eef41002  
Step 2/31 : ENV DEBIAN_FRONTEND=noninteractive  
----> Using cache  
----> ee1ba7c4f5f4  
Step 3/31 : RUN apt-get update &&      apt-get  
install --yes --no-install-recommends locales &&  
apt-get purge &&      apt-get clean &&      rm -rf  
/var/lib/apt/lists/*
```



Step 1: get the **repo**





repo2docker
*what does it
do?*

```
git clone https://github.com/me/myproject
```



Step 2: Identify **requirements**

repo2docker
*what does it
do?*

choldgraf Update README.md	
 README.md	Update README.md
 index.ipynb	first move
 requirements.txt 	Pin requirements.txt to values that actually exist



Step 2: Identify requirements

repo2docker
*what does it
do?*

yuvipanda s/Anaconda/conda/ ...	
README.md	s/Anaconda/conda/
environment.yml ←	Require a stable install of dask
index.ipynb	updating to new syntax



Step 2: Identify requirements

repo2docker
*what does it
do?*

betatim Update README.md	
📁 bus-dashboard	Remove the DESCRIPTION file
📄 README.md	Update README.md
📄 index.ipynb	adding example
📄 install.R	Add example Shiny app
📄 runtime.txt	Remove dockerfile, add runtime.txt



Step 3: generate **Dockerfile**

setup the runtime environment

...

```
COPY conda/install-miniconda.bash /tmp/install-miniconda.bash
```

```
COPY conda/environment.py-3.6.frozen.yml /tmp/environment.yml
```

```
RUN bash /tmp/install-miniconda.bash && \  
rm /tmp/install-miniconda.bash /tmp/environment.yml
```

...

repo2docker
*what does it
do?*



Step 3: generate **Dockerfile**

assemble the environment for the repo

...

Copy and chown stuff.

```
COPY src/ ${HOME}
```

```
RUN chown -R ${NB_USER}:${NB_USER} ${HOME}
```

*# Run assemble scripts! These will actually build the spec
in the repository into the image.*

```
USER ${NB_USER}
```

```
RUN ${KERNEL_PYTHON_PREFIX}/bin/pip install --no-cache-dir \  
-r "requirements.txt"
```

...

repo2docker
what does it
do?



Step 4: **build** (& **push**) image

repo2docker
*what does it
do?*

```
docker build -t myimage  
docker push myimage
```



repo2docker
*what does it
do?*

