

Title: Markov Chain Monte Carlo

Speakers: Dustin Lang

Date: September 17, 2019 - 1:00 PM

URL: <http://pirsa.org/19090120>

Markov Chain Monte Carlo

Dustin Lang

2019-09-17

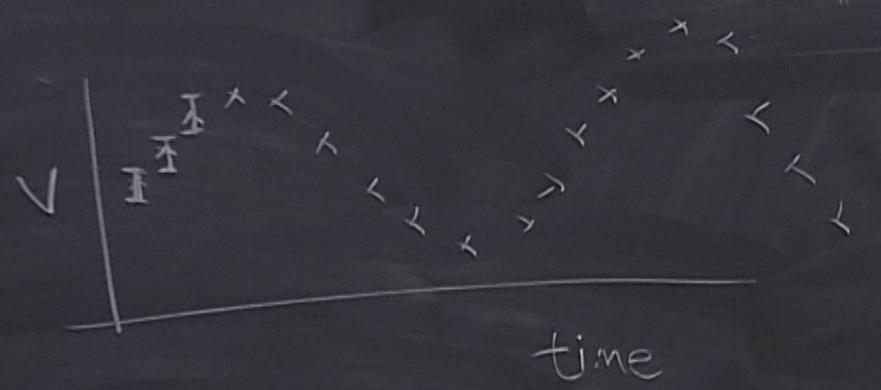
1

Markov Chain Monte Carlo for data analysis

- ▶ MCMC generates samples from a probability distribution function
- ▶ For data analysis: you have some *measurements* X , with uncertainties
- ▶ have a *model* for the thing you're measuring
- ▶ with unknown *parameters* θ
- ▶ the model is *generative* – predicts measurements given parameters
- ▶ MCMC can put constraints on the parameters given your measurements

Markov Chain Monte Carlo for data analysis

- ▶ can think of it as: *model* is a function mapping parameters to observables t : $m(\theta) \rightarrow t$
- ▶ and *measurement process* gives probability of observing X given t : $P(X|t)$
- ▶ (eg, maybe $P(X|t)$ is a Gaussian with zero mean and standard deviation σ)
- ▶ Then, can chain $P(X|t) = P(X|m(\theta))$
- ▶ Usually, we fold m and P together to write $P(X|\theta)$



exp X

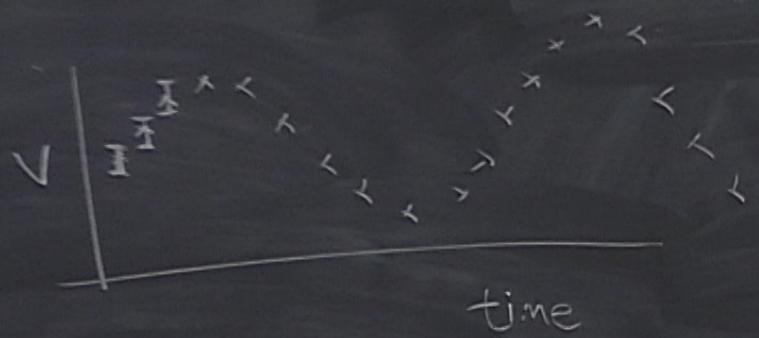
$$P(x|t) = \text{Gaussian}(\text{mean}=0, \sigma=0.1V)$$



exp X

$$P(x|t) = \text{Gaussian}(\text{mean}=0, \sigma=0.1 \text{ V})$$

$$m(\theta) = A \sin(\theta)$$



~~six~~ X

$$p(x|t) = \text{Gaussian}(\text{mean}=0, \sigma=0.1V)$$

$$m(\theta) = A \sin(\theta)$$

$$P(\{x_i\} | \theta) \propto \prod_i \exp\left(-\frac{(x_i - A \sin(\theta_i))^2}{2(0.1V)^2}\right)$$

is conserved.

Markov Chain Monte Carlo for data analysis

- ▶ $P(X|\theta)$ (the “likelihood”) is not what you want!
- ▶ You want $P(\theta|X)$ – the “posterior” – constraints on / measurements of the parameters
- ▶ Good thing you’re Bayesian!
- ▶ “posterior” $P(\theta|X) = \frac{P(X|\theta)P(\theta)}{P(X)}$
- ▶ $P(\theta)$ is the “prior”
- ▶ $P(X)$ is the “evidence”, and in this setting is constant (ignored)

Markov Chain Monte Carlo for data analysis

- ▶ We'll use MCMC to draw samples from the *posterior* probability distribution
- ▶ $P(\theta|X) = P(X|\theta)P(\theta)$
- ▶ The samples are a set of θ values drawn from your parameter constraints distribution

What's in a name

- ▶ *Monte Carlo* – a famous casino in Monaco – alluding to the use of *randomness* in the algorithm
- ▶ *Chain* – a list (here, a list of samples / steps $\theta_1, \theta_2, \theta_3$)
- ▶ *Markov Chain* – a list of “steps” where the decision to step from θ_i to θ_{i+1} depends only on the value of θ_i
- ▶ *Metropolis–Hastings* – first authors of 1953 and 1970 papers (resp) with the simplest, most commonly taught algorithm

Using MCMC

- ▶ “Just write down the probability distributions”
- ▶ hah, the “model to predictions” part, $m(\theta) \rightarrow t$, can be *very* difficult
- ▶ eg, in cosmology, the software *CAMB* does this, in 25,000 lines of FORTRAN
- ▶ writing down the *likelihood* might take some work too!

MCMC: the idea

- ▶ move a “particle” or “walker” around in the parameter space θ , based on the value of the probability function (for us: the posterior probability $P(\theta|X)$)
- ▶ choose new parameters according to a *proposal distribution* – often Gaussian: $\theta_{new} = \theta + \mathcal{N}(0, \Sigma)$
- ▶ compute probability at new θ_{new} and decide whether to move (“jump”) to new parameters
- ▶ *always* keep improvements, *sometimes* keep decreases
- ▶ in the long run, the set of walker positions $\{\theta_1, \theta_2, \theta_3, \dots\}$ are your samples

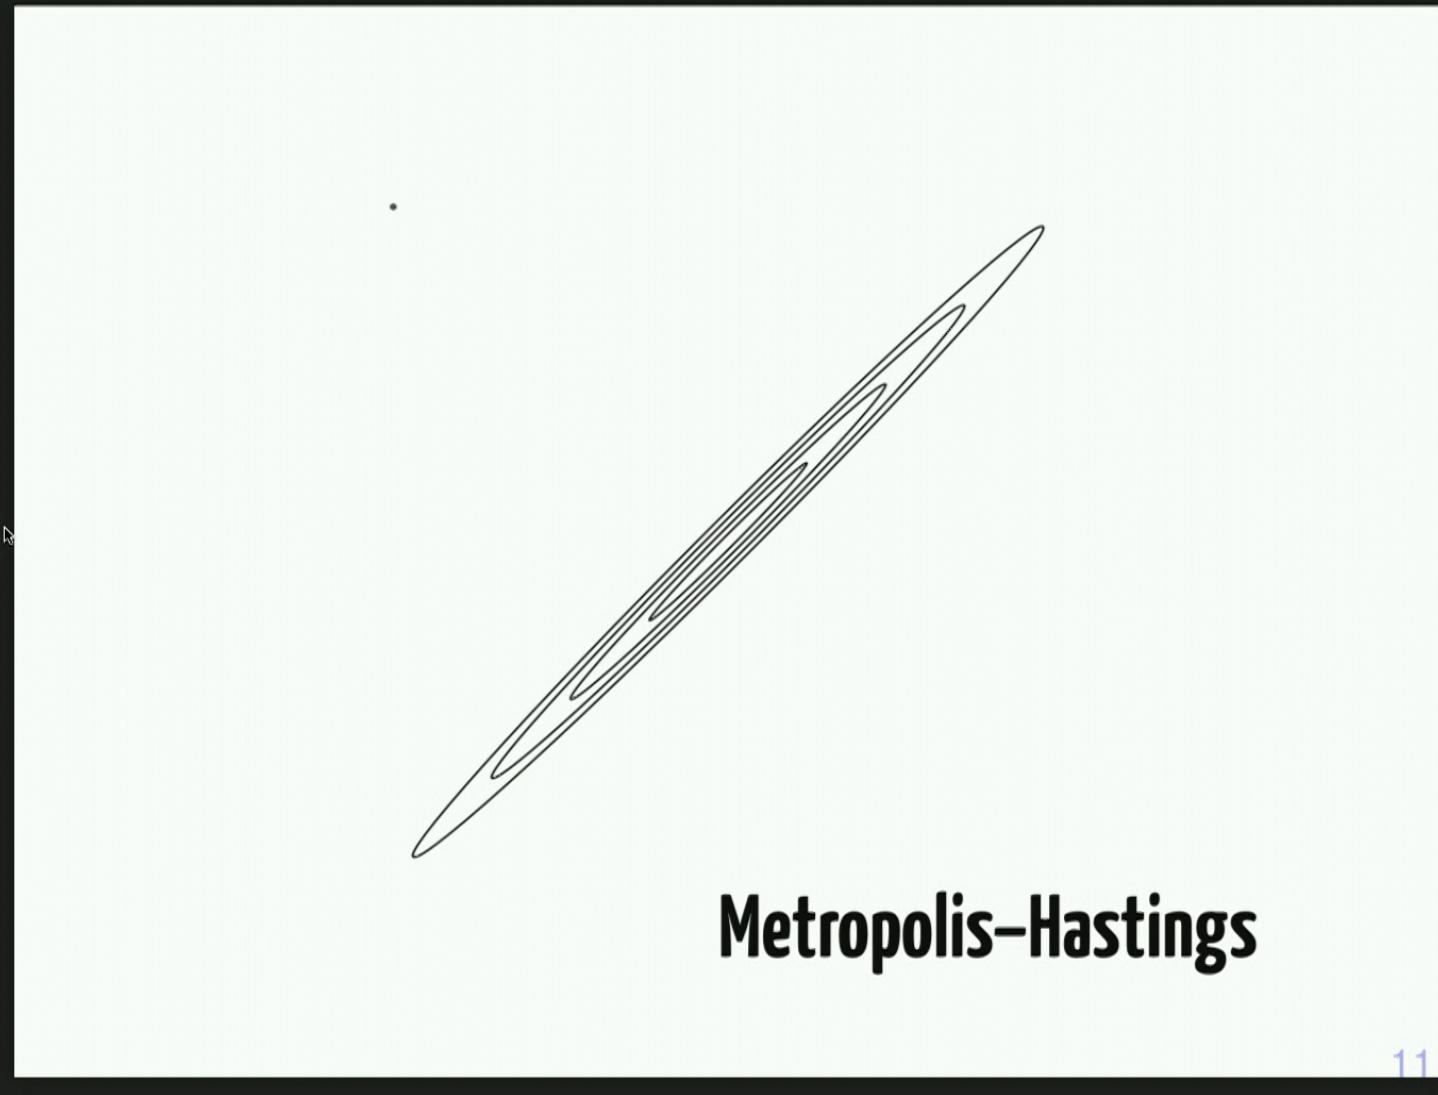
Finally, M–H MCMC code

```
fun MCMC(prob_function, sample_proposal, params, Nsteps):
    chain = []
    prob = prob_function(params)
    for i in 1 to Nsteps:
        # sample from proposal distribution
        params_new = sample_proposal(params)
        prob_new = prob_function(params_new)
        # accept?
        if ((prob_new > prob) OR
            (prob_new / prob) > uniform_random()):
            params = params_new
            prob = prob_new
        chain.append(params)
    return chain
```

Demo

Credit: Dan Foreman-Mackey

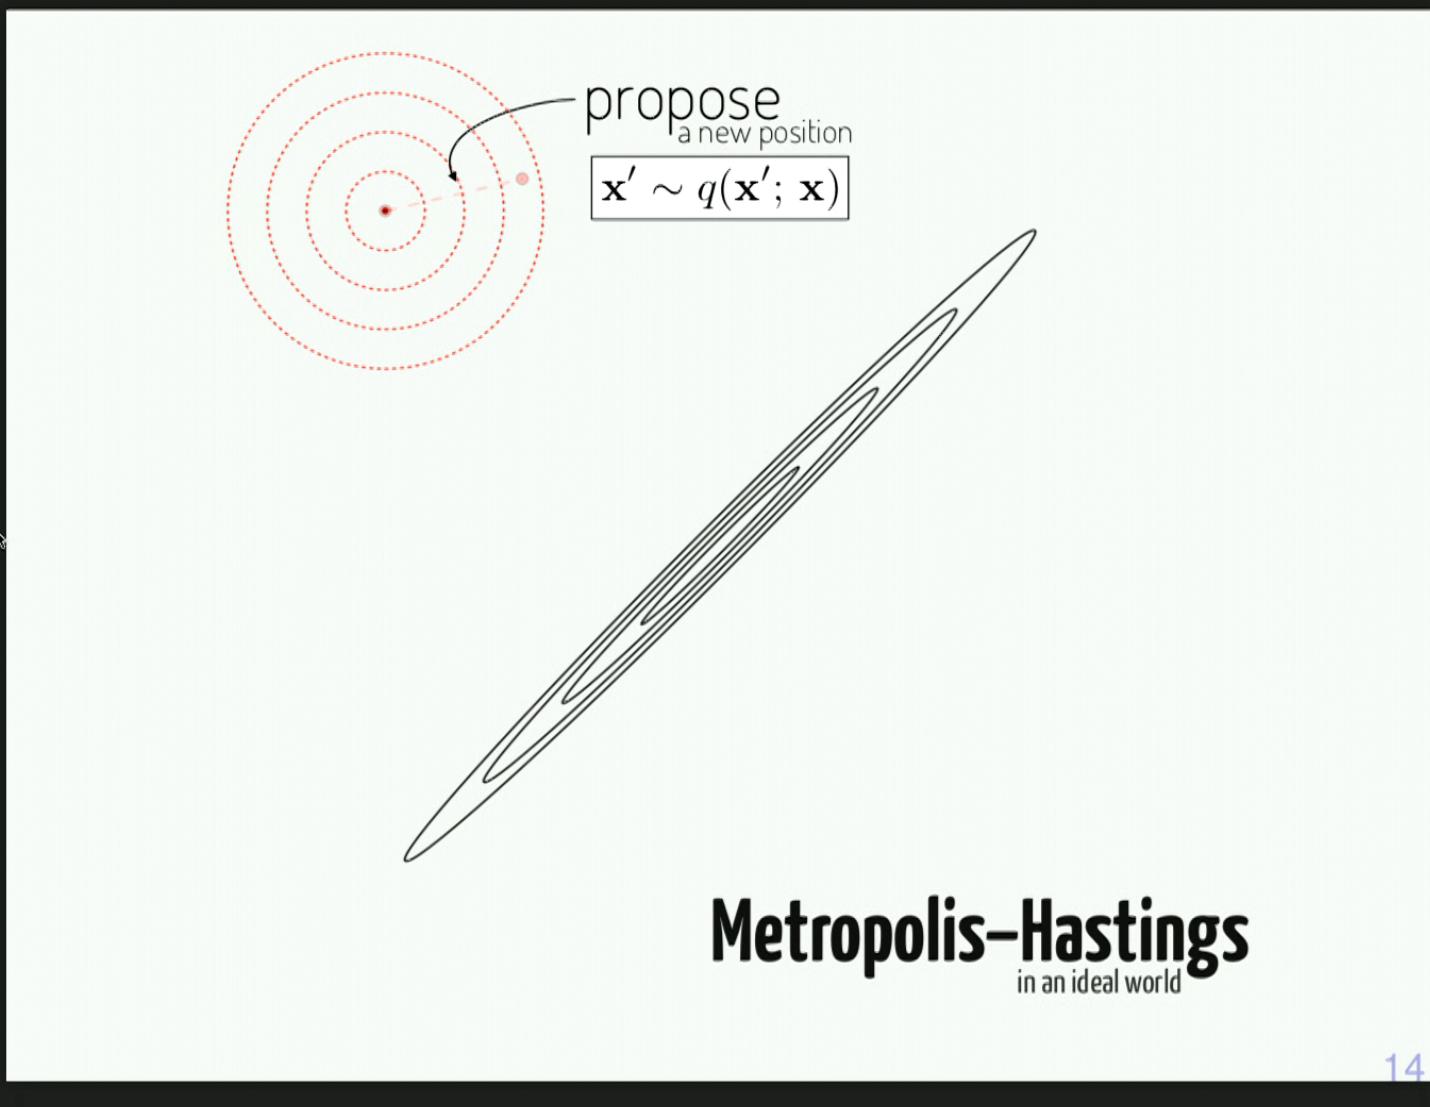
<https://speakerdeck.com/dfm/data-analysis-with-mcmc>



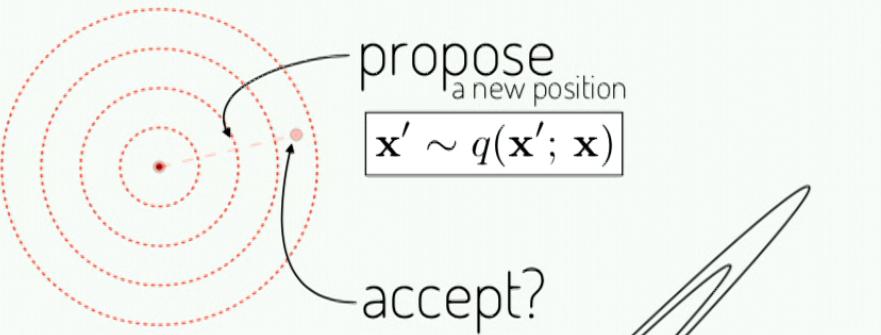
$\leftarrow X$

$$P(x|t) = \text{Gaussian}(\text{mean}=0, \sigma=0.1v)$$
$$m(\theta) = \underbrace{A}_{\{A\}} \sin(\tau - \phi)$$
$$P(\{x_i\} | \theta) \propto \prod_i \exp\left(-\frac{(x_i - A \sin(\tau_i))^2}{2(0.1v)^2}\right)$$

GERVED.



14

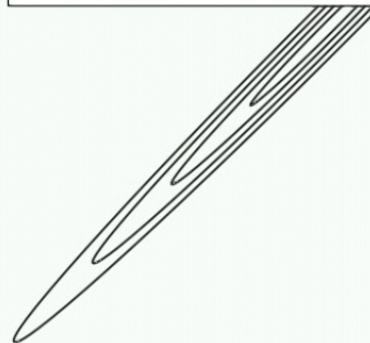


propose
a new position

$$\mathbf{x}' \sim q(\mathbf{x}'; \mathbf{x})$$

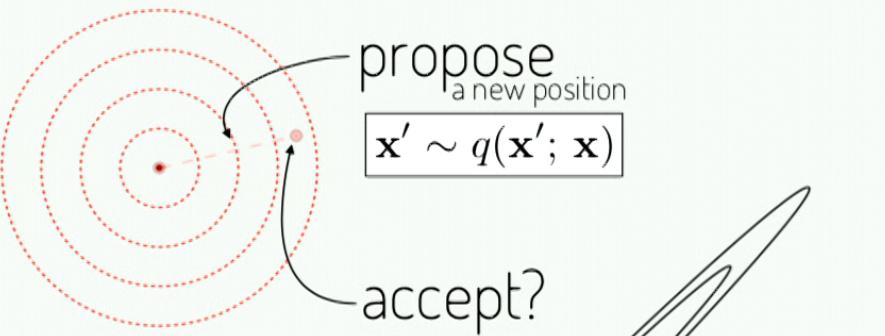
accept?

$$p(\text{accept}) = \min \left(1, \frac{p(\mathbf{x})}{p(\mathbf{x}')}, \frac{q(\mathbf{x}; \mathbf{x}')}{q(\mathbf{x}'; \mathbf{x})} \right)$$



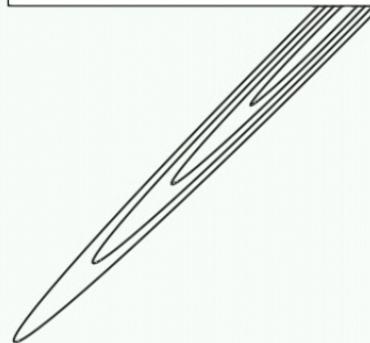
Metropolis-Hastings

in an ideal world



$$p(\text{accept}) = \min \left(1, \frac{p(\mathbf{x})}{p(\mathbf{x}')}, \frac{q(\mathbf{x}; \mathbf{x}')}{q(\mathbf{x}'; \mathbf{x})} \right)$$

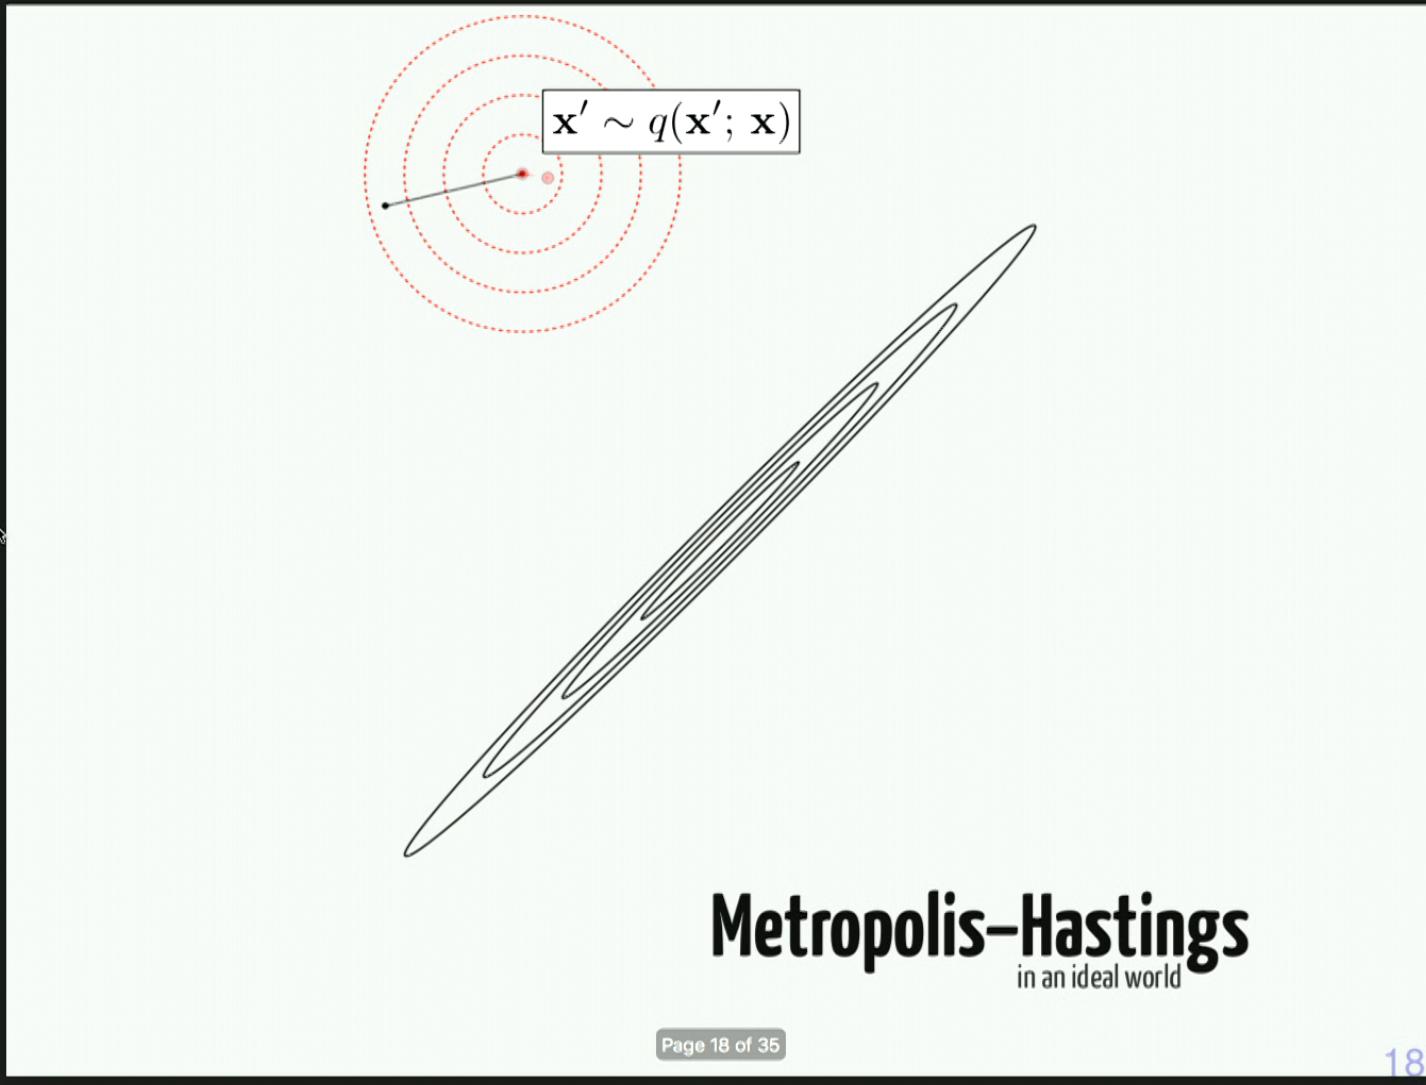
definitely.