```
$ jupyter repo2docker \  
>      https://github.com/minrk/ligo-binder
```

```
Cloning into  
'/var/folders/.../T/repo2dockermu6z66sd' ...  
Using CondaBuildPack builder  
Step 1/31 : FROM buildpack-deps:bionic  
---> 29f4eef41002  
Step 2/31 : ENV DEBIAN_FRONTEND=noninteractive  
---> Using cache  
---> ee1ba7c4f5f4  
Step 3/31 : RUN apt-get update &&      apt-get  
install --yes --no-install-recommends locales &&  
apt-get purge &&      apt-get clean &&      rm -rf  
/var/lib/apt/lists/*
```



Some supported configuration files

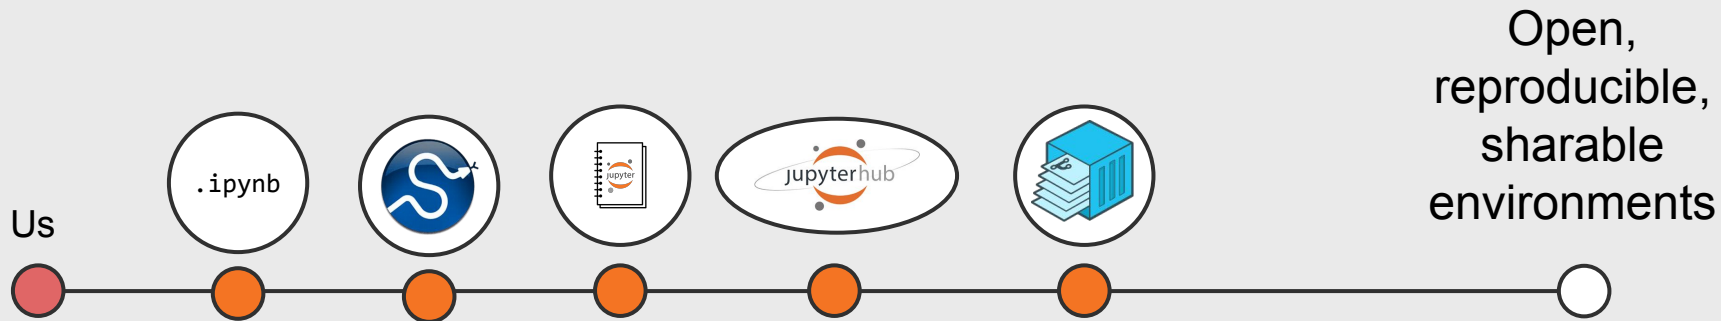
- `environment.yml` - conda
- `requirements.txt` - pip
- `REQUIRE` - julia
- `install.R` - R
- `apt.txt` - apt
- `setup.py` - pip
- `postBuild` - shell
- `runtime.txt`
- `Dockerfile` - docker
- **`yournewbuildpack.txt`**



Guiding principles of repo2docker

- Repos should be **human** and **machine readable**
- Use **existing specifications and standards**
- Support **many languages** and interfaces
- Be lightweight and **tightly-scoped**, but **extendable**





Part 3: tying this together with **BinderHub**



What is BinderHub?


One-click sharable, interactive, reproducible environments from your public git repository



mybinder.org

binderhub.readthedocs.io





=



+



BinderHub is open tech...

- Built on Kubernetes
- Cloud-agnostic
- Scalable
- **Community driven and deployable by anyone**



One example: mybinder.org



Turn a GitHub repo into a collection of
interactive notebooks

Have a repository full of Jupyter notebooks? With Binder, open those notebooks in an executable environment, making your code immediately reproducible by anyone, anywhere.

Build and launch a repository

GitHub repo or URL

Git branch, tag, or commit

Path to a notebook file (optional)

File ▾

launch

Copy the URL below and share your Binder with others:

Copy the text below, then paste into your README to show a binder badge: `launch binder`



Chris Is Trying A Live Demo

Hopefully he doesn't embarrass himself too badly.

- * Press ESC to cancel and return to Windows.
- * Press ENTER to close this application that is not responding. You will lose any unsaved information in this application.
- * Press CTRL+ALT+DEL again to restart your computer. You will lose any unsaved information in all applications.

mybinder.org

binder-examples/requirements





Project Jupyter [Follow](#)

Project Jupyter exists to develop open-source software, open standards, and services for interactive and reproducible computing.

Feb 20 · 5 min read

JupyterLab is Ready for Users

We are proud to announce the beta release series of JupyterLab, the next-generation web-based interface for [Project Jupyter](#).

tl;dr: JupyterLab is ready for daily use ([installation](#), [documentation](#), [try it with Binder](#)) *