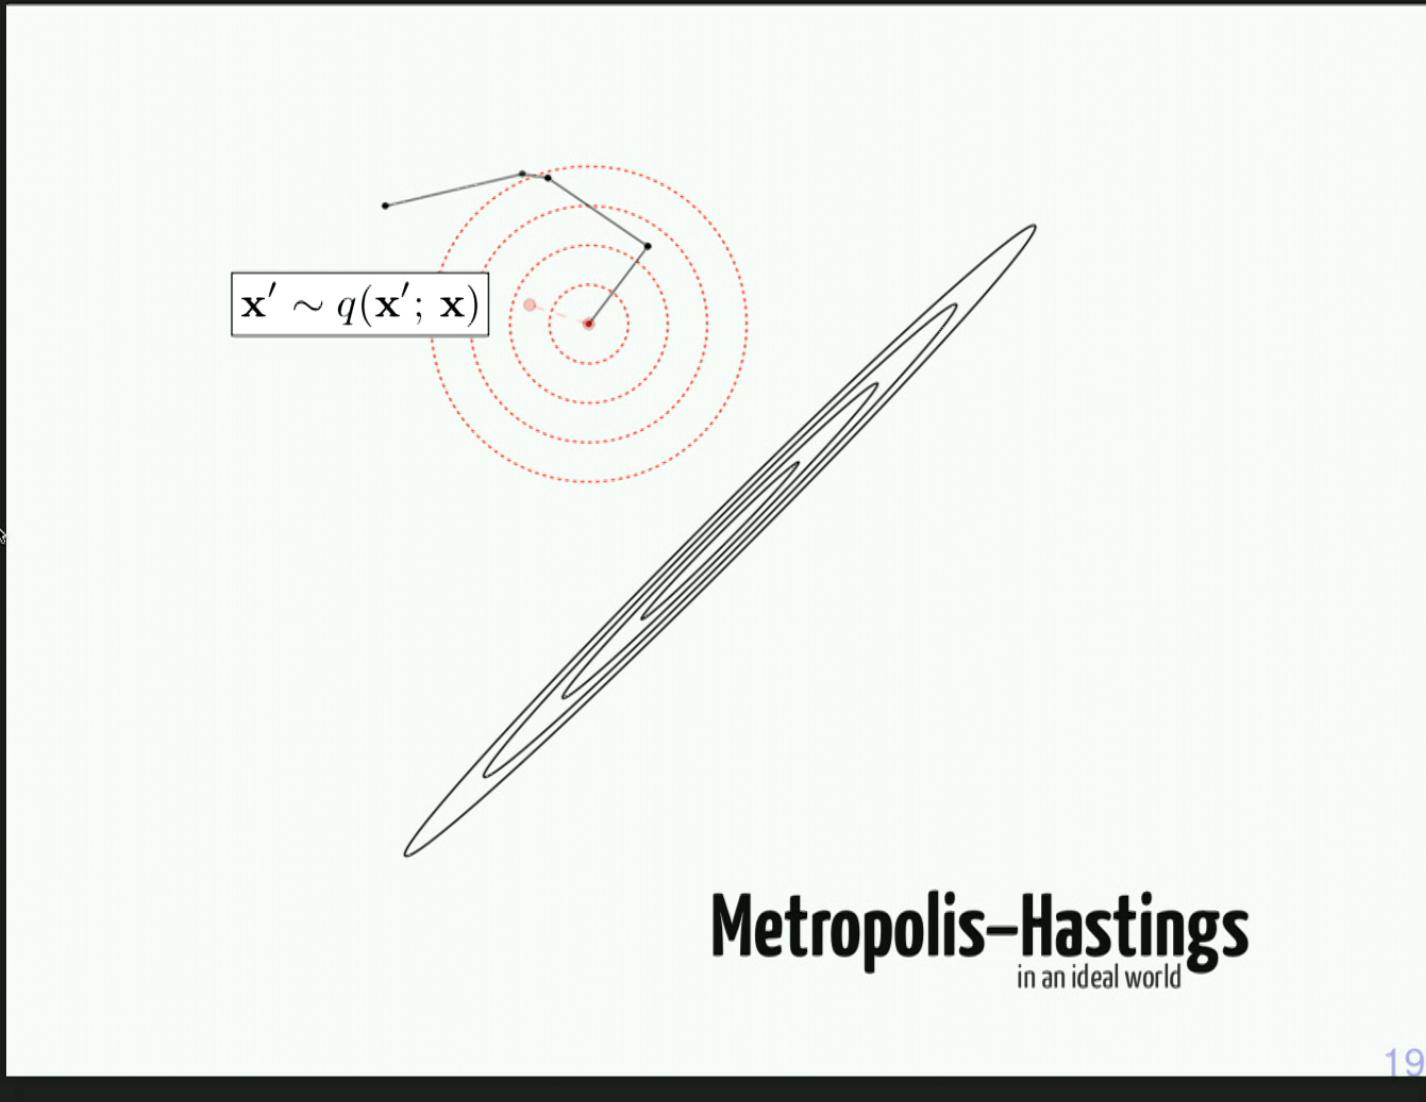


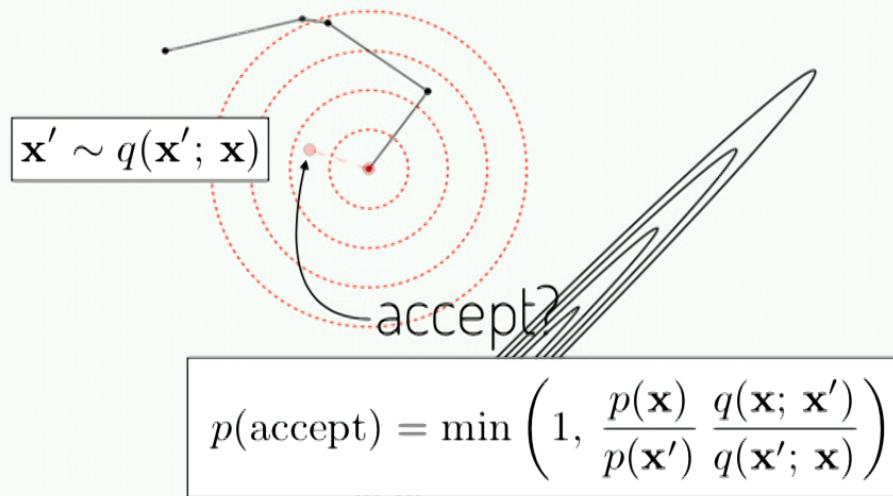
Metropolis-Hastings

in an ideal world

16



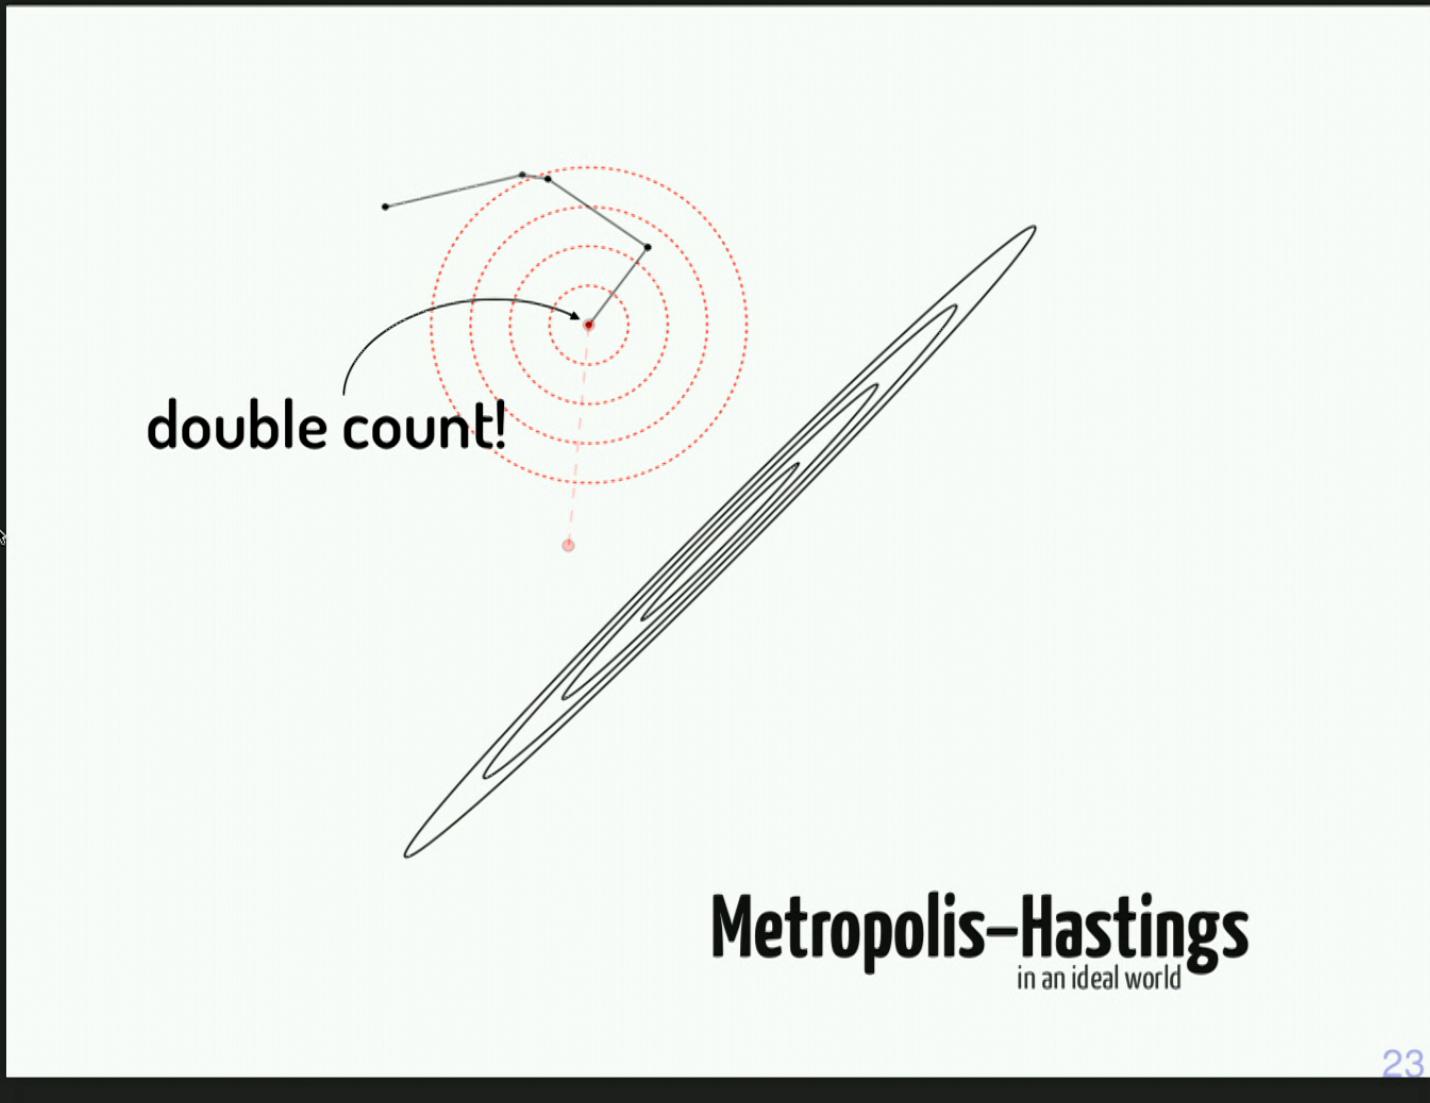




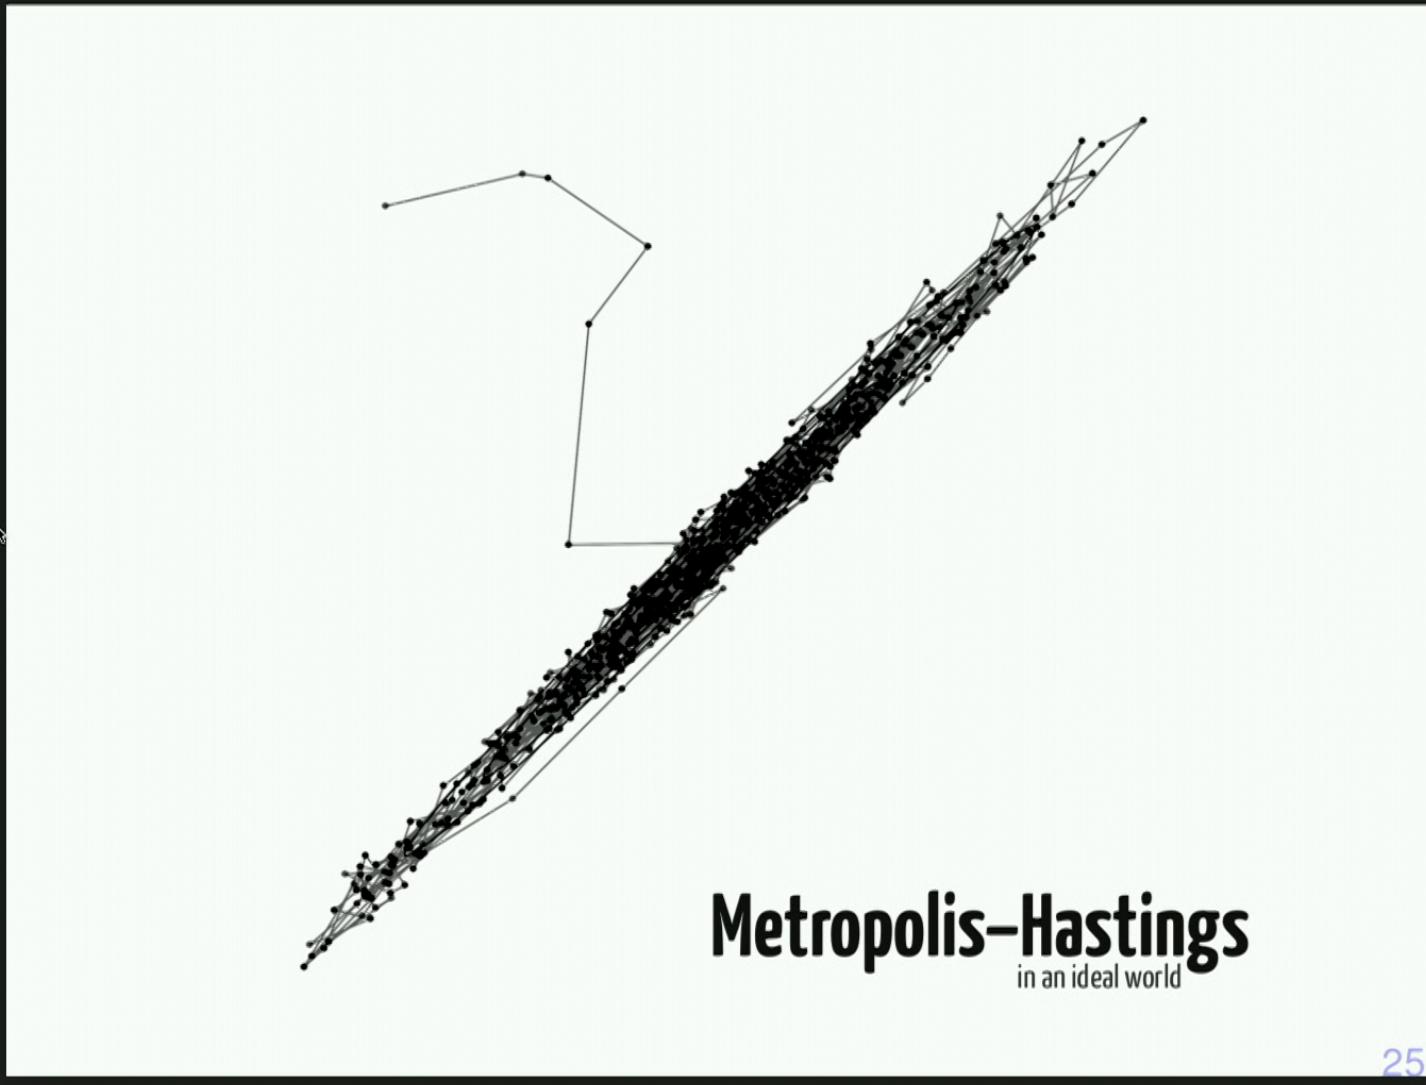
not this time.

Metropolis-Hastings in an ideal world

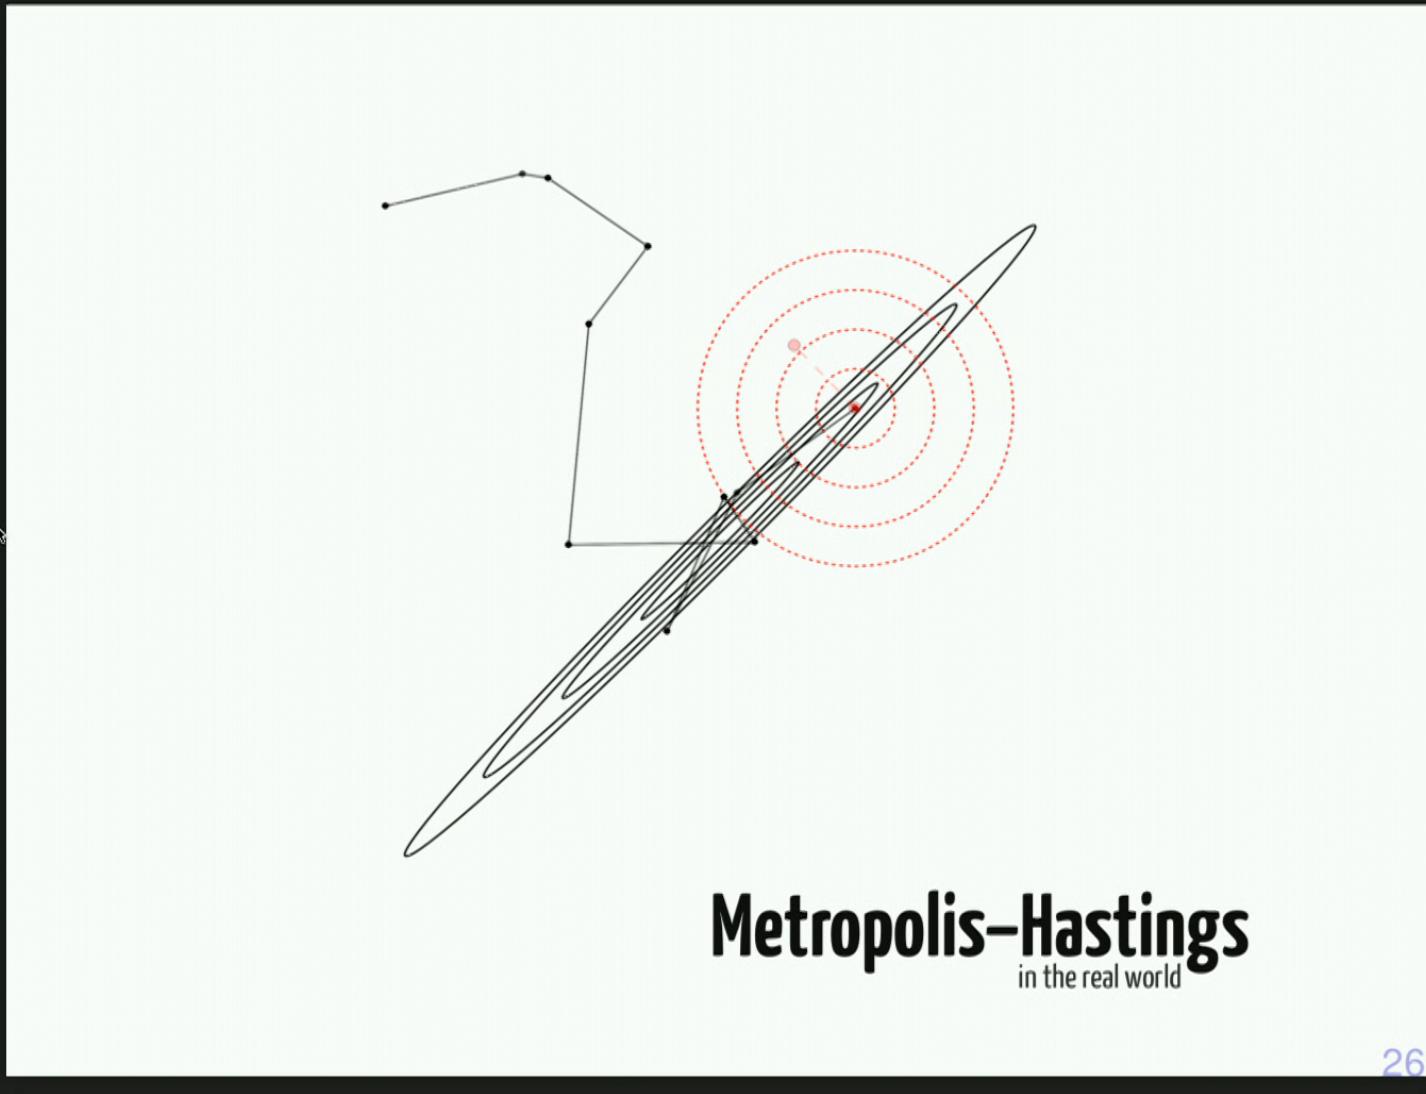
21



23



25



26



Metropolis-Hastings in the real world

Page 31 of 35

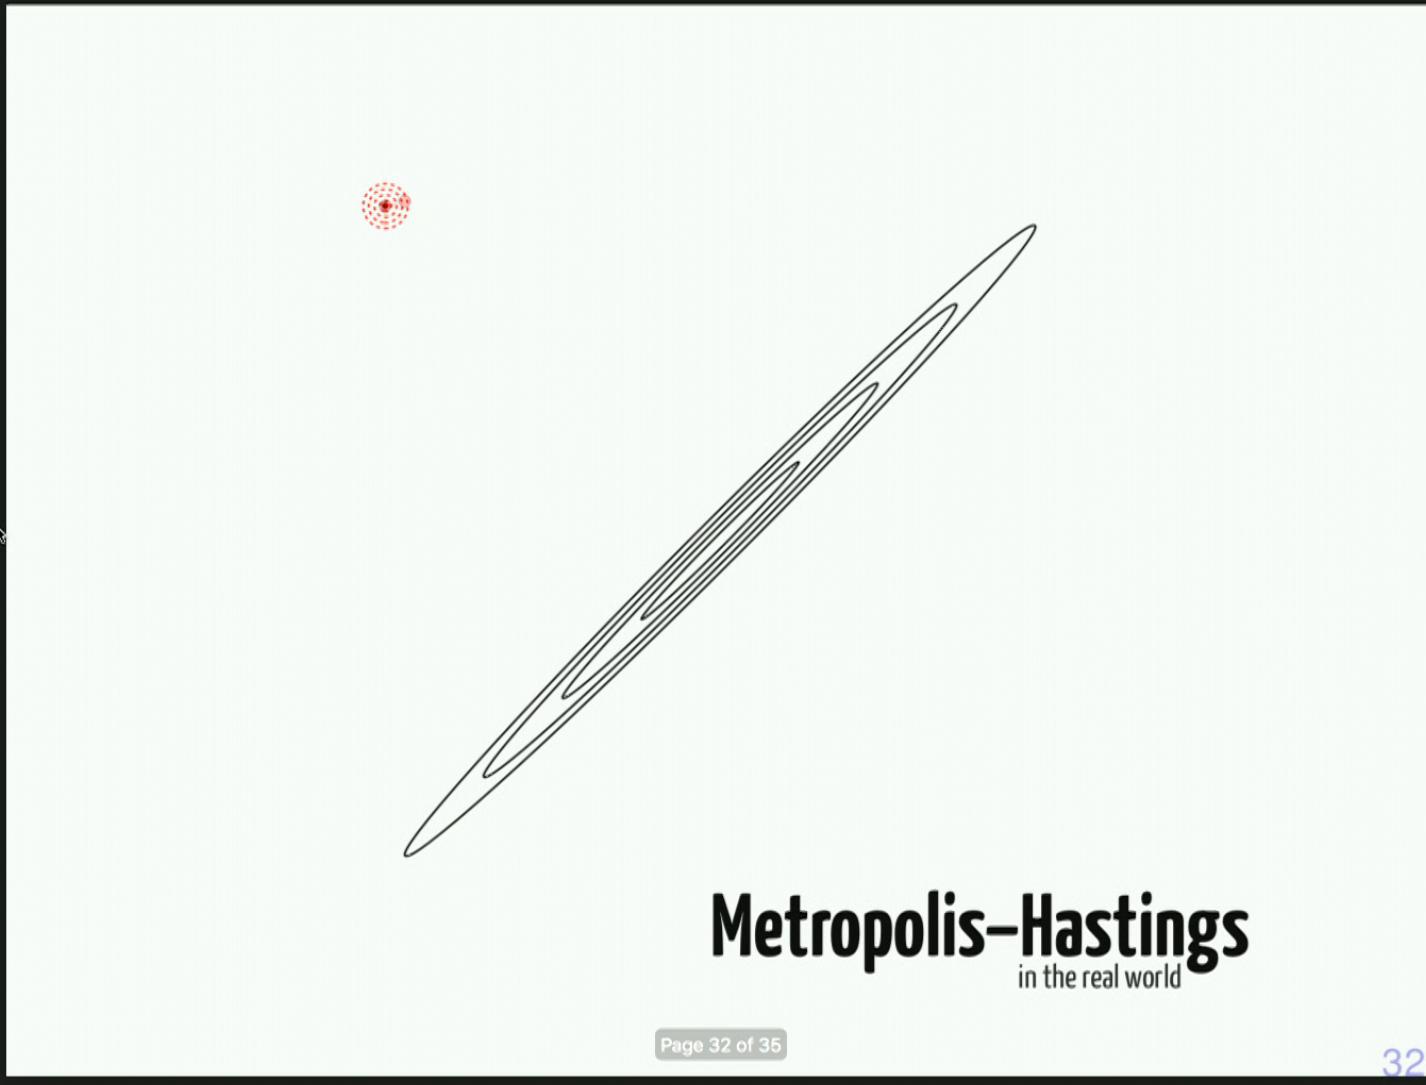
31

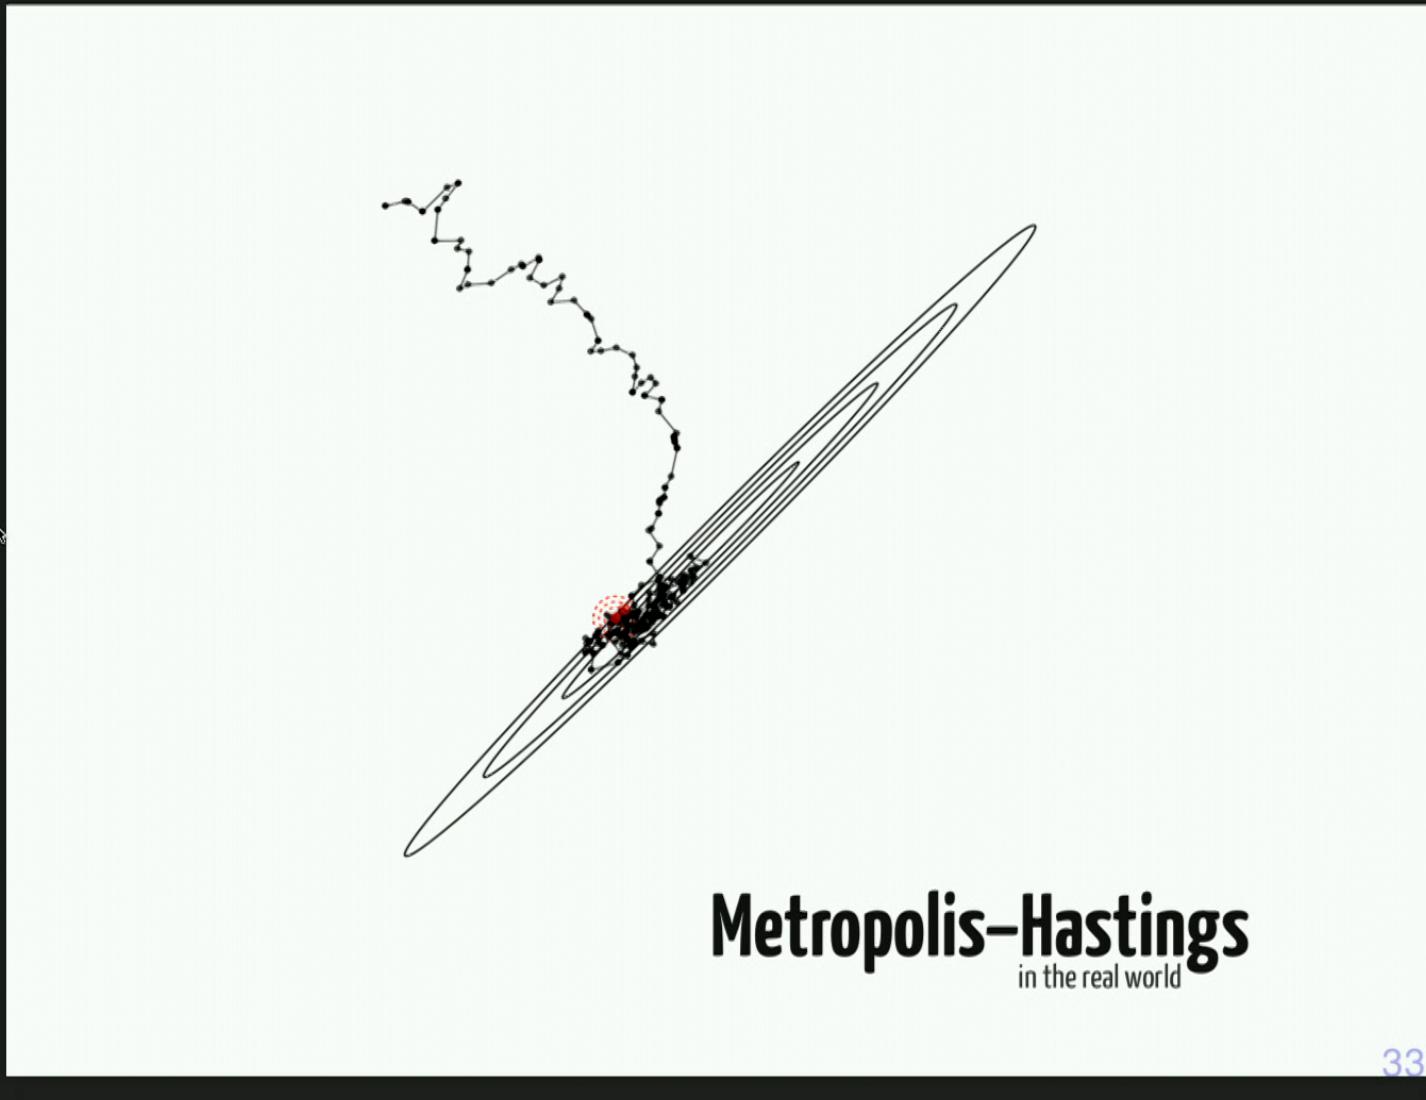


the
Small Acceptance Fraction
problem

Metropolis-Hastings
in the real world

31







Metropolis-Hastings

in the real world

34

File Edit View Run Kernel Tabs Settings Help

mcmc.ipynb × perlmutter99-julia-filled.ipynb × Untitled.ipynb ×

Julia 1.2.0

Name Last Modified

- mcmc.ipynb an hour ago
- p99-data.txt an hour ago
- perlmutter99-julia-filled.ipynb an hour ago
- Untitled.ipynb an hour ago

Imagine we are measuring voltages of a sine wave, and we want to constrain the amplitude and phase.

Our measurements have Gaussian noise of known standard deviation, and occur at known times (assume we can measure time so well that there is practically no uncertainty).

```
[18]: # "y" are our measurements
y = Array([-1.64, 4.63, 6.09, 5.05, 2.66, -0.90, -6.09, -7.16, -2.54, -1.84])
# "noise_std" is the standard deviation of the Gaussian noise on our measurements
noise_std = 1.0
```

```
[18]: nsamples = length(y)
# "times" are when the measurements were taken (evenly spaced)
times = collect(LinRange(0.0, 2.0*pi, nsamples));
```

```
[6]: using Plots
```

```
[7]: plot(times, y)
```

Mode: Command Ln 1, Col 103 mcmc.ipynb

GitHub, Inc. (US) | https://github.com/dstndstn?tab=repositories 120% ... wi

Search or jump to... Pull requests Issues Marketplace Explore

Overview Repositories 71 Projects 0 Packages 0 Stars 9 Followers 123 Following 7

Find a repository... Type: All Language: All New

mcmc-notes
TeX Updated 20 hours ago ★ Star

unwise.me
Web service powering unwise.me ★ Star

tractor
The Tractor: measuring astronomical sources via probabilistic inference ★ Star

astrometry.net
Astrometry.net -- automatic recognition of astronomical images ★ Star

Leaflet
Forked from Leaflet/Leaflet ★ Star

Dustin Lang
dstndstn

Edit profile

Perimeter Institute (@PerimeterInsti...
Waterloo, Canada
<http://dstn.astrometry.net>

Organizations

PI

A screenshot of a web browser displaying a GitHub repository page. The URL in the address bar is <https://github.com/dstndstn/mcmc-notes>. The repository name is **mcmc-notes**. The page shows basic repository statistics: 2 commits, 1 branch, 0 packages, 0 releases, and 1 contributor. The latest commit is `0e06aba` from 20 hours ago, labeled "Initial commit". The repository contains files: `README.md`, `header.tex`, and `mcmc.tex`, all updated 20 hours ago. The `README.md` file is open, showing the text "mcmc-notes".

dstndstn / **mcmc-notes**

Code Issues Pull requests Projects Wiki Security Insights Settings

No description, website, or topics provided.

Manage topics

2 commits 1 branch 0 packages 0 releases 1 contributor

Branch: master New pull request Create new file Upload files Find File Clone or download

dstndstn quick notes Latest commit `0e06aba` 20 hours ago

`README.md` Initial commit 23 hours ago

`header.tex` quick notes 20 hours ago

`mcmc.tex` quick notes 20 hours ago

`README.md`

mcmc-notes

© 2019 GitHub, Inc. Terms Privacy Security Status Help



Contact GitHub Pricing API Training Blog About

File Edit View Run Kernel Tabs Settings Help

/ ... / dlang / mcmc /

Name	Last Modified
mcmc	2 minutes ago
mcmc2	2 minutes ago
mcmc.ipynb	an hour ago
p99-data.txt	an hour ago
perlmutter99-julia-fille...	an hour ago
Untitled.ipynb	an hour ago

```
mypasswd          Untitled93.ipynb      slurm-38979.out    slurm-39055.out  Untitled25.ipynb
mypasswd-         Untitled94.ipynb      slurm-38980.out    slurm-39056.out  Untitled26.ipynb
myproject-mpi.sif Untitled95.ipynb      slurm-38981.out    slurm-39057.out  Untitled27.ipynb
myproject.sif     Untitled96.ipynb      slurm-38982.out    slurm-39058.out  Untitled28.ipynb
orig-binary_neutron_star_example.py Untitled97.ipynb      slurm-38983.out    slurm-39059.out  Untitled29.ipynb
orig-binary_neutron_star_example.py- Untitled98.ipynb      slurm-38984.out    slurm-39060.out  Untitled2.ipynb
outdir           Untitled99.ipynb      slurm-38985.out    slurm-39061.out  Untitled30.ipynb
overlay          Untitled9.ipynb       slurm-38986.out    slurm-39070.out  Untitled31.ipynb
Pictures          Untitled.ipynb       slurm-38987.out    slurm-39071.out  Untitled32.ipynb
Videos            Untitled.ipynb       slurm-38988.out    slurm-39072.out  Untitled33.ipynb
plot.png          Untitled.ipynb       slurm-38989.out    slurm-39073.out  Untitled34.ipynb
Public            writable          slurm-38989.out    slurm-39073.out  Untitled34.ipynb
writable         

dlang@mn001:~$ git clone http://github.com/dstndstn/mcmc-notes
Cloning into 'mcmc-notes'...
warning: redirecting to https://github.com/dstndstn/mcmc-notes/
remote: Enumerating objects: 11, done.
remote: Counting objects: 100% (11/11), done.
remote: Compressing objects: 100% (9/9), done.
remote: Total 11 (delta 1), reused 7 (delta 0), pack-reused 0
Unpacking objects: 100% (11/11), done.
dlang@mn001:~$
```

1 Terminal 1 5

File Edit View Run Kernel Tabs Settings Help

New Launcher / ... / dlang / mcmc /

Name	Last Modified
mcmc	2 minutes ago
mcmc2	2 minutes ago
mcmc.ipynb	an hour ago
p99-data.txt	an hour ago
perlmutter99-julia-fille...	an hour ago
Untitled.ipynb	an hour ago

```
mypasswd- slurm-38980.out slurm-39056.out Untitled26.ipynb
myproject-mpi.sif Untitled95.ipynb slurm-38981.out slurm-39057.out Untitled27.ipynb
myproject.sif Untitled96.ipynb slurm-38982.out slurm-39058.out Untitled28.ipynb
orig-binary_neutron_star_example.py Untitled97.ipynb slurm-38983.out slurm-39059.out Untitled29.ipynb
orig-binary_neutron_star_example.py- Untitled98.ipynb slurm-38984.out slurm-39060.out Untitled2.ipynb
outdir Untitled99.ipynb slurm-38985.out slurm-39061.out Untitled30.ipynb
overlay Untitled9.ipynb slurm-38986.out slurm-39070.out Untitled31.ipynb
Pictures Untitled.ipynb slurm-38987.out slurm-39071.out Untitled32.ipynb
plot.png Videos slurm-38988.out slurm-39072.out Untitled33.ipynb
Public visit slurm-38989.out slurm-39073.out Untitled34.ipynb
writable

dlang@mn001:~$ git clone http://github.com/dstndstn/mcmc-notes
Cloning into 'mcmc-notes'...
warning: redirecting to https://github.com/dstndstn/mcmc-notes/
remote: Enumerating objects: 11, done.
remote: Counting objects: 100% (11/11), done.
remote: Compressing objects: 100% (9/9), done.
remote: Total 11 (delta 1), reused 7 (delta 0), pack-reused 0
Unpacking objects: 100% (11/11), done.
dlang@mn001:~$ cd mcmc-notes
dlang@mn001:~/mcmc-notes$ ls
header.tex mcmc.ipynb mcmc.tex p99-data.txt README.md
dlang@mn001:~/mcmc-notes$
```

1 s 5 Terminal 1

File Edit View Run Kernel Tabs Settings Help

mcmc.ipynb Terminal 1 mcmc.ipynb perlmutter99-ju Untitled.ipynb

Name Last Modified

- header.tex seconds ago
- mcmc.ipynb** seconds ago
- mcmc.tex seconds ago
- p99-data.txt seconds ago
- README.md seconds ago

Imagine we are measuring voltages of a sine wave, and we want to constrain the amplitude and phase.

Our measurements have Gaussian noise of known standard deviation, and occur at known times (assume we can measure time so well that there is practically no uncertainty).

```
[18]: # "y" are our measurements
y = Array([-1.64, 4.63, 6.09, 5.05, 2.66, -0.90, -6.09, -7.16, -2.54, -1.84])
# "noise_std" is the standard deviation of the Gaussian noise on our measurements
noise_std = 1.0

nsamples = length(y)
# "times" are when the measurements were taken (evenly spaced)
times = collect(LinRange(0.0, 2.0*pi, nsamples));
```

```
[6]: using Plots
```

```
[7]: plot(times, y)
```

```
[7]:
```

Mode: Command Ln 1, Col 1 mcmc.ipynb

A screenshot of a Jupyter Notebook interface. On the left is a sidebar with icons for file operations like new file, upload, and settings. Below it is a file browser showing a directory structure under `/ ... / dlang / mcmc-notes /`. The files listed are `header.tex`, `mcmc.ipynb` (which is selected), `mcmc.tex`, `p99-data.txt`, and `README.md`. The `mcmc.ipynb` file has a blue background. To the right is a terminal window titled "Terminal 1" with the tab "mcmc.ipynb". It contains the command `dlang@mn001:~/mcmc-notes$ git clone http://github.com/dstndstn/mcmc-notes`. At the bottom left are navigation icons for the file browser, and at the bottom right is the text "Terminal 1".

File Edit View Run Kernel Tabs Settings Help

/ ... / dlang / mcmc-notes /

Name	Last Modified
header.tex	2 minutes ago
mcmc.ipynb	2 minutes ago
mcmc.tex	2 minutes ago
p99-data.txt	2 minutes ago
README.md	2 minutes ago

Terminal 1 | mcmc.ipynb | Markdown | Julia 1.2.0

Imagine we are measuring voltages of a sine wave, and we want to constrain the amplitude and phase.

Our measurements have Gaussian noise of known standard deviation, and occur at known times (assume we can measure time so well that there is practically no uncertainty).

```
[18]: # "y" are our measurements
y = Array([-1.64, 4.63, 6.09, 5.05, 2.66, -0.90, -6.09, -7.16, -2.54, -1.84])
# "noise_std" is the standard deviation of the Gaussian noise on our measurements
noise_std = 1.0

nsamples = length(y)
# "times" are when the measurements were taken (evenly spaced)
times = collect(LinRange(0.0, 2.0*pi, nsamples));
```

```
[6]: using Plots
```

```
[7]: plot(times, y)
```

Mode: Command | Ln 1, Col 1 | mcmc.ipynb

https://symmetry1/user/dlang/lab

File Edit View Run Kernel Tabs Settings Help

Terminal 1 mcmc.ipynb

Julia 1.2.0

Name

header.tex

mcmc.ipynb

mcmc.tex

p99-data.txt

README.md

[6]: `using Plots`

[7]: `plot(times, y)`

[7]:

Mode: Command Ln 1, Col 1 mcmc.ipynb

File Edit View Run Kernel Tabs Settings Help

Terminal 1 mcmc.ipynb

Name

- header.tex
- mcmc.ipynb
- mcmc.tex
- p99-data.txt
- README.md

Julia 1.2.0

Now, we have to write our *model* function. Our model is that this is a sine wave with parameters of *phase offset* and *amplitude*.

```
[15]: function sine_model(t, phase, amplitude)
    return amplitude * sin.(t .- phase)
end;
```

```
[17]: tplot = collect(LinRange(0.0, 2.0*pi, 200));
# I'm going to guess at parameters: phase = 0, amplitude = 6
plot(tplot, sine_model(tplot, 0., 6))
plot!(times, y)
```

[17]:

Mode: Command Ln 1, Col 1 mcmc.ipynb

File Edit View Run Kernel Tabs Settings Help

Terminal 1 mcmc.ipynb

[17]: tplot = collect(LinRange(0.0, 2.0*pi, 200));
I'm going to guess at parameters: phase = 0, amplitude = 6
plot(tplot, sine_model(tplot, 0., 6))
plot!(times, y)

[17]:

Mode: Command Ln 1, Col 1 mcmc.ipynb

File Edit View Run Kernel Tabs Settings Help

Terminal 1 mcmc.ipynb

Name

header.tex

mcmc.ipynb

mcmc.tex

p99-data.txt

README.md

Now, in order to use MCMC, we must write down our *likelihood*. Given model parameters, how probable are the measurements we made?

```
[29]: function sine_log_likelihood(t, y, phase, amplitude)
    ymodel = sine_model(t, phase, amplitude)
    log_likelihood = -0.5 * sum((y - ymodel).^2 / noise_std.^2)
    return log_likelihood
end;
```

```
[30]: sine_log_likelihood(times, y, 0., 6.)
```

```
[30]: -6.261172416426372
```

```
[31]: sine_log_likelihood(times, y, 0.05, 6.)
```

```
[31]: -5.560820776311154
```

Now, since this is a low-dimensional problem, we can just grid up the parameter space and compute the likelihood at each position! In real problems, you can't do this....

```
[28]: plo,phi = 0., 0.5
alo,ahi = 5., 8.
pp = range(plo, stop=phi, length=100)
aa = range(alo, stop=ahi, length=100)
LL = zeros((length(pp), length(aa)))
```

The screenshot shows a Jupyter Notebook interface running on a web browser. The top bar includes standard browser controls like back, forward, and search, along with tabs for various open documents. The main window has a toolbar with icons for file operations, followed by a menu bar with File, Edit, View, Run, Kernel, Tabs, Settings, and Help. A sidebar on the left lists files in the current directory: header.tex, mcmc.ipynb (which is selected and highlighted in blue), mcmc.tex, p99-data.txt, and README.md.

The central area contains two panes: Terminal 1 and mcmc.ipynb. The mcmc.ipynb pane displays the following text:

Now, in order to use MCMC, we must write down our *likelihood*. Given model parameters, how probable are the measurements we made?

```
[29]: function sine_log_likelihood(t, y, phase, amplitude)
    ymodel = sine_model(t, phase, amplitude)
    log_likelihood = -0.5 * sum((y - ymodel).^2 / noise_std.^2)
    return log_likelihood
end;
```

```
[30]: sine_log_likelihood(times, y, 0., 6.)
[30]: -6.261172416426372
```

```
[31]: sine_log_likelihood(times, y, 0.05, 6.)
[31]: -5.560820776311154
```

Now, since this is a low-dimensional problem, we can just grid up the parameter space and compute the likelihood at each position! In real problems, you can't do this....

```
[28]: plo,phi = 0., 0.5
alo,ahi = 5., 8.
pp = range(plo, stop=phi, length=100)
aa = range(alo, stop=ahi, length=100)
LL = zeros((length(pp), length(aa)))
```

At the bottom, status indicators show "Julia 1.2.0 | Idle", "Mode: Edit", "Ln 1, Col 18", and the file name "mcmc.ipynb".

The screenshot shows a Jupyter Notebook interface running on a web browser. The top bar includes standard browser controls like back, forward, and search, along with tabs for various documents and a Julia 1.2.0 kernel.

The left sidebar displays a file tree:

- / ... / dlang / mcmc-notes /
- Name
- header.tex
- mcmc.ipynb (selected)
- mcmc.tex
- p99-data.txt
- README.md

The main area has two tabs: "Terminal 1" and "mcmc.ipynb". The "mcmc.ipynb" tab is active, showing the following content:

Now, in order to use MCMC, we must write down our *likelihood*. Given model parameters, how probable are the measurements we made?

```
[29]: function sine_log_likelihood(t, y, phase, amplitude)
    ymodel = sine_model(t, phase, amplitude)
    log_likelihood = -0.5 * sum((y - ymodel).^2 / noise_std.^2)
    return log_likelihood
end;
```

```
[30]: sine_log_likelihood(times, y, 0., 6.)
[30]: -6.261172416426372
```

```
[31]: sine_log_likelihood(times, y, 0.05, 6.)
[31]: -5.560820776311154
```

Now, since this is a low-dimensional problem, we can just grid up the parameter space and compute the likelihood at each position! In real problems, you can't do this....

```
[28]: plo,phi = 0., 0.5
alo,ahi = 5., 8.
pp = range(plo, stop=phi, length=100)
aa = range(alo, stop=ahi, length=100)
LL = zeros((length(pp), length(aa)))
```

At the bottom, the status bar shows "Julia 1.2.0 | Idle", "Mode: Edit", "Ln 3, Col 36", and the file name "mcmc.ipynb".

File Edit View Run Kernel Tabs Settings Help

Terminal 1 mcmc.ipynb Julia 1.2.0

Name

header.tex

mcmc.ipynb

mcmc.tex

p99-data.txt

README.md

Now, in order to use MCMC, we must write down our *likelihood*. Given model parameters, how probable are the measurements we made?

```
[29]: function sine_log_likelihood(t, y, phase, amplitude)
    ymodel = sine_model(t, phase, amplitude)
    log_likelihood = -0.5 * sum((y - ymodel).^2 / noise_std.^2)
    return log_likelihood
end;
```

```
[1]: sine_log_likelihood(times, y, 0., 6.)
```

```
UndefVarError: sine_log_likelihood not defined
Stacktrace:
 [1] top-level scope at In[1]:1
```

```
[31]: sine_log_likelihood(times, y, 0.05, 6.)
```

```
[31]: -5.560820776311154
```

Now, since this is a low-dimensional problem, we can just grid up the parameter space and compute the likelihood at each position! In real problems, you can't do this....

```
[28]: plo,phi = 0., 0.5
```

Mode: Command Ln 1, Col 1 mcmc.ipynb

The screenshot shows a Jupyter Notebook interface running on a web browser. The top bar includes standard browser controls like back, forward, and search, along with tabs for various open documents. The main area has a toolbar with icons for file operations, followed by a sidebar containing a file tree and a list of files. The central workspace contains a terminal window and a notebook cell.

Terminal 1:

```
mcmc.ipynb
```

File Tree:

- / ... / dlang / mcmc-notes /
- Name
- header.tex
- mcmc.ipynb**
- mcmc.tex
- p99-data.txt
- README.md

Notebook Cell 1:

```
Imagine we are measuring voltages of a sine wave, and we want to constrain the amplitude and phase.  
Our measurements have Gaussian noise of known standard deviation, and occur at known times (assume we can measure time so well that there is practically no uncertainty).
```

Notebook Cell 2:

```
[2]: # "y" are our measurements  
y = Array([-1.64, 4.63, 6.09, 5.05, 2.66, -0.90, -6.09, -7.16, -2.54, -1.84])  
# "noise_std" is the standard deviation of the Gaussian noise on our measurements  
noise_std = 1.0  
  
nsamples = length(y)  
# "times" are when the measurements were taken (evenly spaced)  
times = collect(LinRange(0.0, 2.0*pi, nsamples));
```

Notebook Cell 3:

```
[3]: using Plots
```

Notebook Cell 4:

```
[*]: plot(times, y)
```

Notebook Cell 5:

```
Now, we have to write our model function. Our model is that this is a sine wave with parameters of phase offset and amplitude.
```

Notebook Cell 6:

```
[15]: function sine_model(t, phase, amplitude)  
    return amplitude * sin.(t .- phase)  
end;
```

1 s 6 Mode: Command \otimes Ln 1, Col 1 mcmc.ipynb

File Edit View Run Kernel Tabs Settings Help

Terminal 1 mcmc.ipynb Julia 1.2.0

nsamples = length(y)
"times" are when the measurements were taken (evenly spaced)
times = collect(LinRange(0.0, 2.0*pi, nsamples));

[3]: using Plots

[4]: plot(times, y)

[4]:

The plot shows a single blue line labeled 'y1'. The x-axis is labeled 'e' and ranges from 0 to approximately 3.14. The y-axis ranges from -5 to 6. The curve starts at (0, -1.5), rises to a peak of about 6 at x ≈ pi/2, falls to a minimum of about -5 at x = pi, and then rises again towards x = 3pi/2.

Mode: Command Ln 1, Col 1 mcmc.ipynb

File Edit View Run Kernel Tabs Settings Help

Terminal 1 mcmc.ipynb Julia 1.2.0

```
[7]: function sine_log_likelihood(t, y, phase, amplitude)
    ymodel = sine_model(t, phase, amplitude)
    log_likelihood = -0.5 * sum((y - ymodel).^2 / noise_std.^2)
    return log_likelihood
end;

[8]: sine_log_likelihood(times, y, 0., 6.)
[8]: -6.261172416426372

[31]: sine_log_likelihood(times, y, 0.05, 6.)
[31]: -5.560820776311154
```

Now, since this is a low-dimensional problem, we can just grid up the parameter space and compute the likelihood at each position! In real problems, you can't do this....

```
[28]: plo,phi = 0., 0.5
alo,ahi = 5., 8.
pp = range(plo, stop=phi, length=100)
aa = range(alo, stop=ahi, length=100)
LL = zeros((length(pp), length(aa)))
for i in 1:length(pp)
    for j in 1:length(aa)
        LL[i, j] = sine_log_likelihood(times, y, pp[j], aa[i])
    end
end
```

File Edit View Run Kernel Tabs Settings Help

Terminal 1 mcmc.ipynb Julia 1.2.0

Name

header.tex

mcmc.ipynb

mcmc.tex

p99-data.txt

README.md

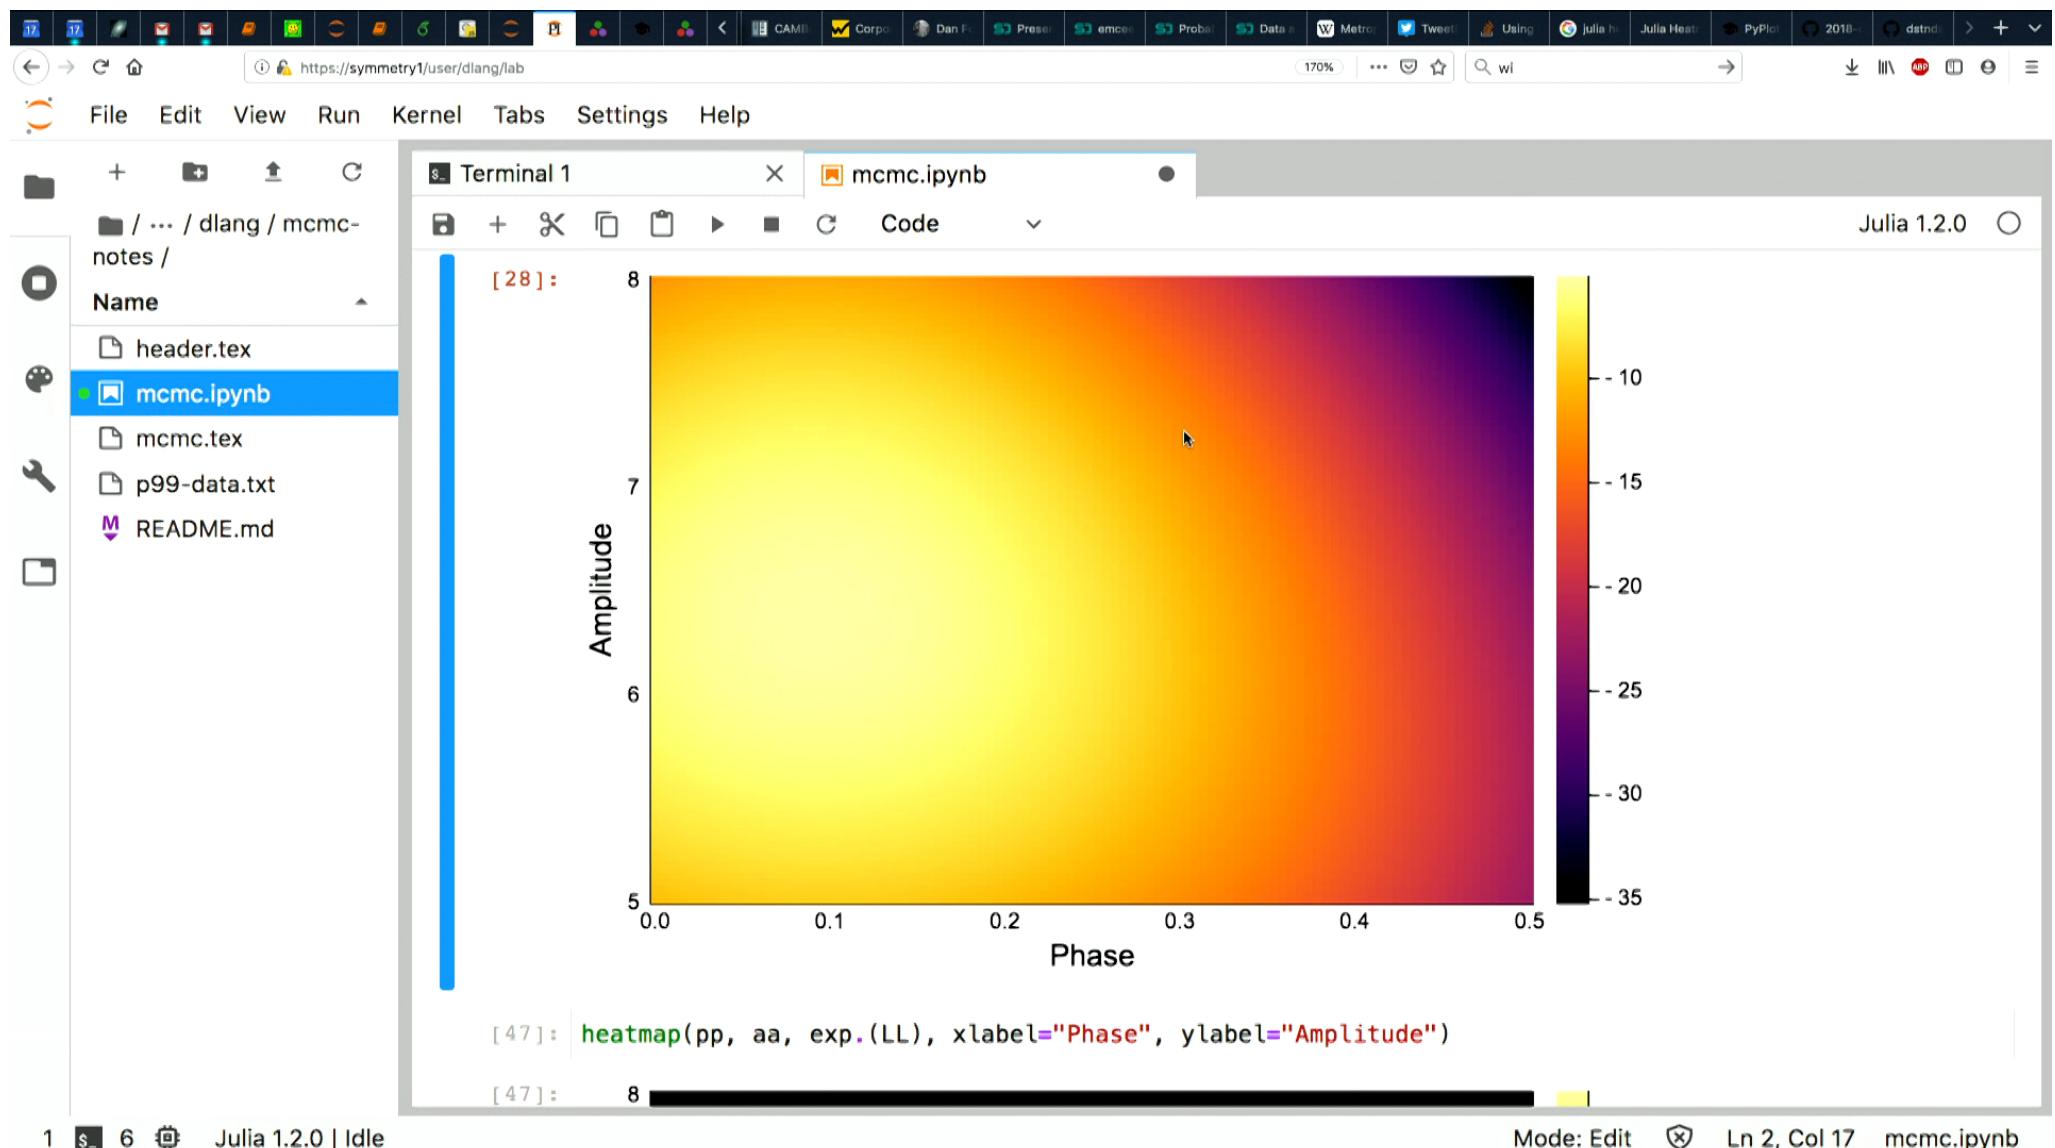
[28]:

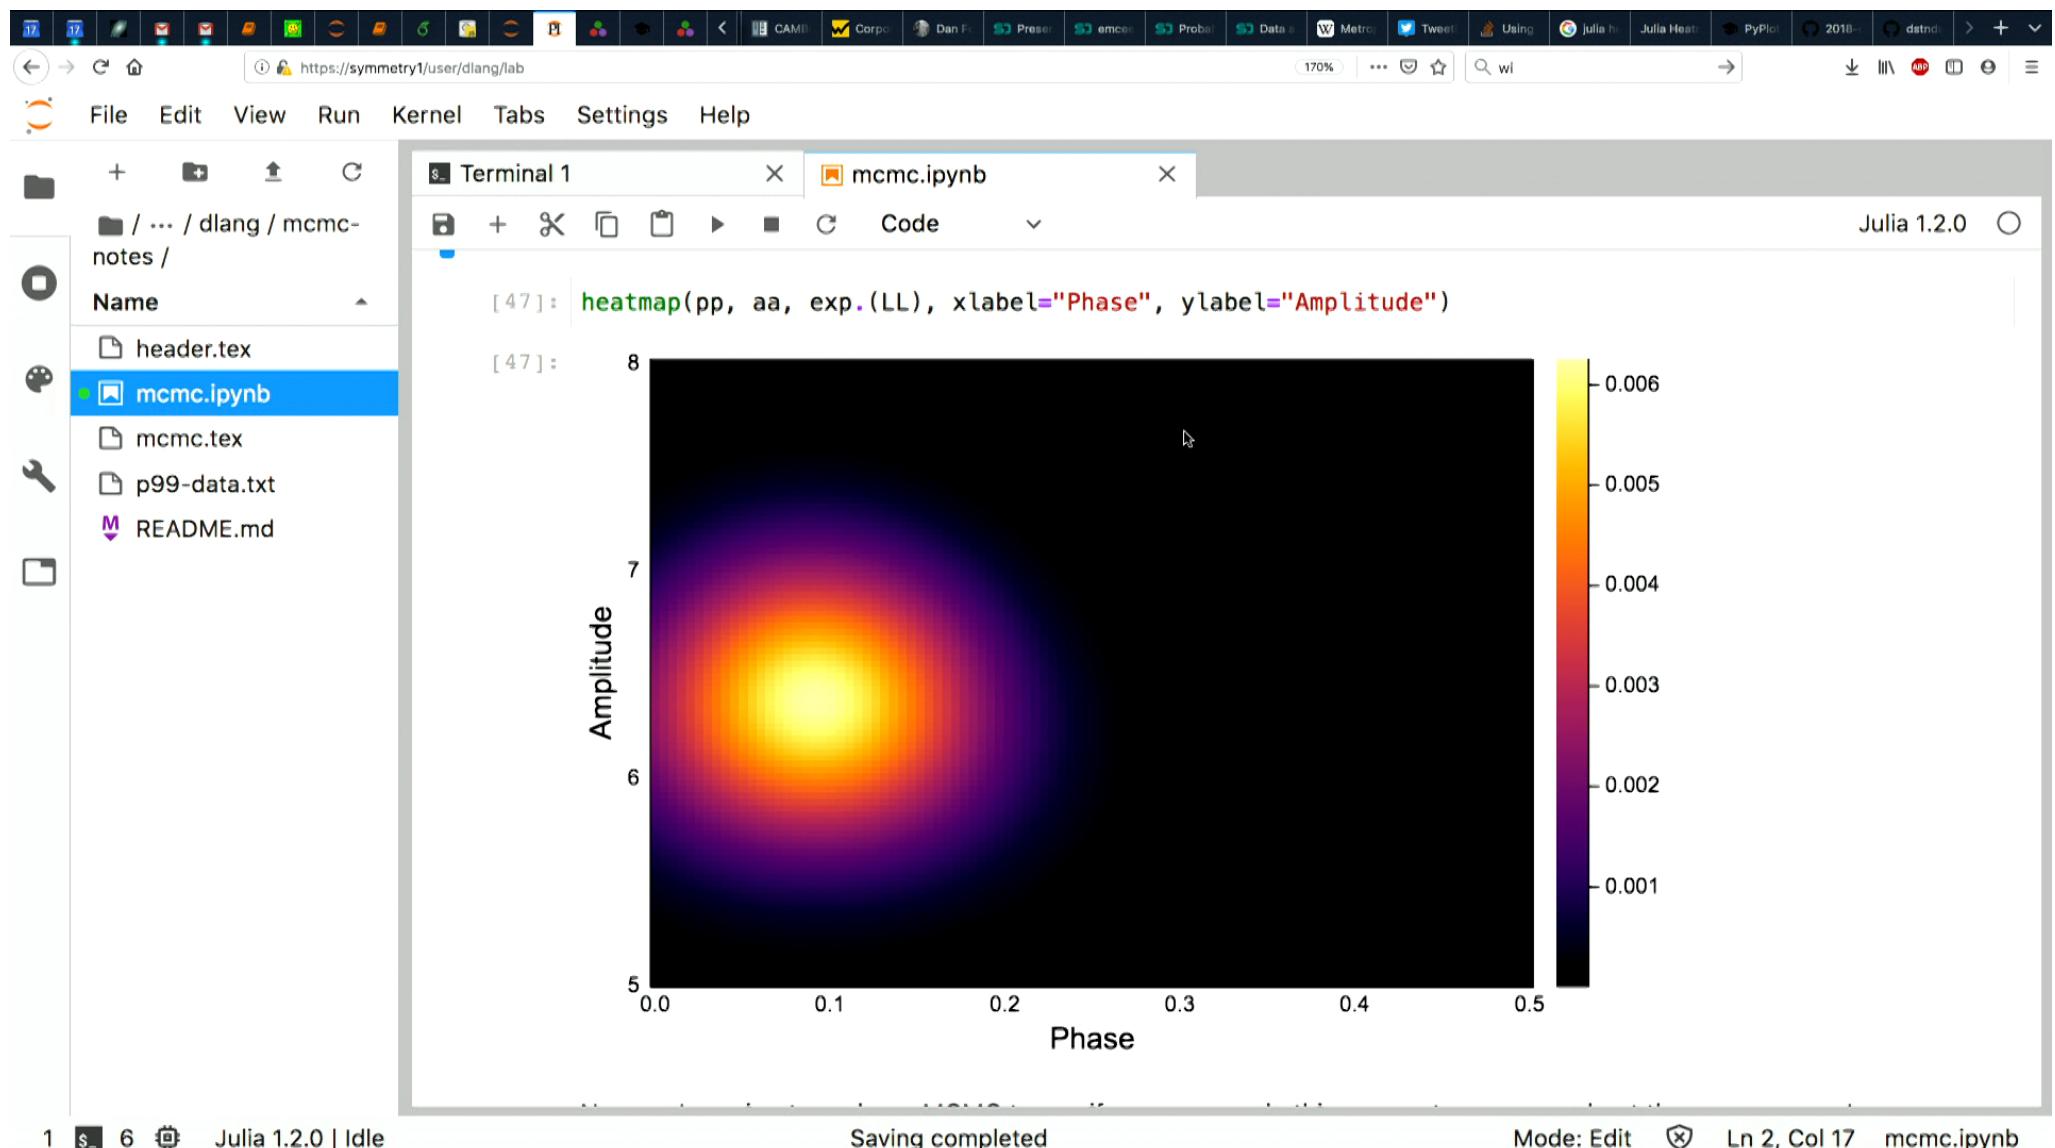
```
plo,phi = 0., 0.5
alo,ahi = 5., 8.
pp = range(plo, stop=phi, length=100)
aa = range(alo, stop=ahi, length=100)
LL = zeros((length(pp), length(aa)))
for i in 1:length(pp)
    for j in 1:length(aa)
        LL[i, j] = sine_log_likelihood(times, y, pp[j], aa[i])
    end
end
heatmap(pp, aa, LL, xlabel="Phase", ylabel="Amplitude")
```

[28]:

Mode: Edit

Ln 2, Col 17 mcmc.ipynb





The screenshot shows a Jupyter Notebook interface running on a Mac OS X desktop. The top menu bar includes File, Edit, View, Run, Kernel, Tabs, Settings, and Help. The toolbar has icons for file operations like New, Open, Save, and Run. The address bar shows the URL <https://symmetry1/user/dlang/lab>. The status bar at the bottom indicates "Julia 1.2.0 | Idle".

The left sidebar displays a file tree:

- / ... / dlang / mcmc-notes /
- Name
- header.tex
- mcmc.ipynb (selected)
- mcmc.tex
- p99-data.txt
- README.md

The main area contains two tabs: Terminal 1 and mcmc.ipynb. The mcmc.ipynb tab is active and shows the following content:

Now, we're going to code up MCMC to see if we can sample this parameter space and get the same answer!

First, we need to go from a log-likelihood to a log-posterior. This *requires* specifying a *prior* on the parameters.

A common thing to do is to choose a "flat" prior: the prior is constant over the entire range. This is technically wrong ("improper prior").

```
[32]: function sine_log_posterior(t, y, phase, amplitude)
    log_like = sine_log_likelihood(t, y, phase, amplitude)
    # Flat prior! No preference for phase, amplitude values.
    log_prior = 0.
    return log_like + log_prior
end;
```

* Now the actual MCMC algorithm! We need to choose a "proposal distribution" -- how to select the next candidate parameters to jump to. We will make these Gaussians with fixed variances.

```
[36]: chain = []

# initial parameter guess
param_phase = 0.
param_amp = 6.

# proposal jump sizes
jump_phase = 0.01
```

Mode: Edit Ln 2, Col 17 mcmc.ipynb

The screenshot shows a Jupyter Notebook interface with a Julia kernel. The left sidebar displays a file tree with files like header.tex, mcmc.ipynb (selected), mcmc.tex, p99-data.txt, and README.md. The main area has two tabs: Terminal 1 and mcmc.ipynb. The mcmc.ipynb tab contains a code cell with the following Julia code:

```
[36]: chain = []

# initial parameter guess
param_phase = 0.
param_amp = 6.

# proposal jump sizes
jump_phase = 0.01
jump_amp = 0.1

# initial log-posterior.
logprob = sine_log_posterior(times, y, param_phase, param_amp)

# how many steps to take
nsteps = 1_000

for i in 1:nsteps

    # propose new parameter values
    param_phase_new = param_phase + randn() * jump_phase
    param_amp_new   = param_amp   + randn() * jump_amp

    # compute log-posterior at new parameters
    logprob_new = sine_log_posterior(times, y, param_phase_new, param_amp_new)

    if logprob_new > logprob
```

1 s 6 Julia 1.2.0 | Idle

Mode: Edit Ln 2, Col 17 mcmc.ipynb

The screenshot shows a Jupyter Notebook interface with a Julia kernel. The left sidebar displays a file tree with files like `header.tex`, `mcmc.ipynb` (which is selected), `mcmc.tex`, `p99-data.txt`, and `README.md`. The main area has two tabs: "Terminal 1" and "mcmc.ipynb". The "mcmc.ipynb" tab contains the following Julia code:

```
[36]: chain = []

# initial parameter guess
param_phase = 0.
param_amp = 6.

# proposal jump sizes
jump_phase = 0.01
jump_amp = 0.1

# initial log-posterior.
logprob = sine_log_posterior(times, y, param_phase, param_amp)

# how many steps to take
nsteps = 1_000

for i in 1:nsteps

    # propose new parameter values
    param_phase_new = param_phase + randn() * jump_phase
    param_amp_new   = param_amp   + randn() * jump_amp

    # compute log-posterior at new parameters
    logprob_new = sine_log_posterior(times, y, param_phase_new, param_amp_new)

    if (exp(logprob_new - logprob) >= rand(Float64))
```

1 s_ 6 🏃 Julia 1.2.0 | Idle

Mode: Edit ✎ Ln 2, Col 17 mcmc.ipynb

The screenshot shows a Jupyter Notebook interface with a Julia kernel. The left sidebar displays a file tree with files like `mcmc.ipynb`, `header.tex`, and `p99-data.txt`. The main area has two tabs: "Terminal 1" and "mcmc.ipynb". The "mcmc.ipynb" tab contains the following Julia code:

```
nsteps = 1_000

for i in 1:nsteps

    # propose new parameter values
    param_phase_new = param_phase + randn() * jump_phase
    param_amp_new   = param_amp   + randn() * jump_amp

    # compute log-posterior at new parameters
    logprob_new = sine_log_posterior(times, y, param_phase_new, param_amp_new)

    if (exp(logprob_new - logprob) >= rand(Float64))
        logprob = logprob_new
        param_phase = param_phase_new
        param_amp = param_amp_new
    end
    append!(chain, (param_phase, param_amp))
end
# append! makes "chain" a 1-d vector; reshape to a matrix
chain = reshape(chain, (2,Int64(length(chain)/2)))';
```

Below the code, the output of cell [37] is shown:

```
[37]: chain
```

```
[37]: 1000×2 LinearAlgebra.Adjoint{Any,Array{Any,2}}:
 -0.00109936  5.98017
 -0.0115605   6.00415
 -0.00408136  5.93793
```

The status bar at the bottom indicates "Julia 1.2.0 | Idle".

File Edit View Run Kernel Tabs Settings Help

Terminal 1 mcmc.ipynb

Name

- header.tex
- mcmc.ipynb
- mcmc.tex
- p99-data.txt
- README.md

```
logprob = sine_log_posterior(times, y, param_phase, param_amp)

# how many steps to take
nsteps = 10_000

for i in 1:nsteps

    # propose new parameter values
    param_phase_new = param_phase + randn() * jump_phase
    param_amp_new   = param_amp   + randn() * jump_amp

    # compute log-posterior at new parameters
    logprob_new = sine_log_posterior(times, y, param_phase_new, param_amp_new)

    if (exp(logprob_new - logprob) >= rand(Float64))
        logprob = logprob_new
        param_phase = param_phase_new
        param_amp = param_amp_new
    end
    append!(chain, (param_phase, param_amp))
end
# append! makes "chain" a 1-d vector; reshape to a matrix
chain = reshape(chain, (2,Int64(length(chain)/2)))';

UndefVarError: sine_log_posterior not defined

Stacktrace:
 [1] top-level scope at Tn[12]:10
```