Mybinder.org is an open service...

Mybinder.org activity and status:

grafana.mybinder.org/

Our billing:

github.com/jupyterhub/binder-billing

Our operations guide:


mybinder-sre.readthedocs.io/

Our fires:

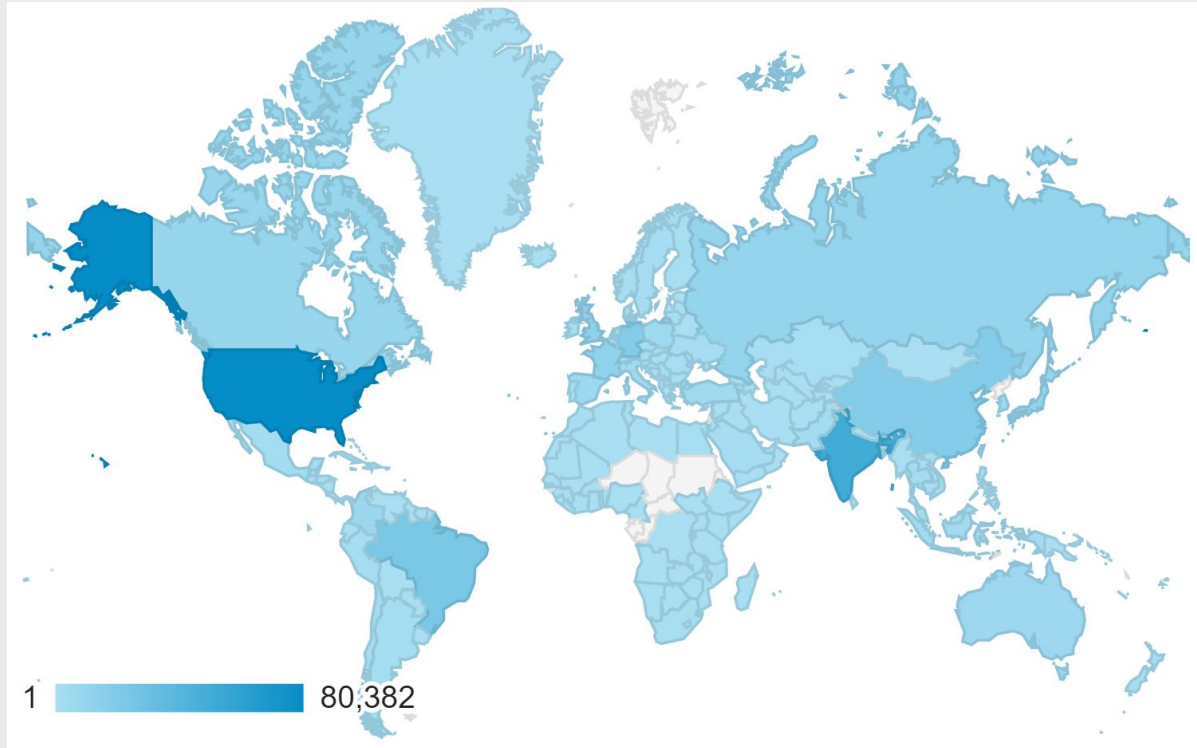
mybinder-sre.readthedocs.io/en/latest/#incident-reports



mybinder.org weekly sessions, last ~year



mybinder.org sessions, last month





BinderHubs around the world



Some cool new projects
(or, stuff you can help out with)



Interactive books with

github.com/jupyter/jupyter-book

Jupyter-Book



Jupyter Book

Home

Search

Getting started

Create your book

Build and publish your book

FAQ

How-to and advanced topics

Features and customization

Markdown files

← TOGGLE SIDEBAR

Download

Thebelab

Interact

Interactive code in your book

Sometimes you'd rather let people interact with code *directly on the page* instead of sending them off to a Binder or a JupyterHub. There are currently a few ways to make this happen in Jupyter Book (both of which are experimental).

This page describes how to bring interactivity to your book. Both of these tools use [MyBinder](#) to provide a remote kernel.

Making your page inputs interactive

🌟 experimental 🌟

If you'd like to provide interactivity for your content without making your readers leave the Jupyter Book site, you can use a project called [Thebelab](#).

This provides you a button that, when clicked, will convert each code cell into an **interactive** cell.

☰ ON THIS PAGE

[MAKING YOUR PAGE INPUTS INTERACTIVE](#)

[USING INTERACTIVE WIDGETS ON YOUR PAGE](#)

HTML dashboards with voila

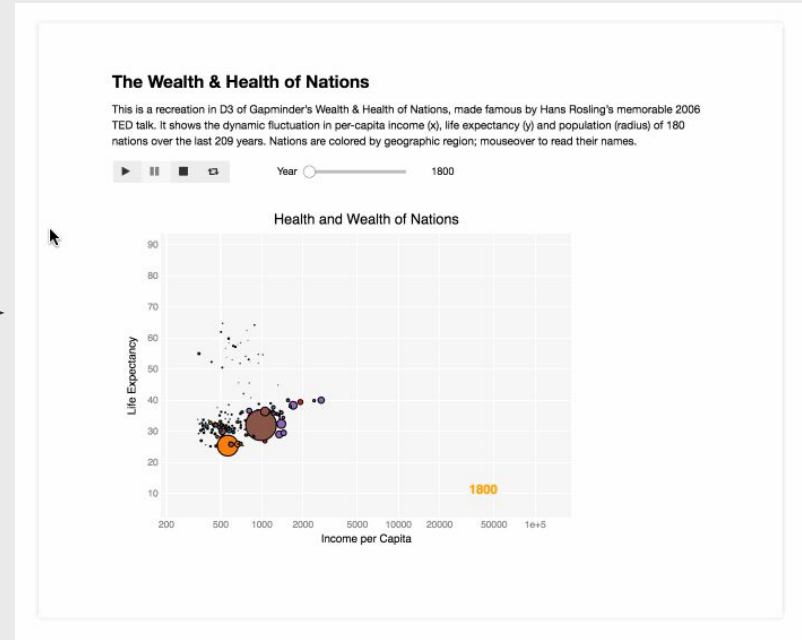
github.com/QuantStack/voila

```
jupyter  
nbviewer
```

voila / notebooks

```
In [ ]: import warnings  
warnings.filterwarnings('ignore')
```

```
In [ ]: import numpy as np  
from bqplot import pyplot as plt  
  
plt.figure(1, title='Line Chart')  
np.random.seed(0)  
n = 200  
x = np.linspace(0.0, 10.0, n)  
y = np.cumsum(np.random.randn(n))  
plt.plot(x, y)  
plt.show()
```



In summary

Jupyter makes the “last-mile” problem as small as possible by building modular, open tools.




jupyter.org




In summary

JupyterHub lets you create a shared, interactive analytics environment




In summary

repo2docker creates reproducible Docker images from a repository



In summary

BinderHub is an open web application to create shareable, reproducible coding environments



Mybinder.org



Get involved with Jupyter

- All of these projects are open source, run by open communities
- Jupyter is a place where **anybody** can participate
- If you'd like to get involved:

jupyterhub-team-compass.readthedocs.io
discourse.jupyter.org