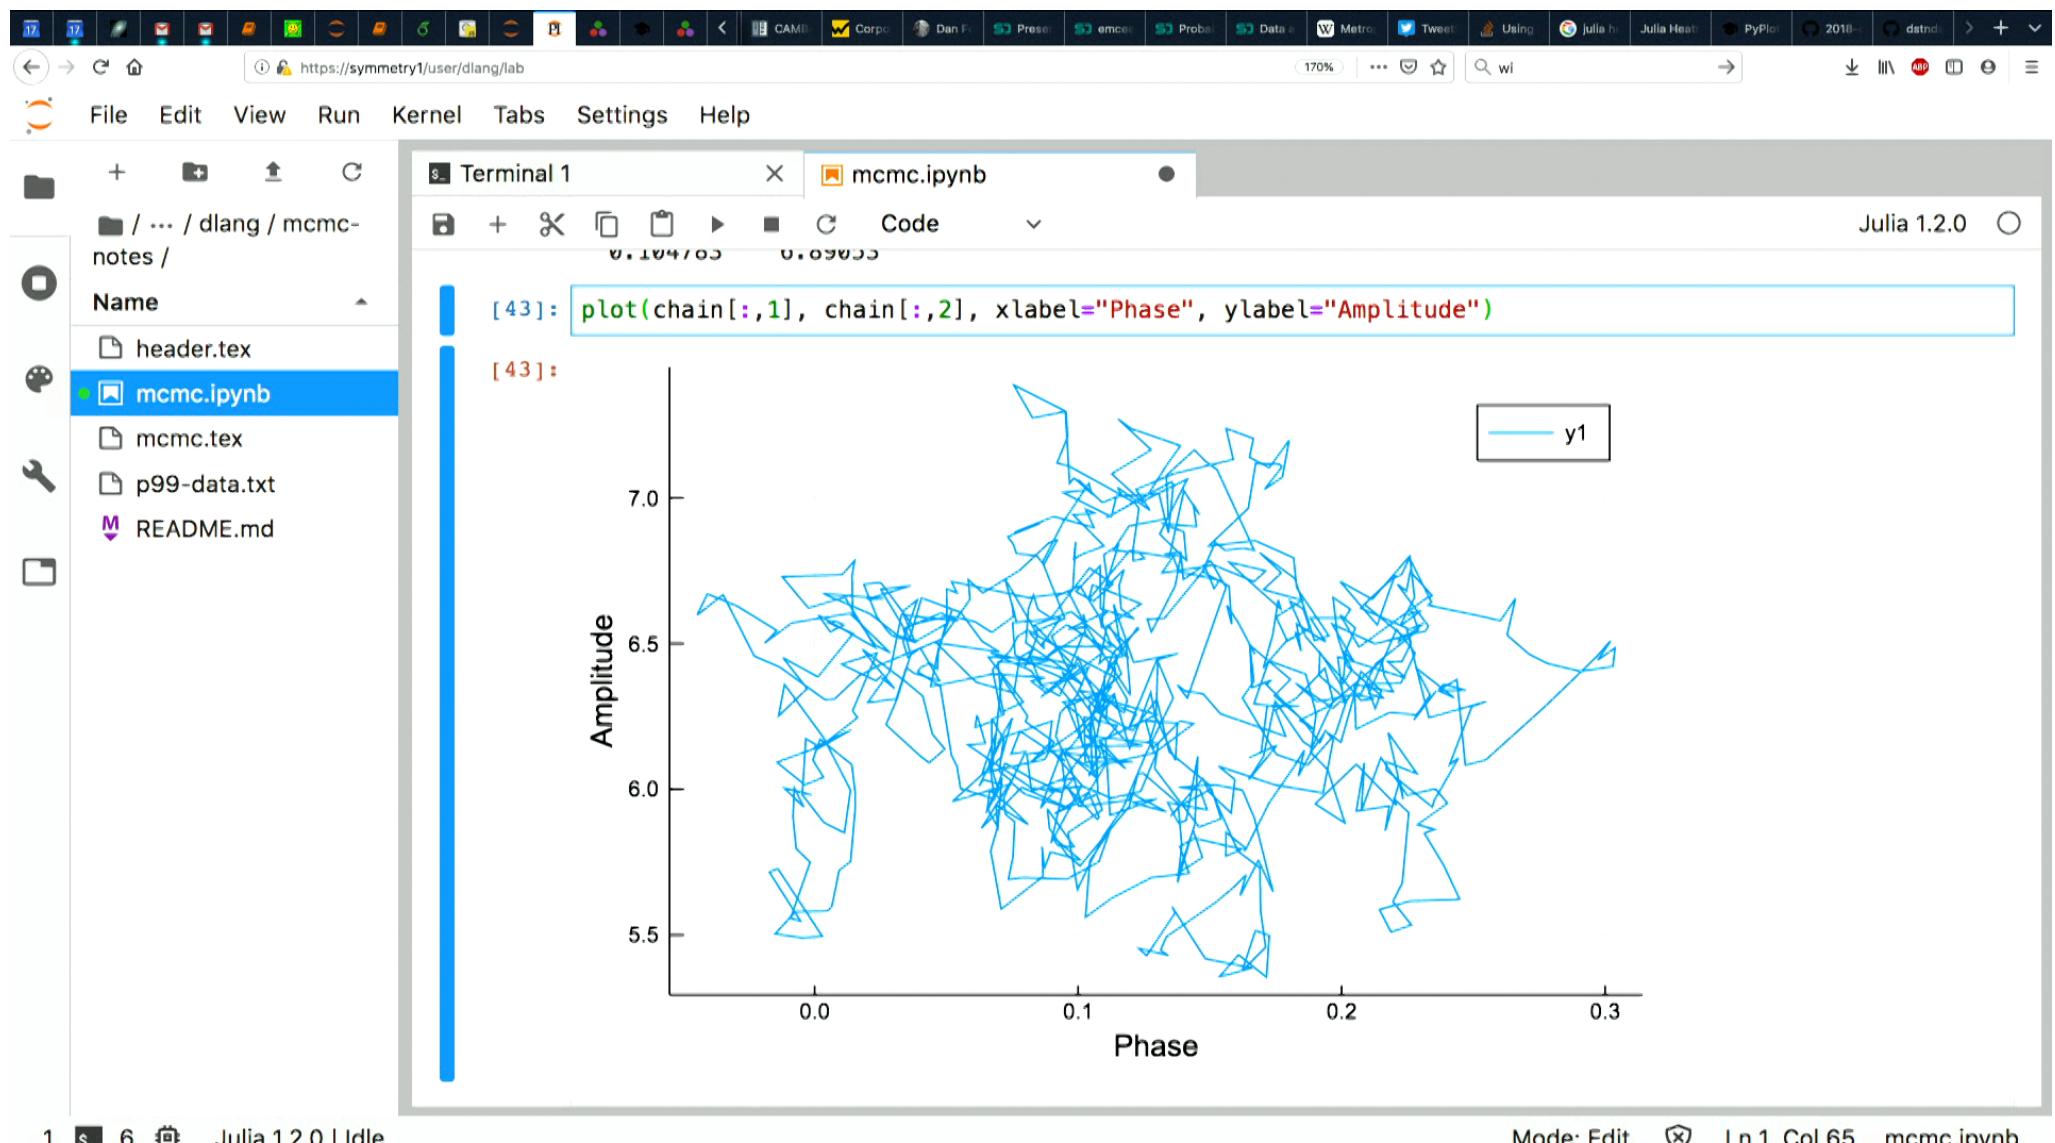
Mode: Command Ln 1, Col 1 mcmc.ipynb

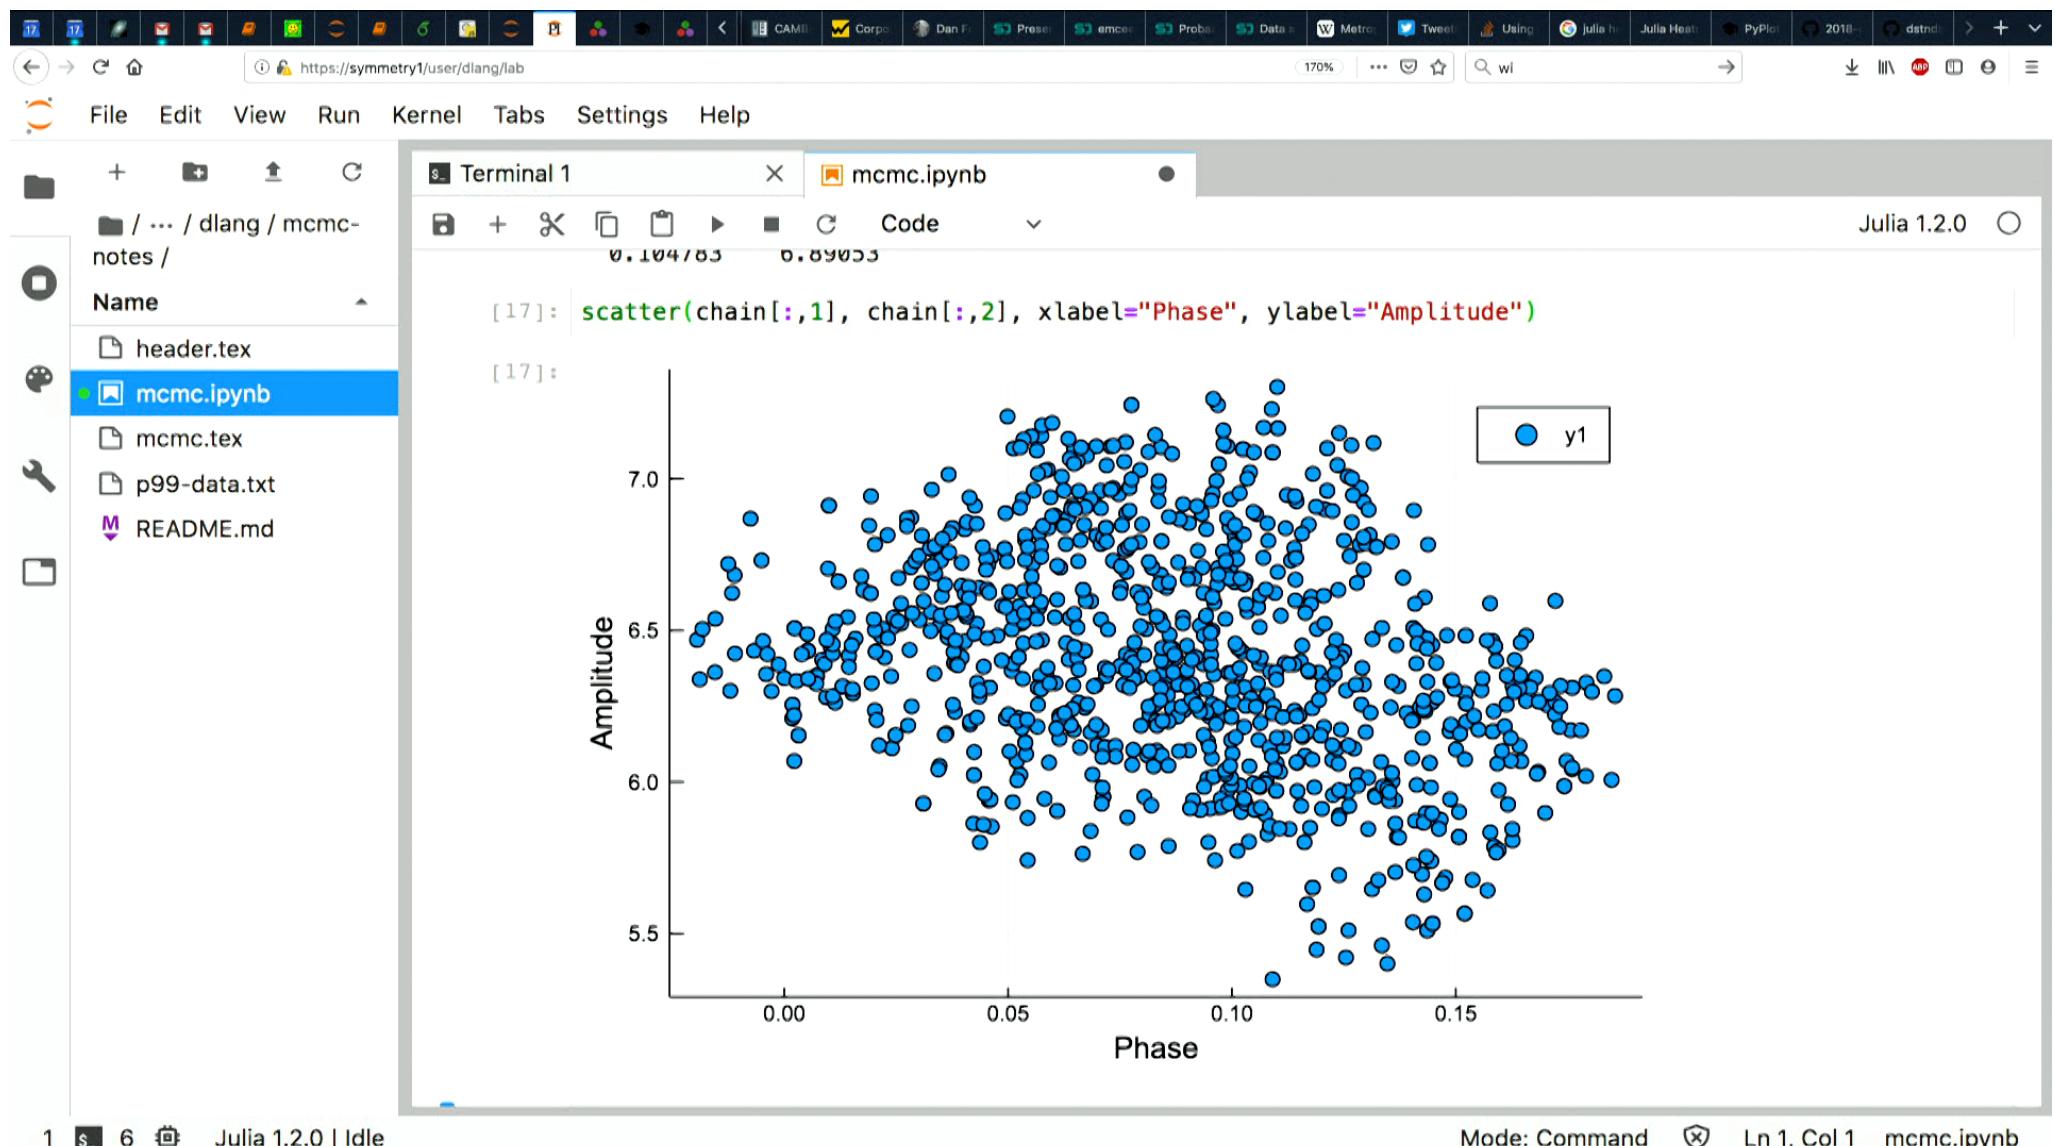
The screenshot shows a Jupyter Notebook interface with a Julia kernel. The left sidebar displays a file tree with files like header.tex, mcmc.ipynb (selected), mcmc.tex, p99-data.txt, and README.md. The main area has two tabs: Terminal 1 and mcmc.ipynb. The mcmc.ipynb tab contains a code cell output labeled [15] showing a 1000x2 array named chain. The array consists of pairs of floating-point numbers.

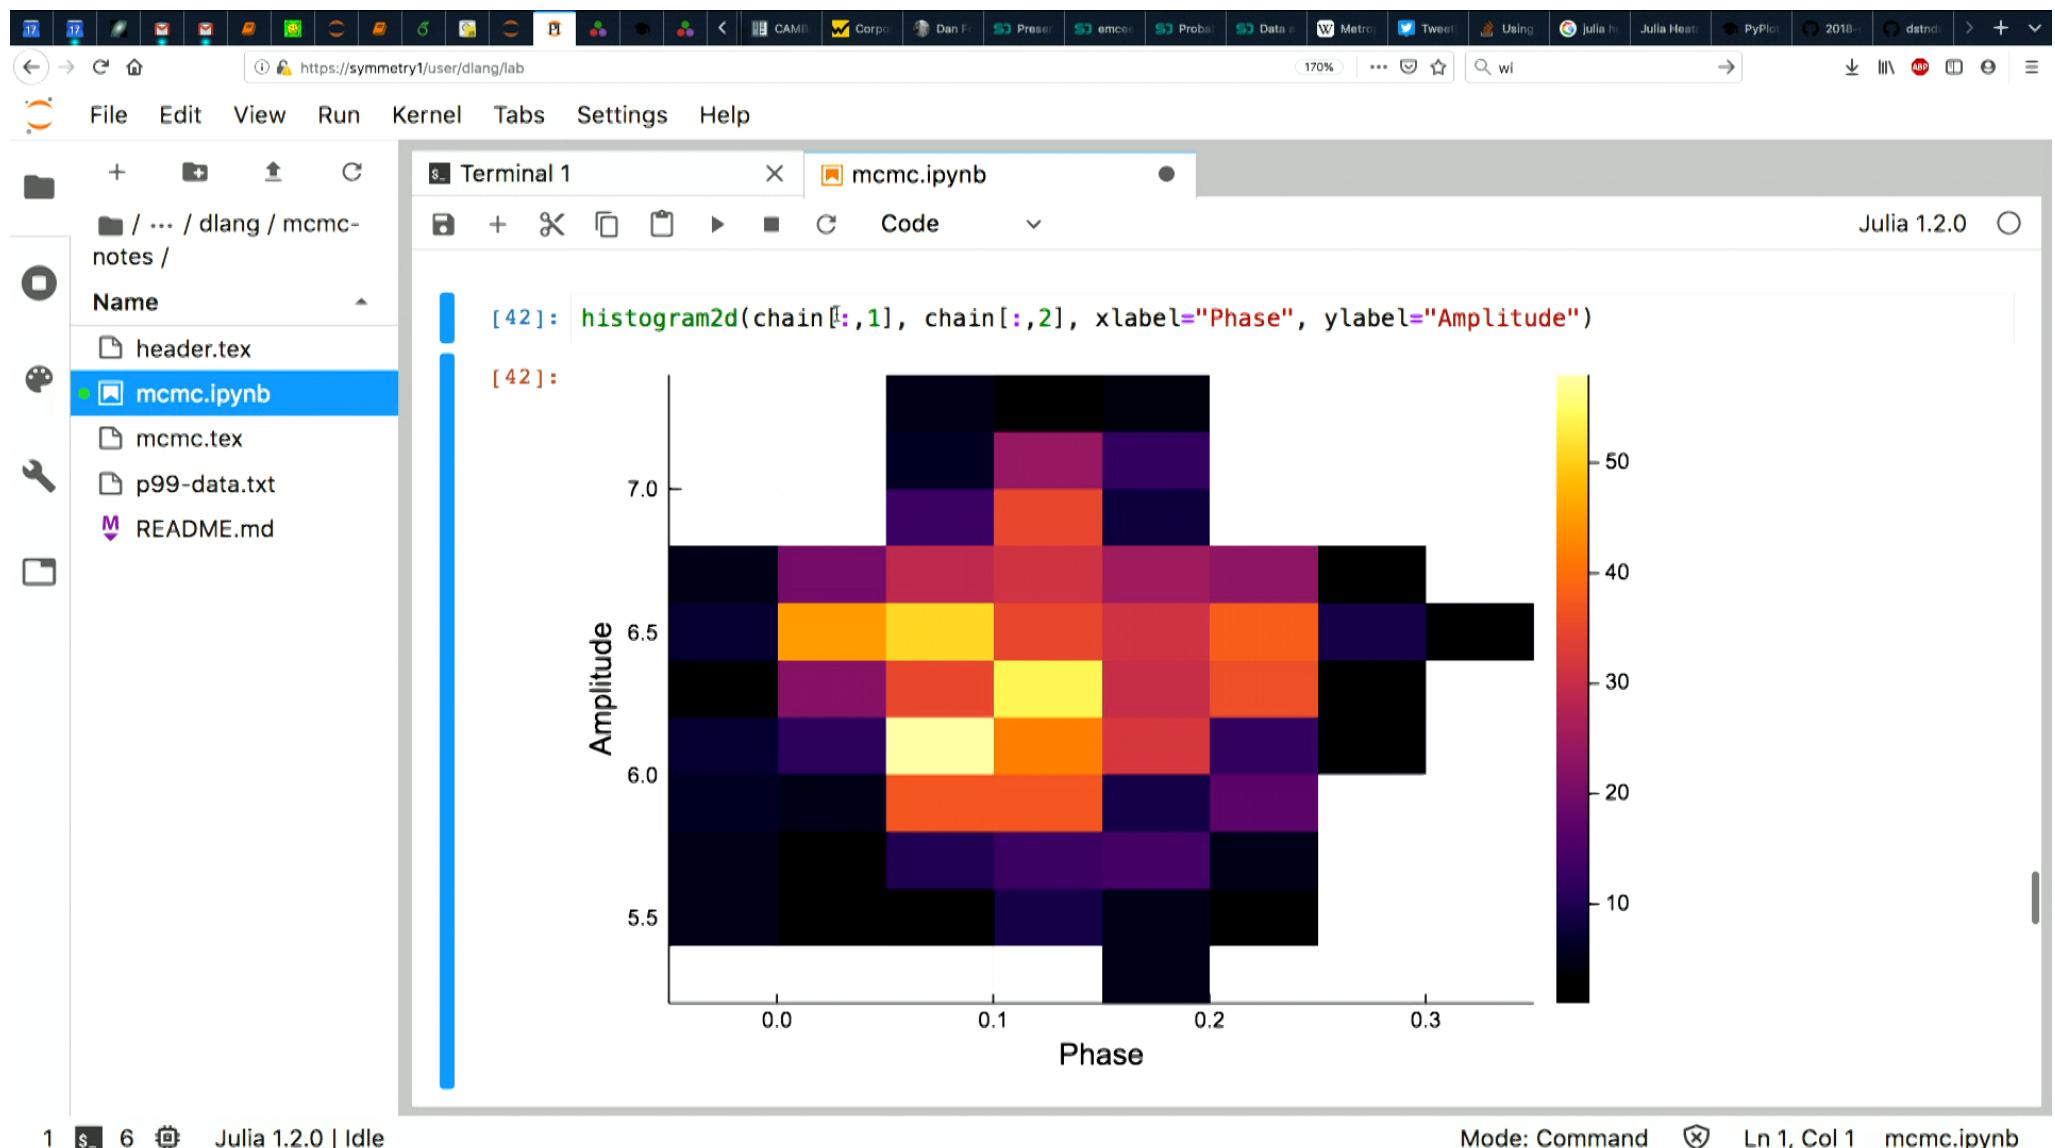
	chain
0.00224217	6.06881
0.00320339	6.15366
0.00320339	6.15366
0.00936336	6.27915
0.00179649	6.20945
0.0084974	6.40282
0.0202574	6.23622
0.00731598	6.35182
-0.00283468	6.29928
-0.00403954	6.35532
-0.0120522	6.30071
-0.0188619	6.33836
-0.0194753	6.46848
:	
0.0814683	6.72598
0.0768506	6.77446
0.0774634	6.78338
0.0859465	6.87445
0.100682	6.84697
0.107923	6.85274
0.129359	6.80687
0.130526	6.8978
0.135696	6.79232
0.1263	6.74453

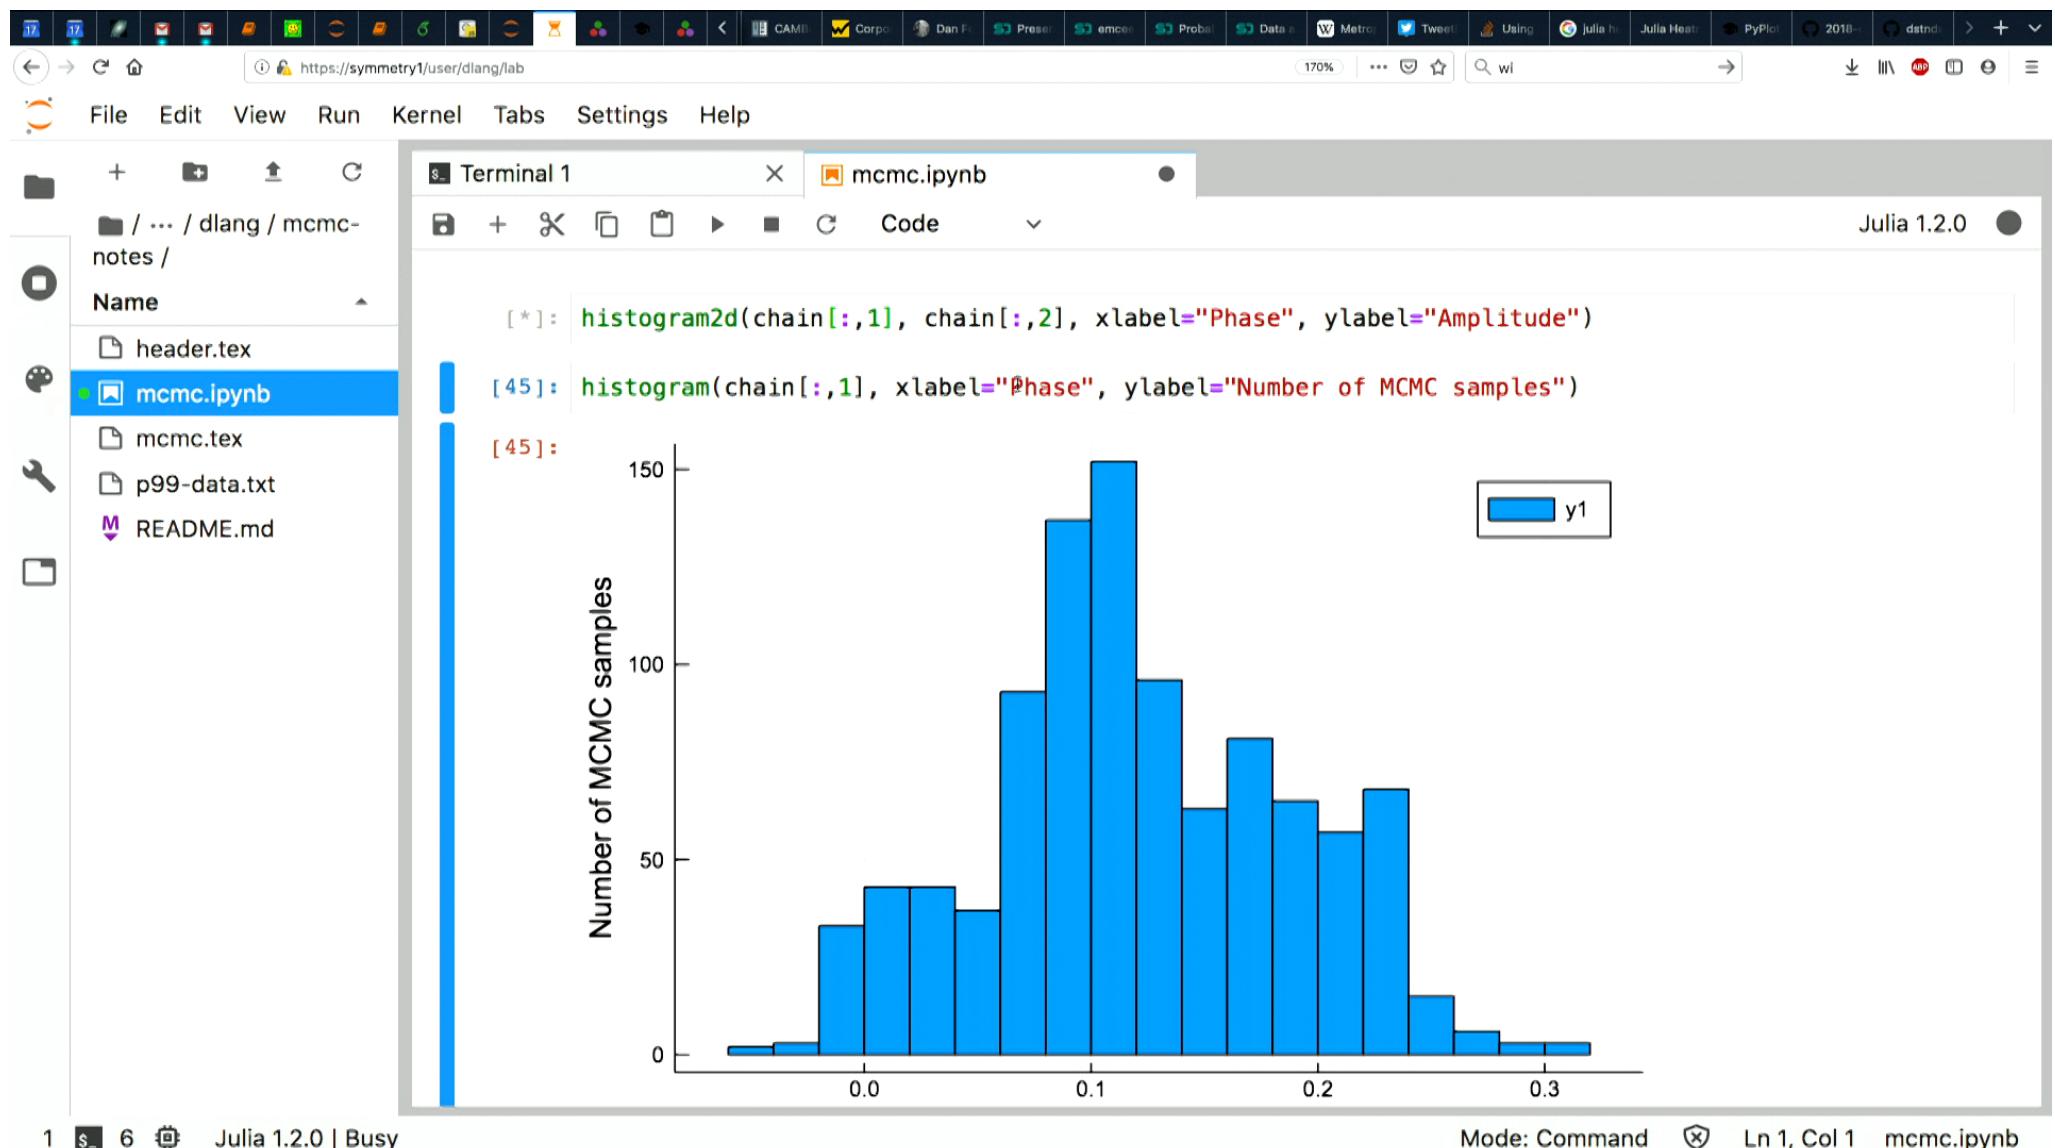
1 s_ 6 🏃 Julia 1.2.0 | Idle

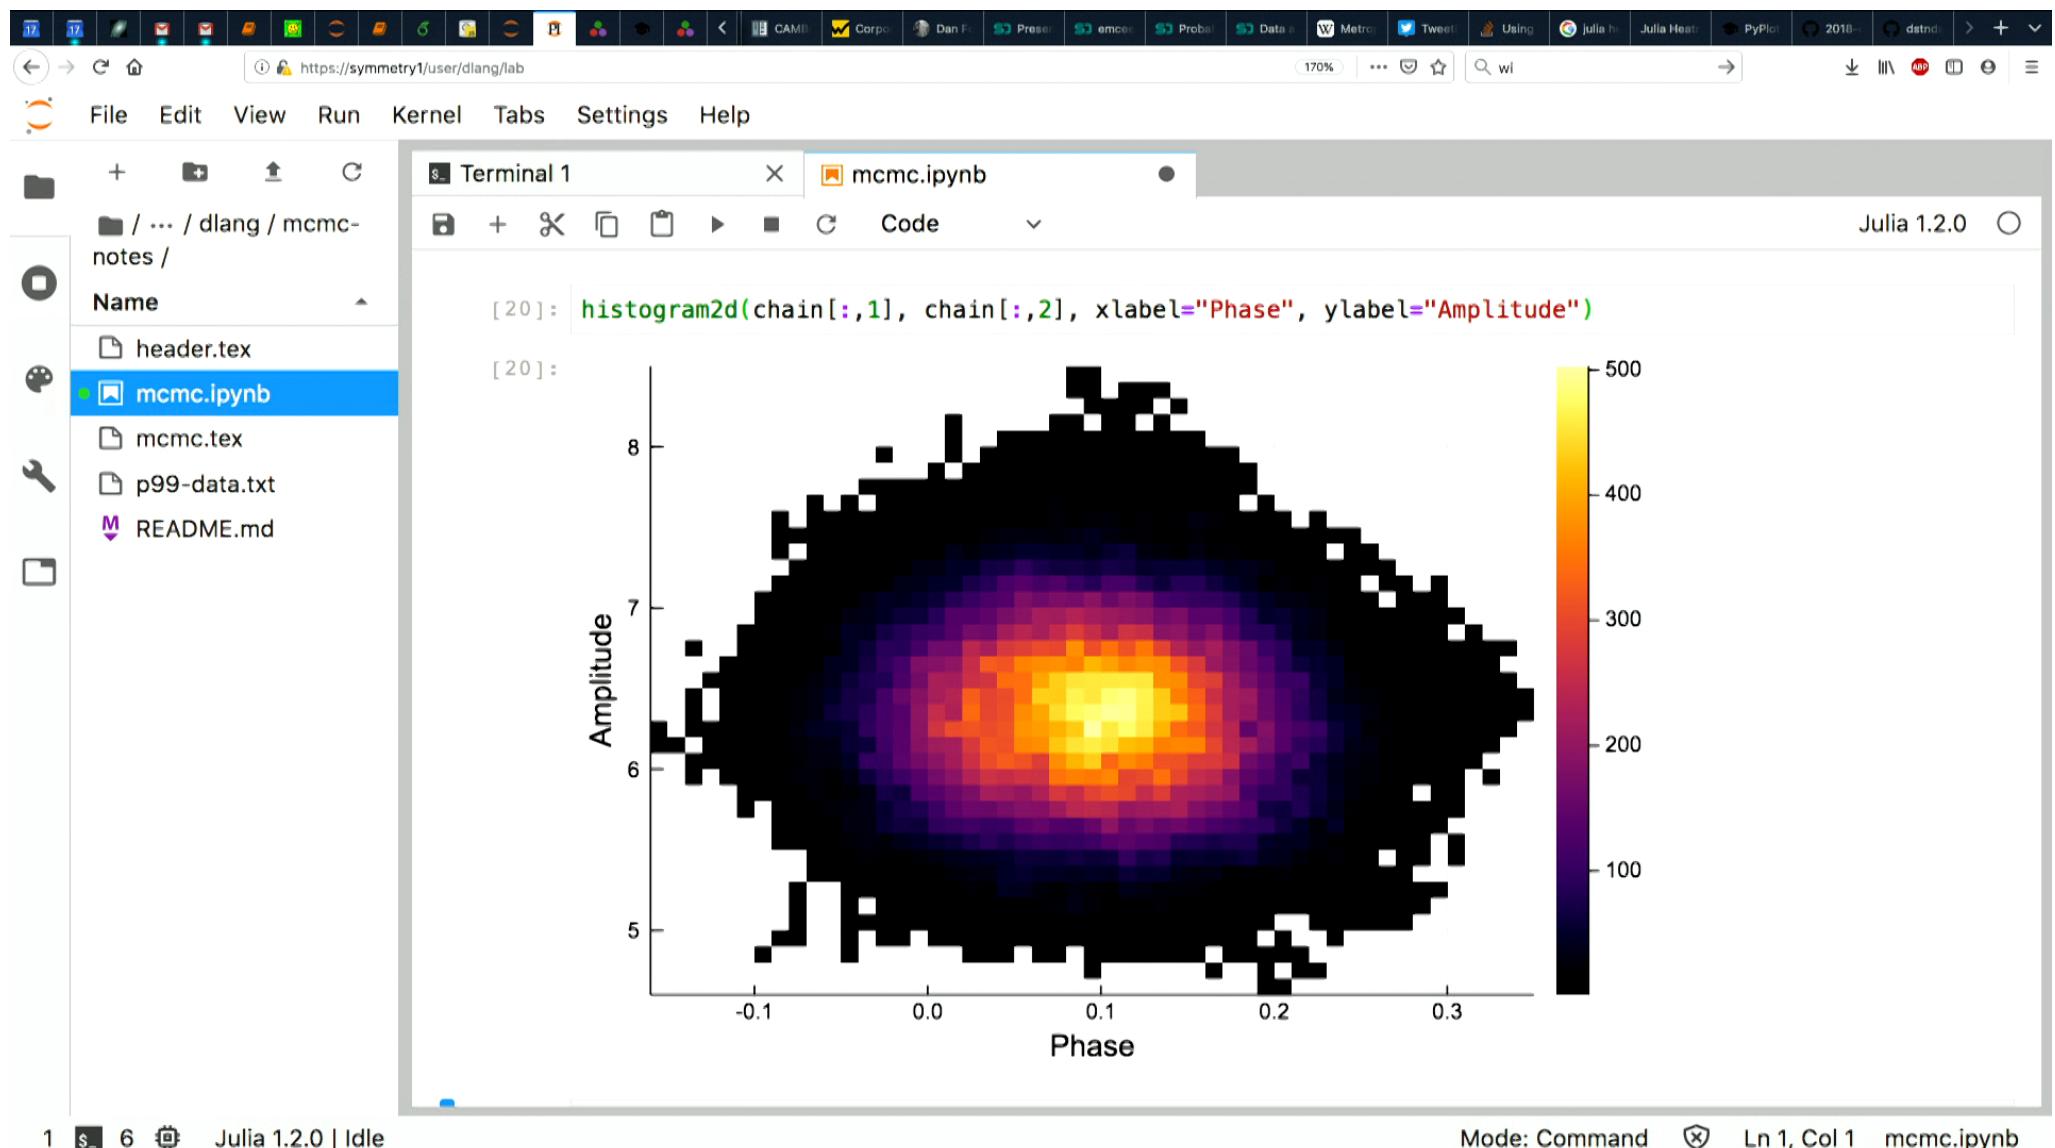
Mode: Command ✎ Ln 1, Col 1 mcmc.ipynb

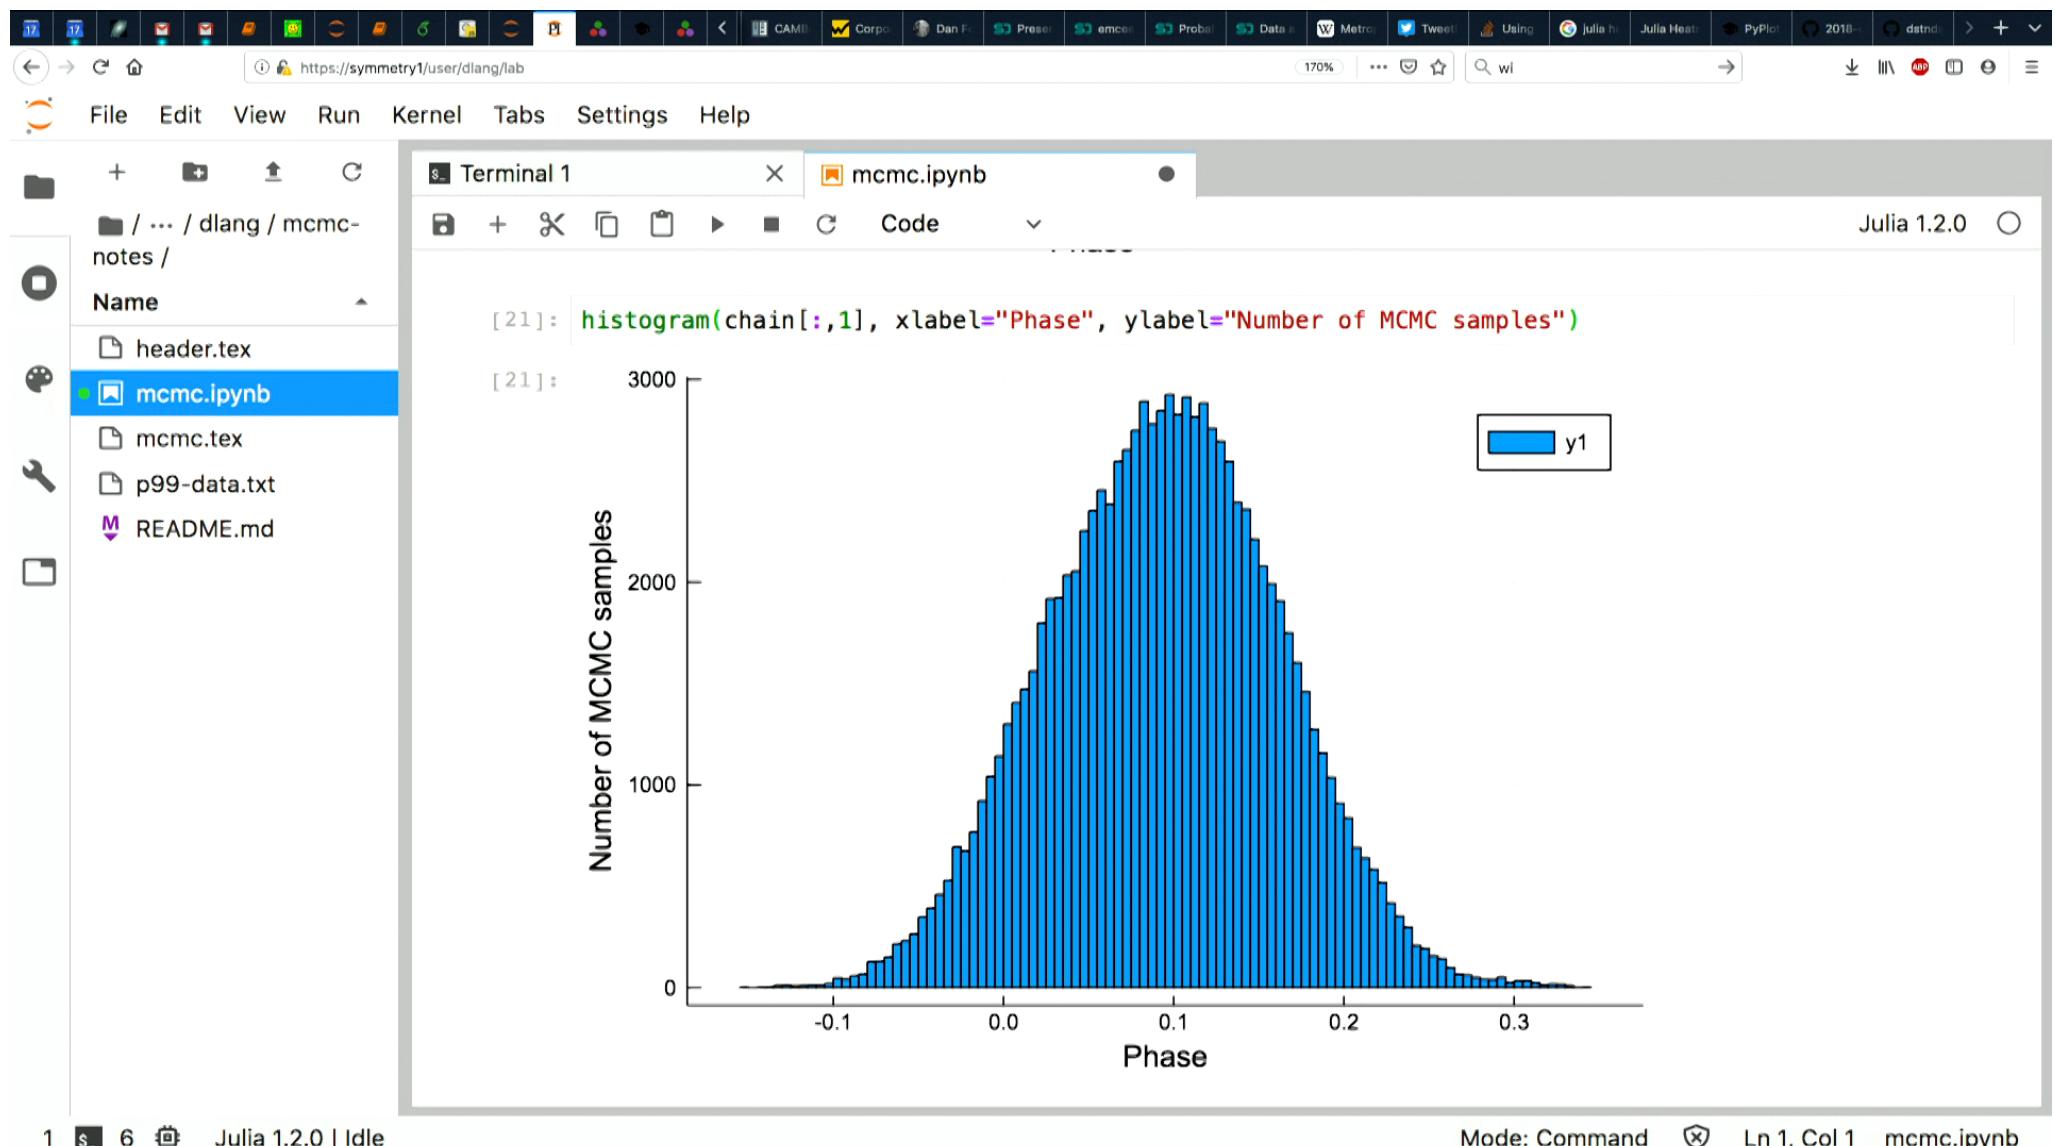


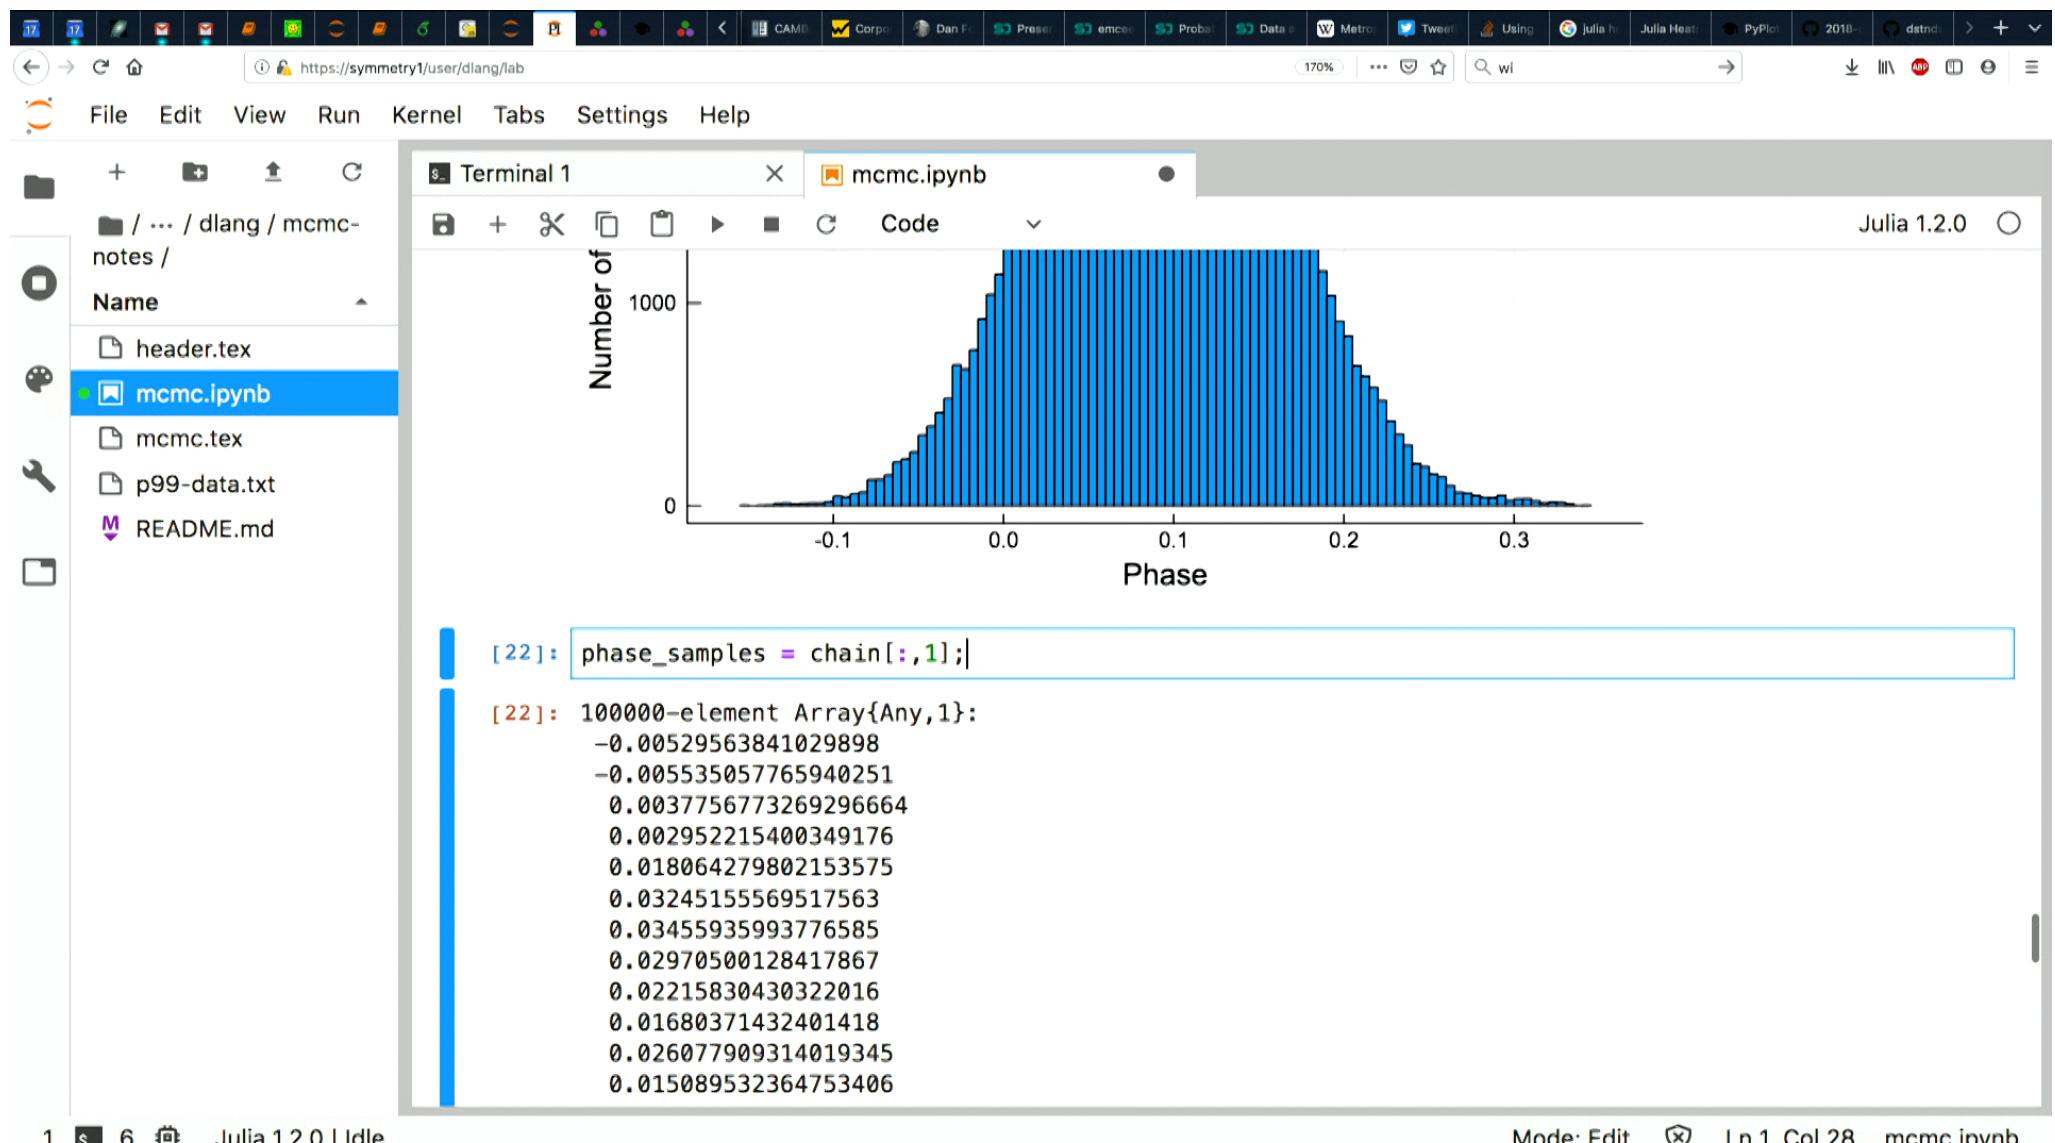


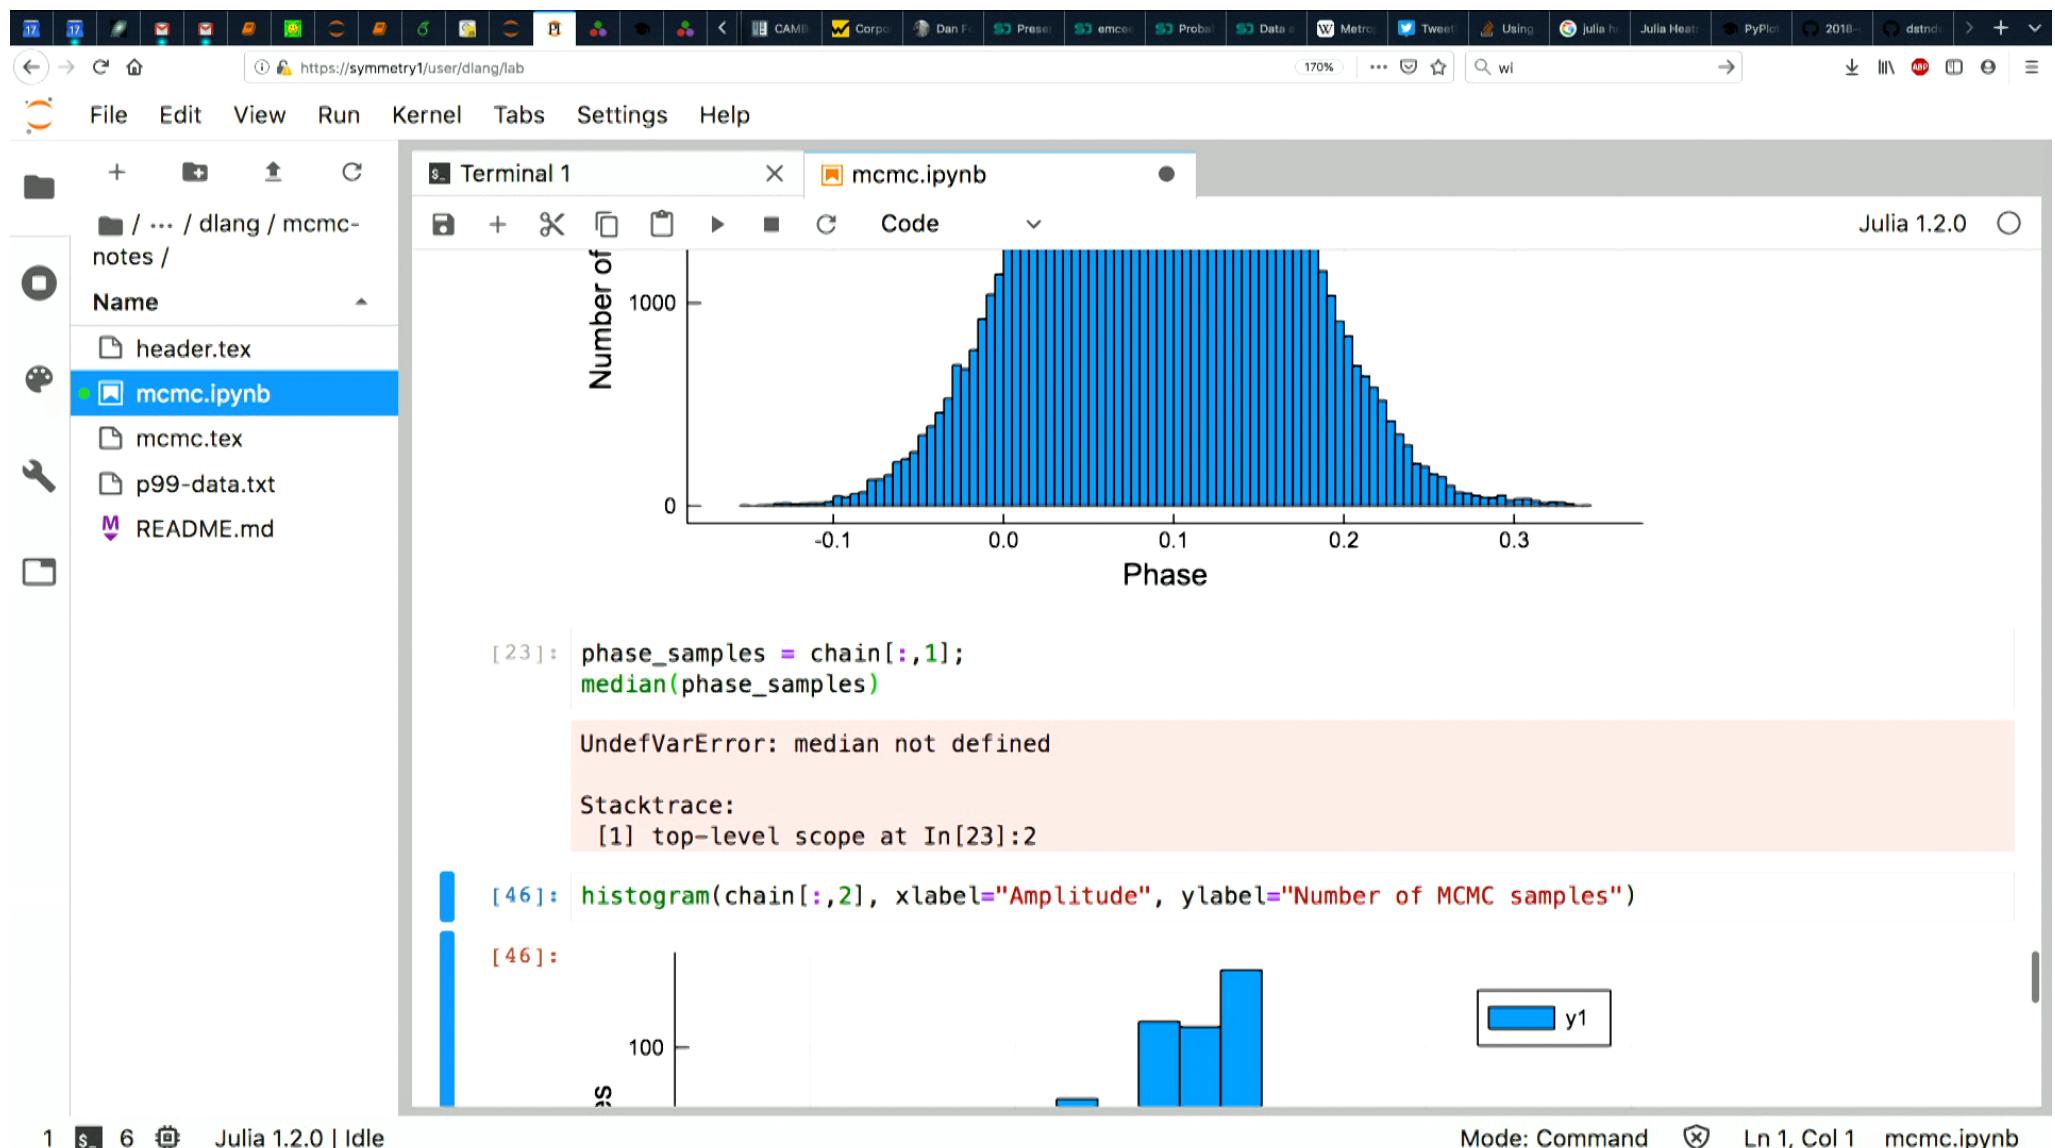


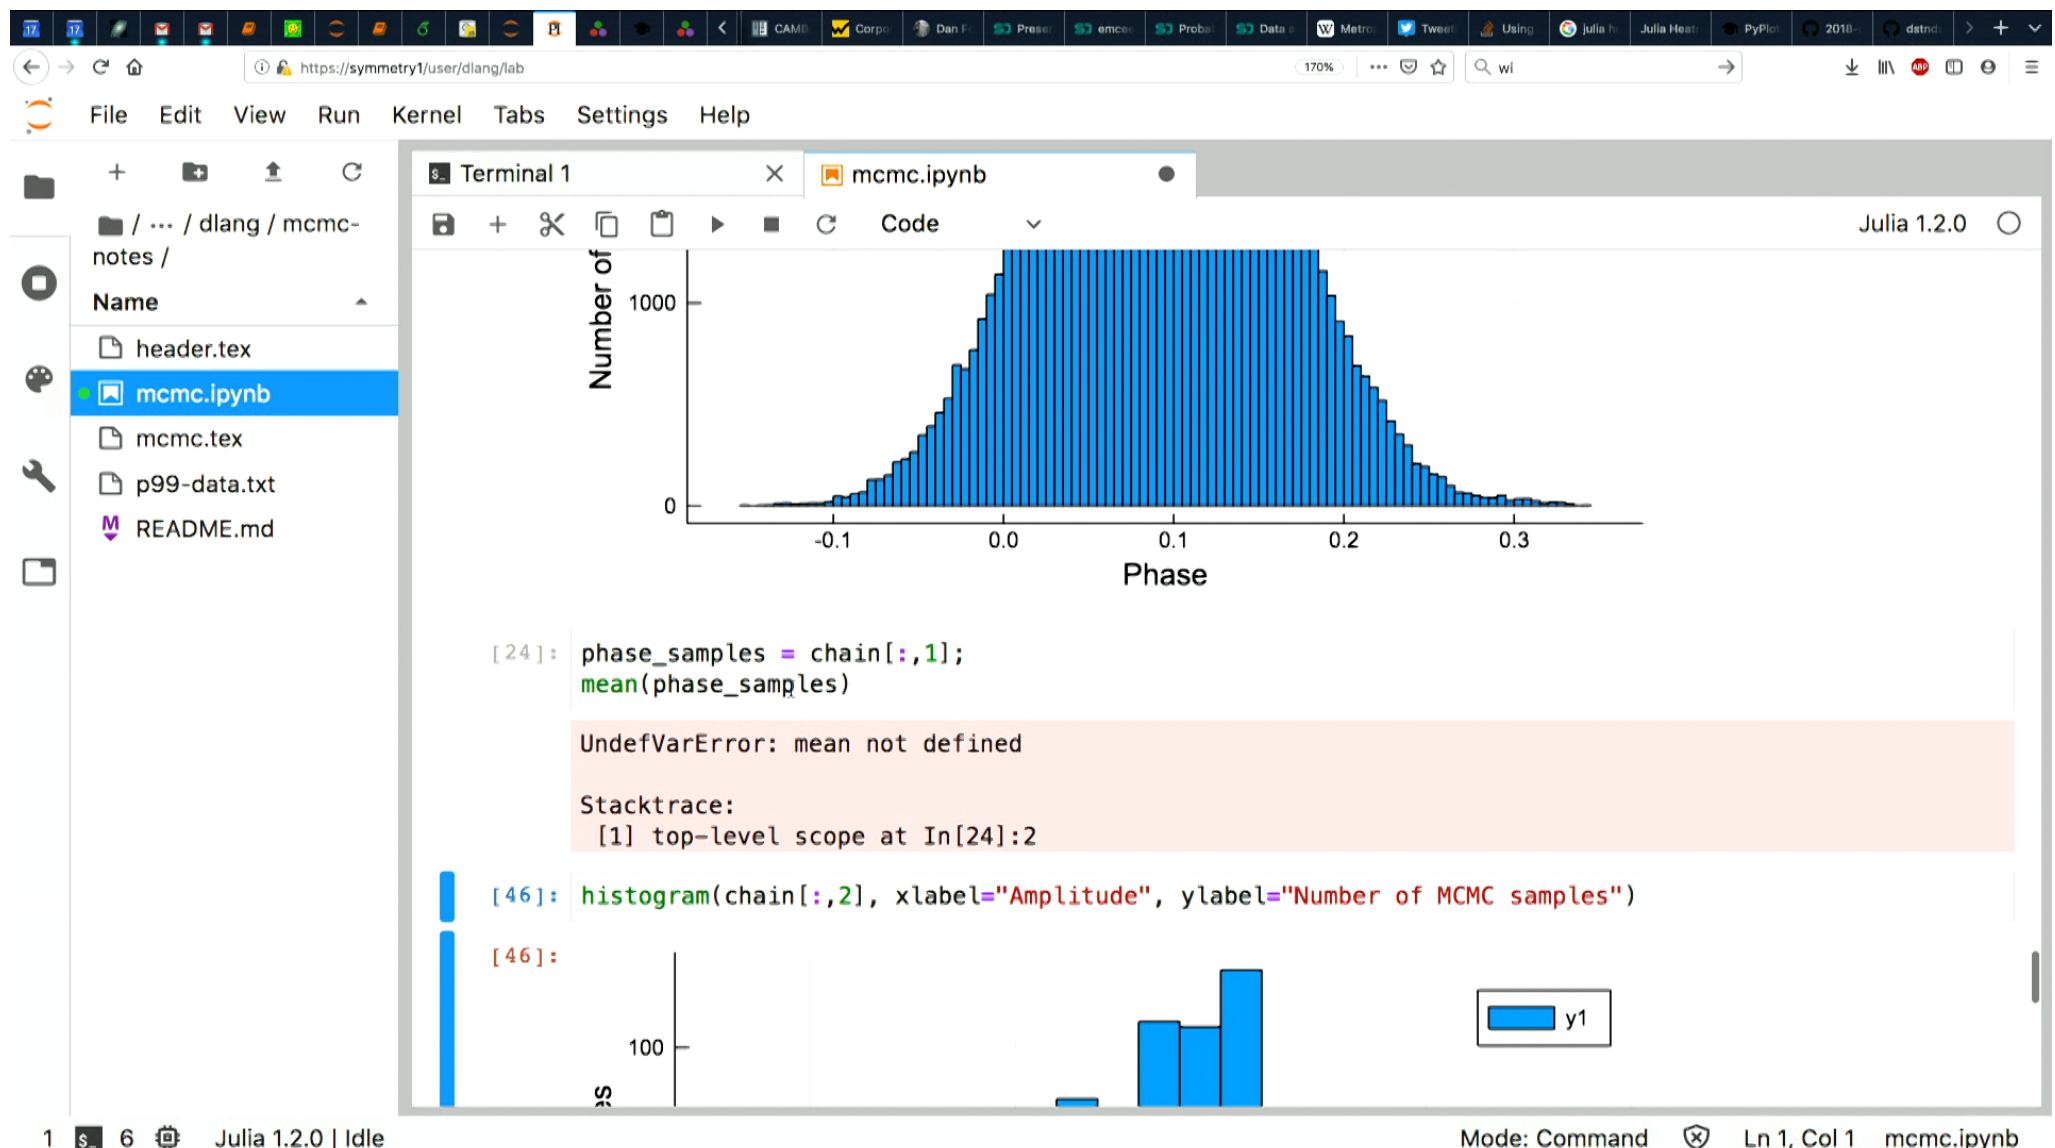


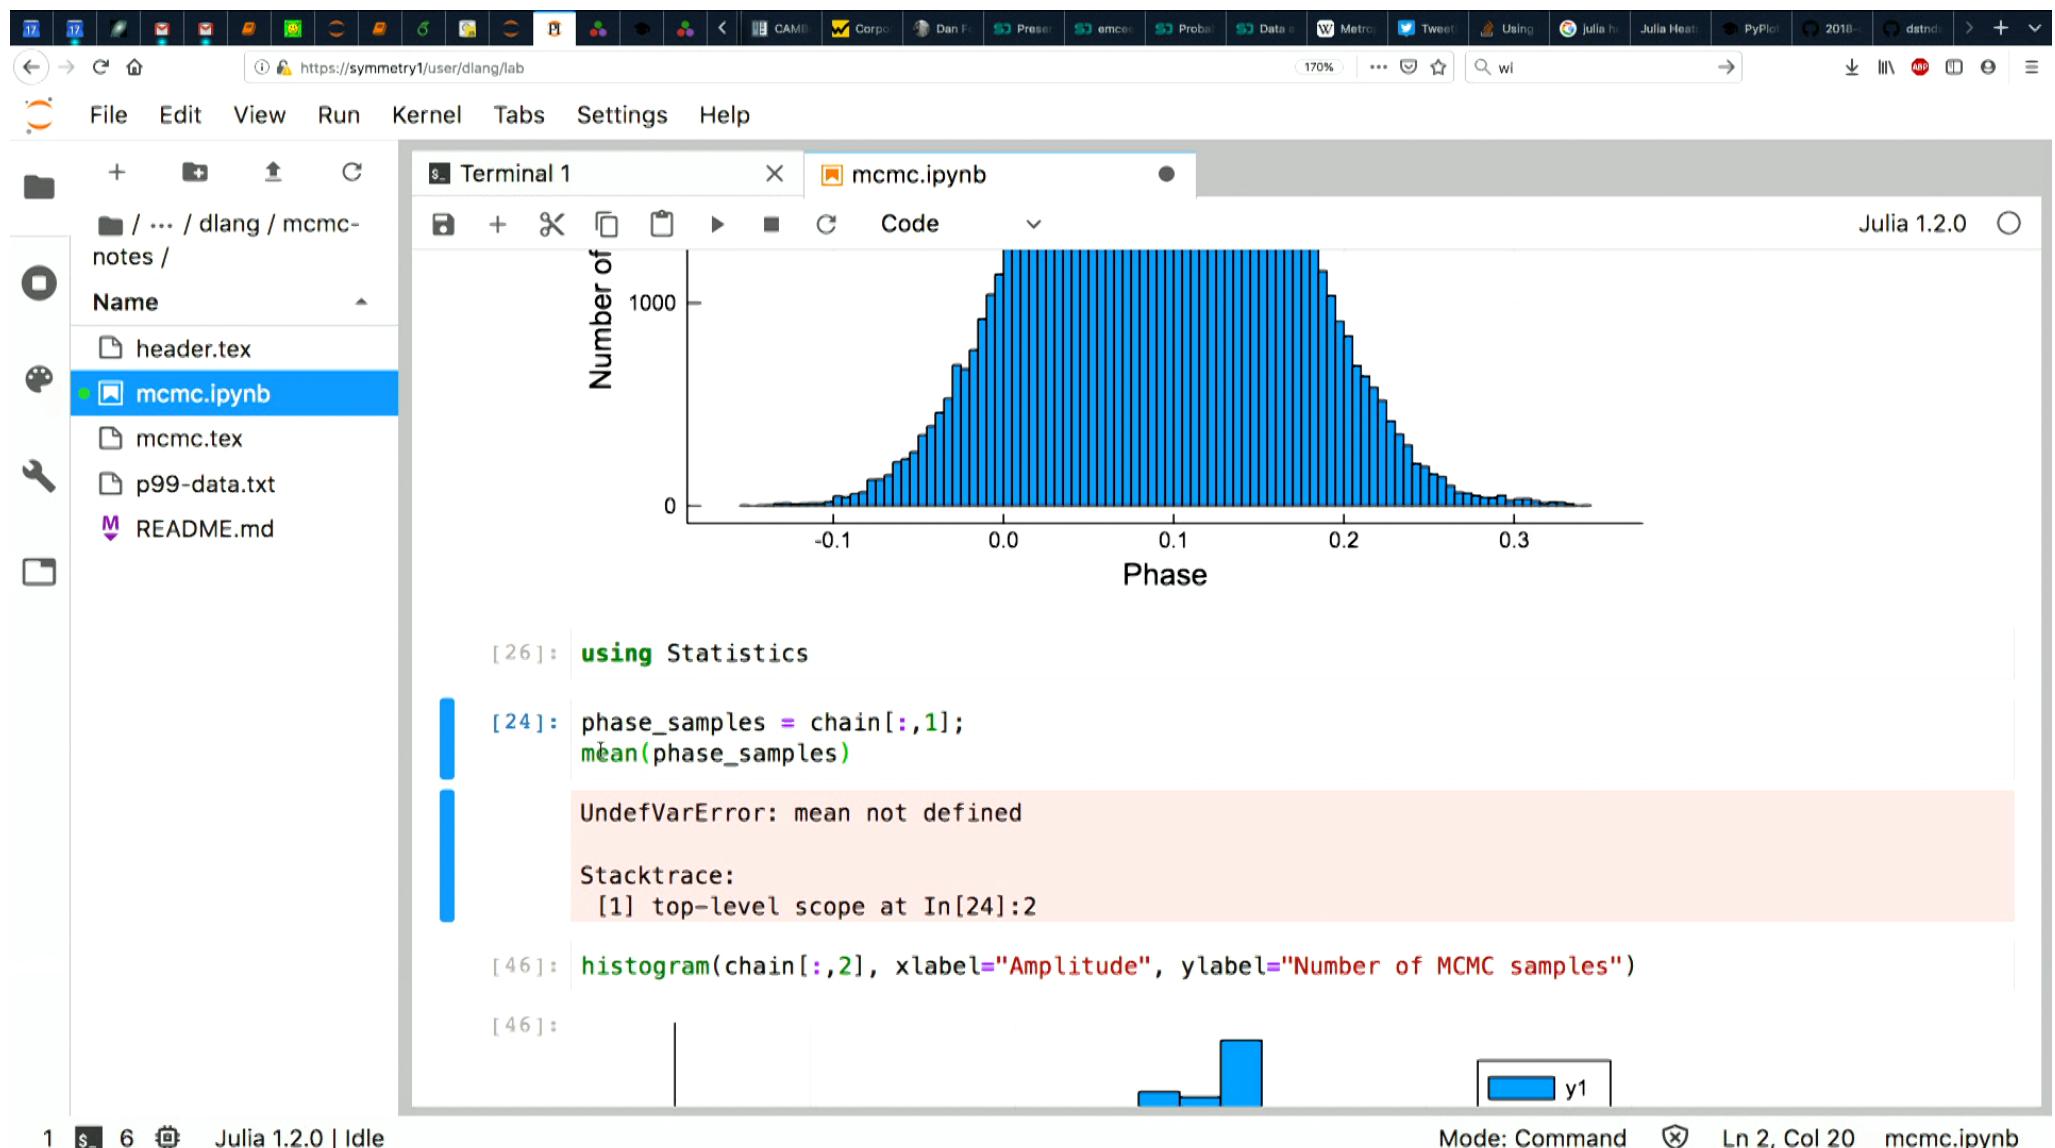


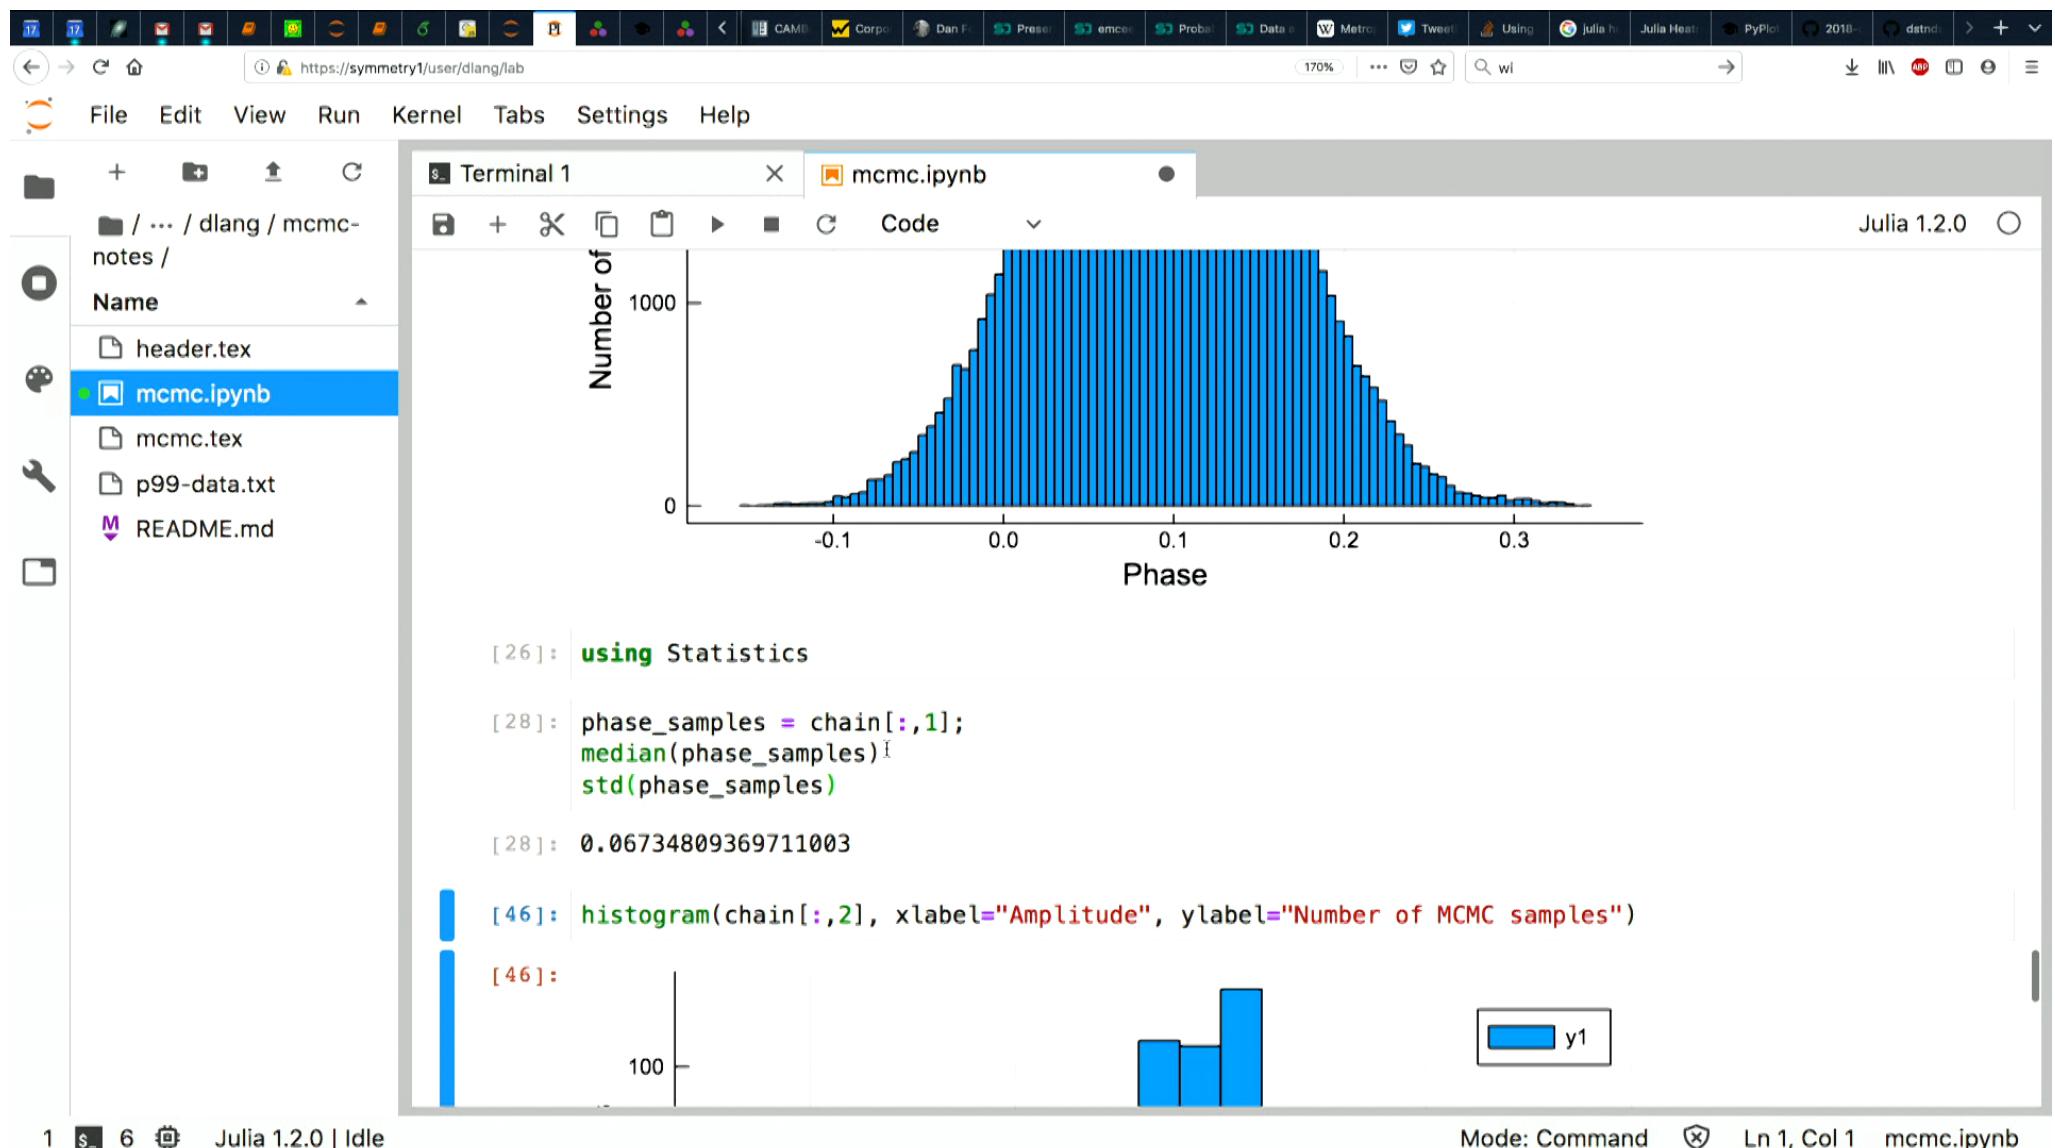


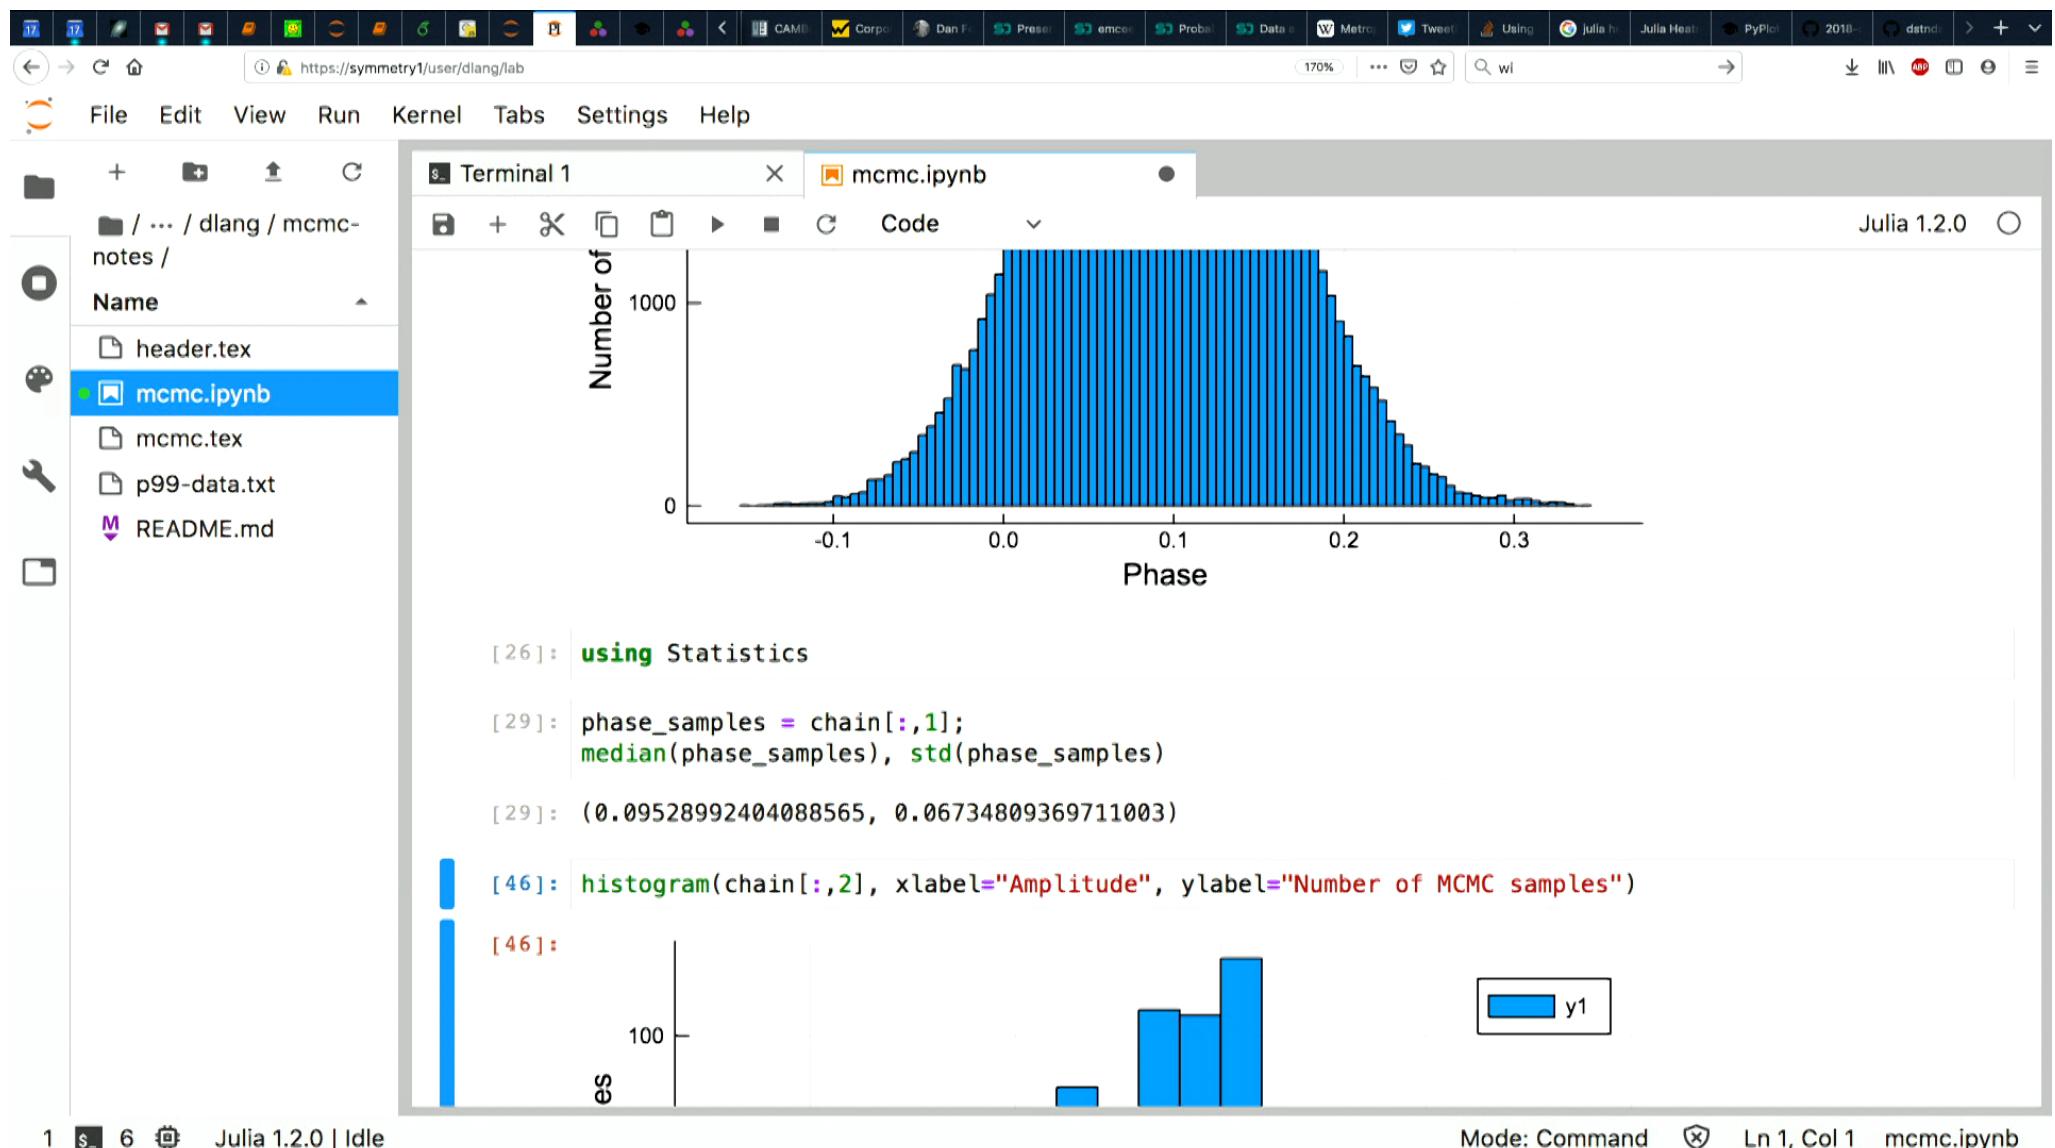








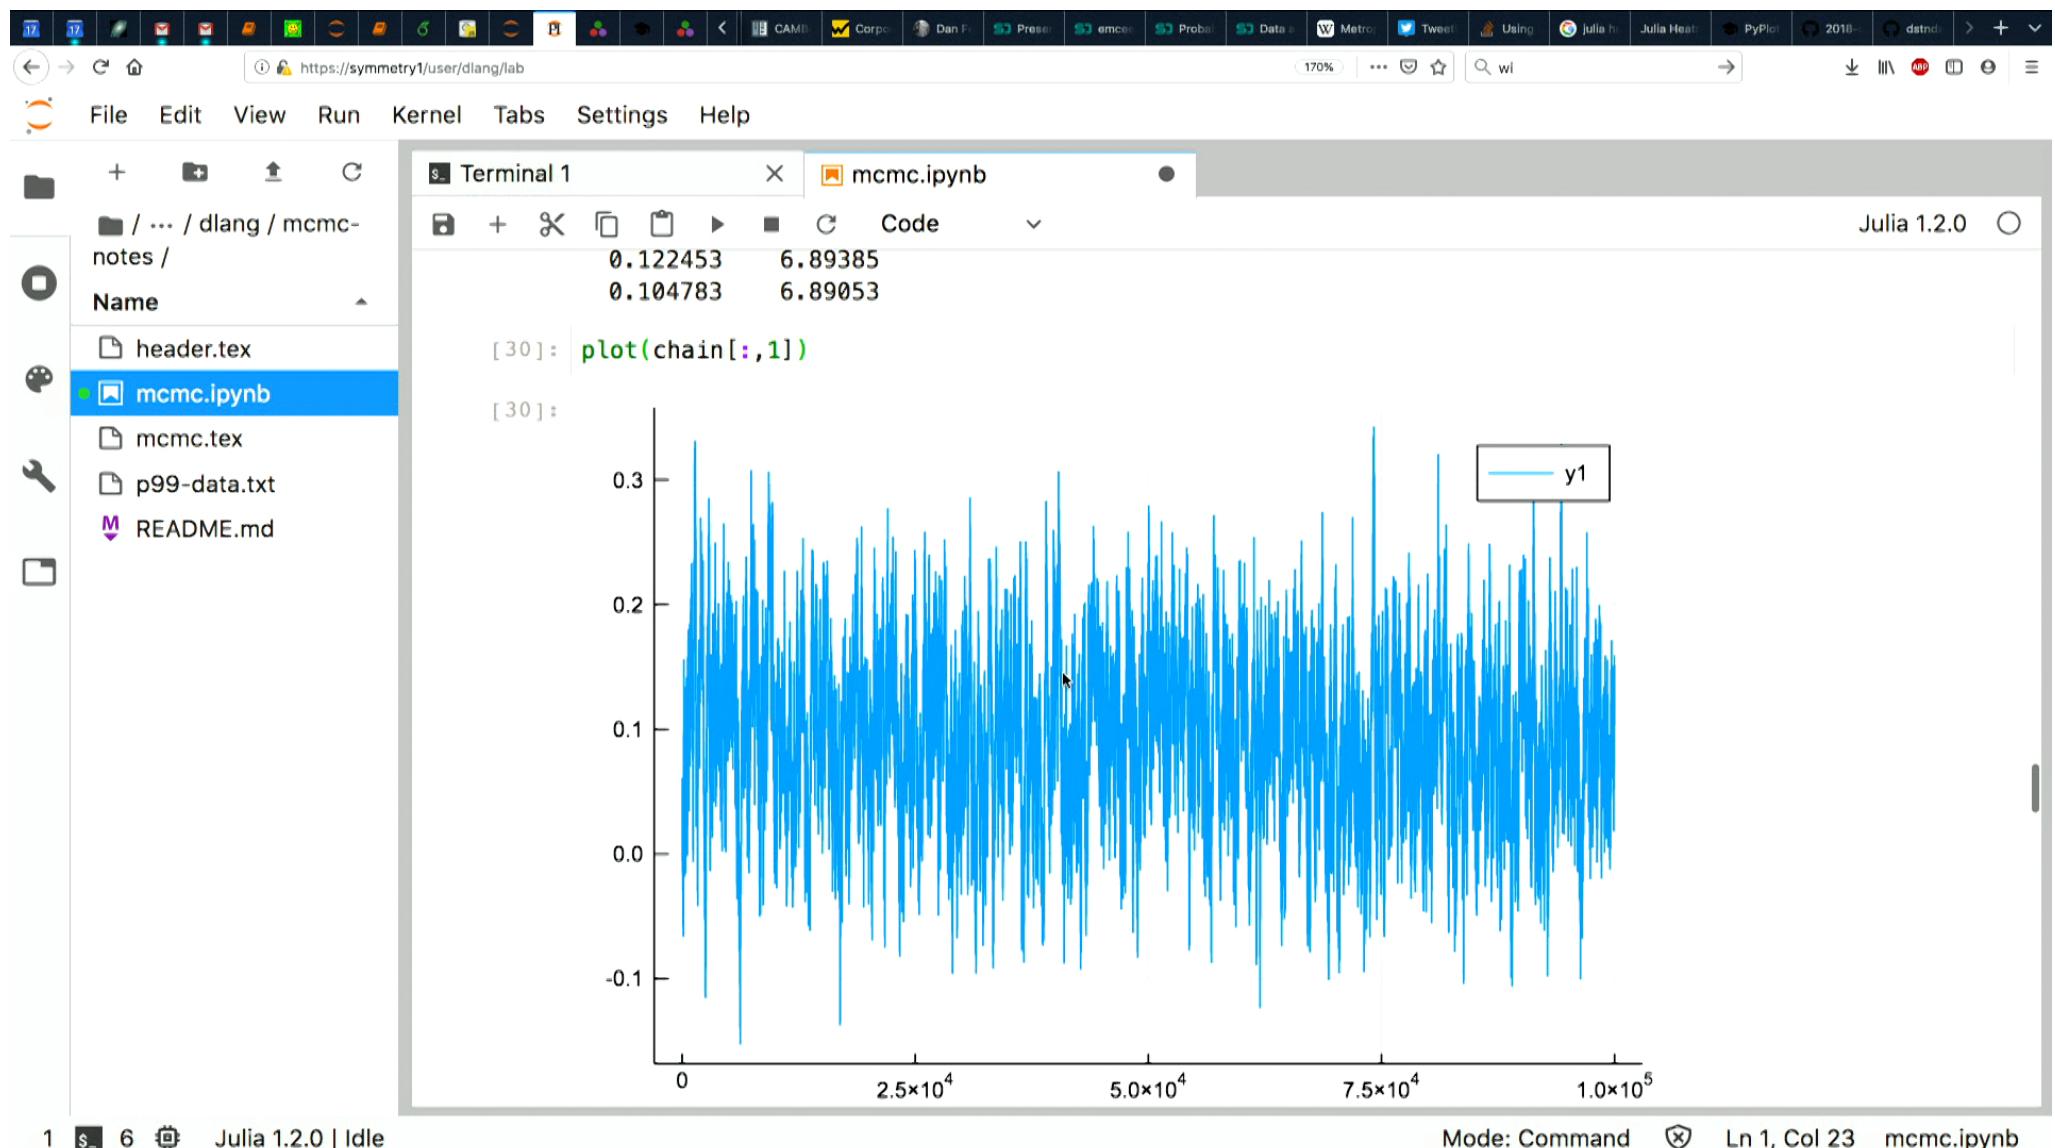


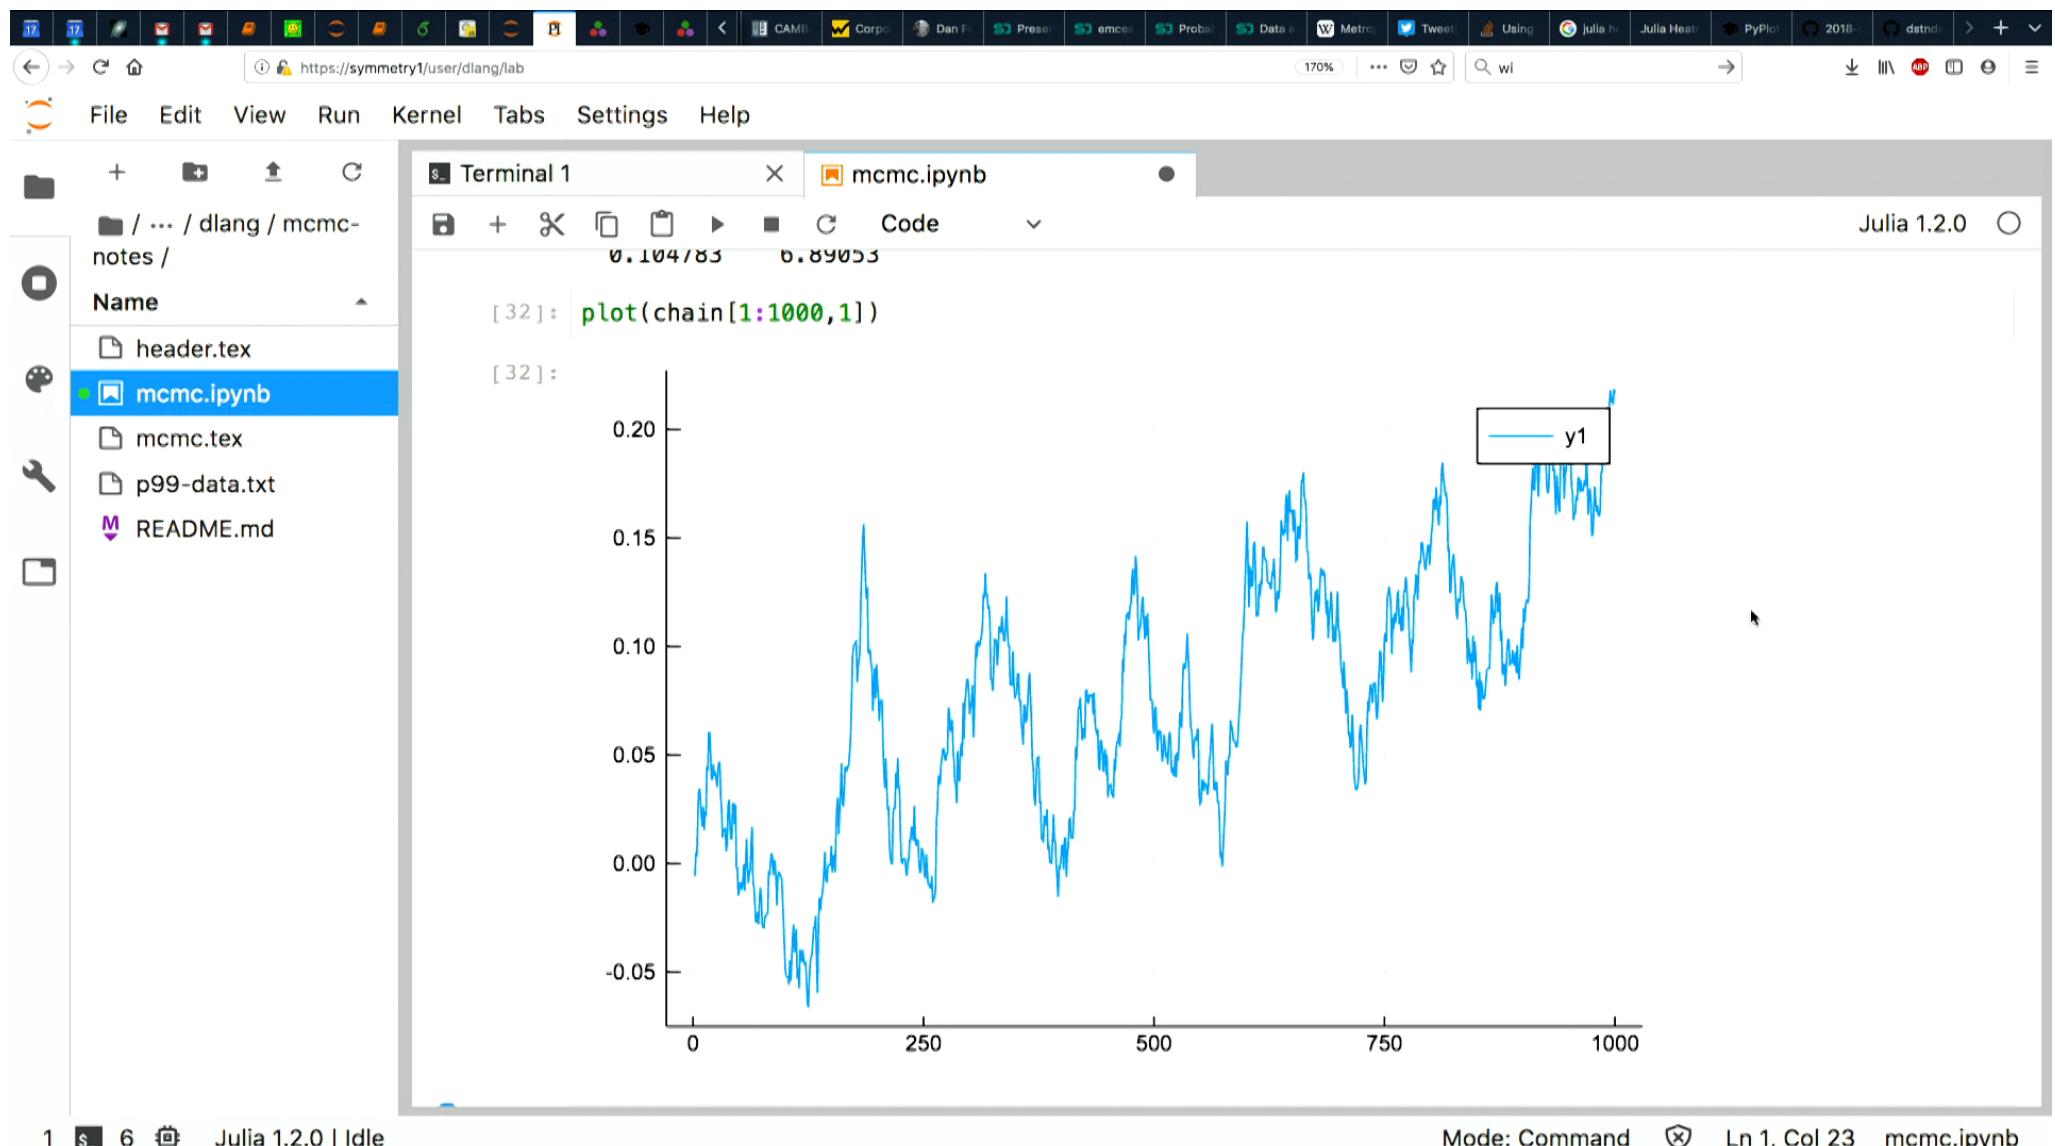


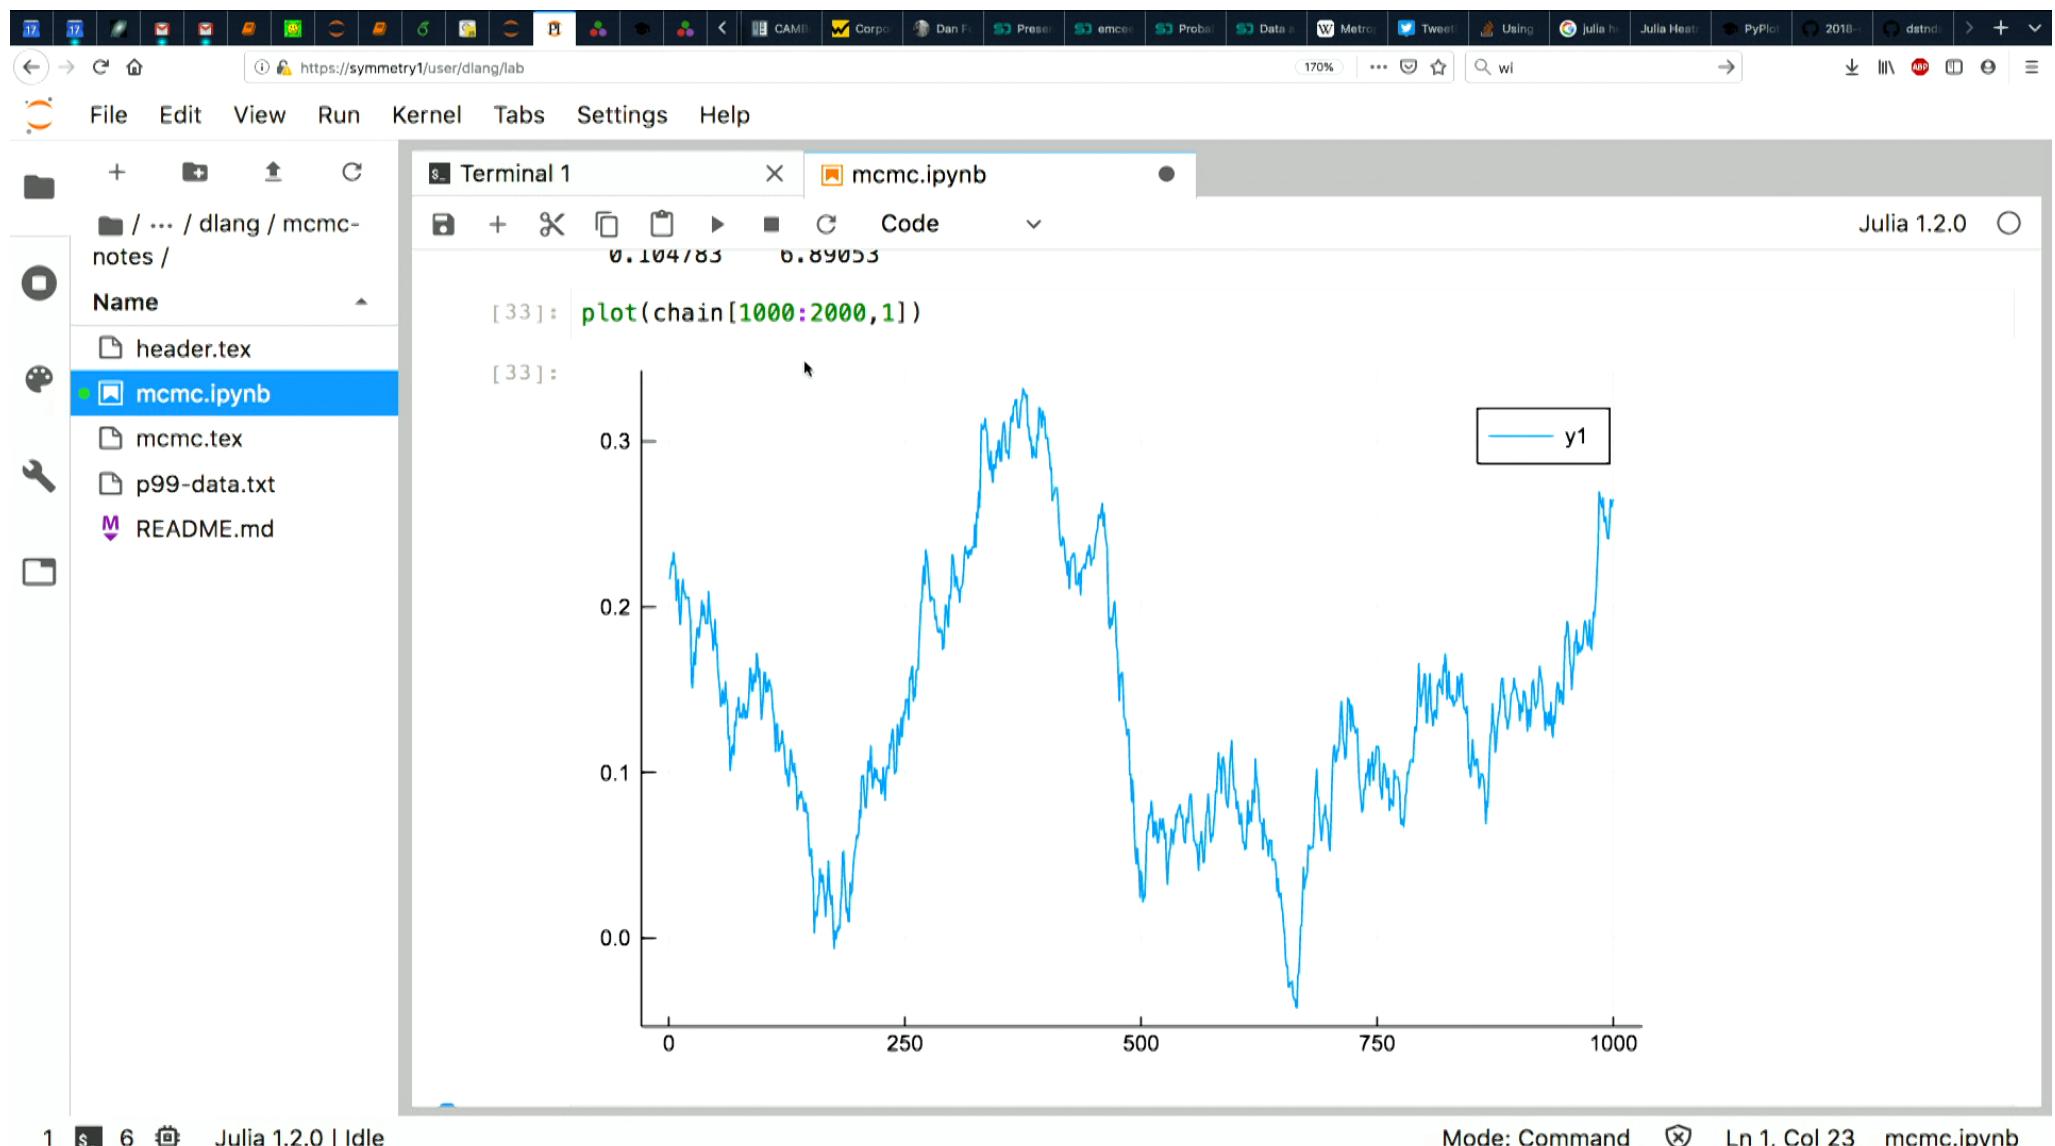
The screenshot shows a Jupyter Notebook interface with a Julia kernel. The left sidebar displays a file tree under the path `/ ... / dlang / mcmc-notes /`. The file `mcmc.ipynb` is selected and highlighted in blue. The main area contains a terminal window titled "Terminal 1" and a code editor window titled "mcmc.ipynb". The terminal window shows the following Julia code and its output:

```
[15]: 1000x2 LinearAlgebra.Adjoint{Any,Array{Any,2}}:
 0.00224217 6.06881
 0.00320339 6.15366
 0.00320339 6.15366
 0.00936336 6.27915
 0.00179649 6.20945
 0.0084974 6.40282
 0.0202574 6.23622
 0.00731598 6.35182
 -0.00283468 6.29928
 -0.00403954 6.35532
 -0.0120522 6.30071
 -0.0188619 6.33836
 -0.0194753 6.46848
 :
 0.0814683 6.72598
 0.0768506 6.77446
 0.0774634 6.78338
 0.0859465 6.87445
 0.100682 6.84697
 0.107923 6.85274
 0.129359 6.80687
 0.130526 6.8978
 0.135696 6.79232
 0.1263 6.74453
 0.122453 6.89385
 0.104783 6.89053
```

The code editor window shows the same output. The status bar at the bottom indicates "Julia 1.2.0 | Idle" and "Mode: Edit".







The screenshot shows a Jupyter Notebook interface with a Julia kernel. The left sidebar displays a file tree under the path `/ ... / dlang / mcmc-notes /`. The current file, `mcmc.ipynb`, is selected and highlighted in blue. The main area contains a code cell with the following Julia code:

```
# proposal jump sizes
jump_phase = 0.01
jump_amp = 0.1

# initial log-posterior.
logprob = sine_log_posterior(times, y, param_phase, param_amp)

# how many steps to take
nsteps = 100_000

# how many times we accepted a move
haccepts = 0

for i in 1:nsteps

    # propose new parameter values
    param_phase_new = param_phase + randn() * jump_phase
    param_amp_new = param_amp + randn() * jump_amp

    # compute log-posterior at new parameters
    logprob_new = sine_log_posterior(times, y, param_phase_new, param_amp_new)

    if (exp(logprob_new - logprob) >= rand(Float64))
        logprob = logprob_new
        param_phase = param_phase_new
        param_amp = param_amp_new
    end
end
```

The status bar at the bottom indicates "Julia 1.2.0 | Idle".

The screenshot shows a Jupyter Notebook interface running on a web browser. The top navigation bar includes links to various websites like CAMB, Corp, Dan F., Preser, emcos, Probe, Data, Metro, Tweet, Using, Julia h., Julia Heath, PyPics, 2018-, distnd, and a search bar for 'wl'. The main window has a file menu with File, Edit, View, Run, Kernel, Tabs, Settings, Help. A sidebar on the left shows a file tree with 'mcmc.ipynb' selected. The central area contains a terminal window titled 'Terminal 1' and a code editor window titled 'mcmc.ipynb'. The code in the editor is written in Julia and performs MCMC sampling:

```
# how many steps to take
nsteps = 100_000

# how many times we accepted a move
naccepts = 0

for i in 1:nsteps

    # propose new parameter values
    param_phase_new = param_phase + randn() * jump_phase
    param_amp_new   = param_amp   + randn() * jump_amp

    # compute log-posterior at new parameters
    logprob_new = sine_log_posterior(times, y, param_phase_new, param_amp_new)

    if (exp(logprob_new - logprob) >= rand(Float64))
        naccepts += 1
        logprob = logprob_new
        param_phase = param_phase_new
        param_amp = param_amp_new
    end
    append!(chain, (param_phase, param_amp))
end
# append! makes "chain" a 1-d vector; reshape to a matrix
chain = reshape(chain, (2,Int64(length(chain)/2)));
```

The bottom status bar indicates 'Julia 1.2.0 | Idle', 'Mode: Edit', 'Ln 18, Col 8', and 'mcmc.ipynb'.

The screenshot shows a Jupyter Notebook interface with a Julia kernel. The left sidebar displays a file tree for a directory named 'mcmc-notes' containing files like 'header.tex', 'mcmc.ipynb' (which is selected), 'mcmc.tex', 'p99-data.txt', and 'README.md'. The main area has two tabs: 'Terminal 1' and 'mcmc.ipynb'. The 'mcmc.ipynb' tab contains the following Julia code:

```
# compute log-posterior at new parameters
logprob_new = sine_log_posterior(times, y, param_phase_new, param_amp_new)

if (exp(logprob_new - logprob) >= rand(Float64))
    naccepts += 1
    logprob = logprob_new
    param_phase = param_phase_new
    param_amp = param_amp_new
end
append!(chain, (param_phase, param_amp))
end
# append! makes "chain" a 1-d vector; reshape to a matrix
chain = reshape(chain, (2,Int64(length(chain)/2)))';

[36]: naccepts/nsteps
[36]: 0.9063

[15]: chain

[15]: 1000×2 LinearAlgebra.Adjoint{Any,Array{Any,2}}:
 0.00224217  6.06881
 0.00320339  6.15366
 0.00320339  6.15366
 0.00936336  6.27915
 0.00179649  6.20945
 0.0084974   6.40282
```

The terminal output shows the result of the division [36]: naccepts/nsteps, which is 0.9063. The final cell [15]: chain displays a 1000x2 matrix of numerical values.

1 s_ 6 🌐 Julia 1.2.0 | Idle

Mode: Edit ✎ Ln 1, Col 16 mcmc.ipynb

The screenshot shows a Jupyter Notebook interface with a Julia kernel. The left sidebar displays a file tree with files like `mcmc.ipynb`, `header.tex`, and `p99-data.txt`. The main area has two tabs: `Terminal 1` and `mcmc.ipynb`. The `mcmc.ipynb` tab contains the following Julia code:

```
# compute log-posterior at new parameters
logprob_new = sine_log_posterior(times, y, param_phase_new, param_amp_new)

if (exp(logprob_new - logprob) >= rand(Float64))
    naccepts += 1
    logprob = logprob_new
    param_phase = param_phase_new
    param_amp = param_amp_new
end
append!(chain, (param_phase, param_amp))
end
# append! makes "chain" a 1-d vector; reshape to a matrix
chain = reshape(chain, (2,Int64(length(chain)/2)))';

[37]: 100.*naccepts/nsteps
[37]: 90.63

[15]: chain

[15]: 1000×2 LinearAlgebra.Adjoint{Any,Array{Any,2}}:
 0.00224217  6.06881
 0.00320339  6.15366
 0.00320339  6.15366
 0.00936336  6.27915
 0.00179649  6.20945
 0.0084974   6.40282
```

The notebook is running on Julia 1.2.0. The status bar at the bottom indicates "Mode: Command" and "Ln 1, Col 6 mcmc.ipynb".

The screenshot shows a Jupyter Notebook interface running on a web browser. The top bar includes standard browser controls like back, forward, and search, along with a URL bar showing <https://symmetry1/user/dlang/lab>. The main window has a toolbar with icons for file operations, followed by a menu bar with File, Edit, View, Run, Kernel, Tabs, Settings, and Help. On the left is a sidebar with a file tree showing a directory structure under `/ ... / dlang / mcmc-notes /`, including files like `header.tex`, `mcmc.ipynb` (which is selected), `mcmc.tex`, `p99-data.txt`, and `README.md`. The central area contains a terminal window titled "Terminal 1" and a code editor window titled "mcmc.ipynb". The code in the editor is written in Julia and performs MCMC sampling:

```
# compute log-posterior at new parameters
logprob_new = sine_log_posterior(times, y, param_phase_new, param_amp_new)

if (exp(logprob_new - logprob) >= rand(Float64))
    naccepts += 1
    logprob = logprob_new
    param_phase = param_phase_new
    param_amp = param_amp_new
end
append!(chain, (param_phase, param_amp))
end
# append! makes "chain" a 1-d vector; reshape to a matrix
chain = reshape(chain, (2,Int64(length(chain)/2)))';

[39]: 100. * naccepts/nsteps
[39]: 58.176

[15]: chain

[15]: 1000×2 LinearAlgebra.Adjoint{Any,Array{Any,2}}:
 0.00224217  6.06881
 0.00320339  6.15366
 0.00320339  6.15366
 0.00936336  6.27915
 0.00179649  6.20945
 0.0084974   6.40282
```

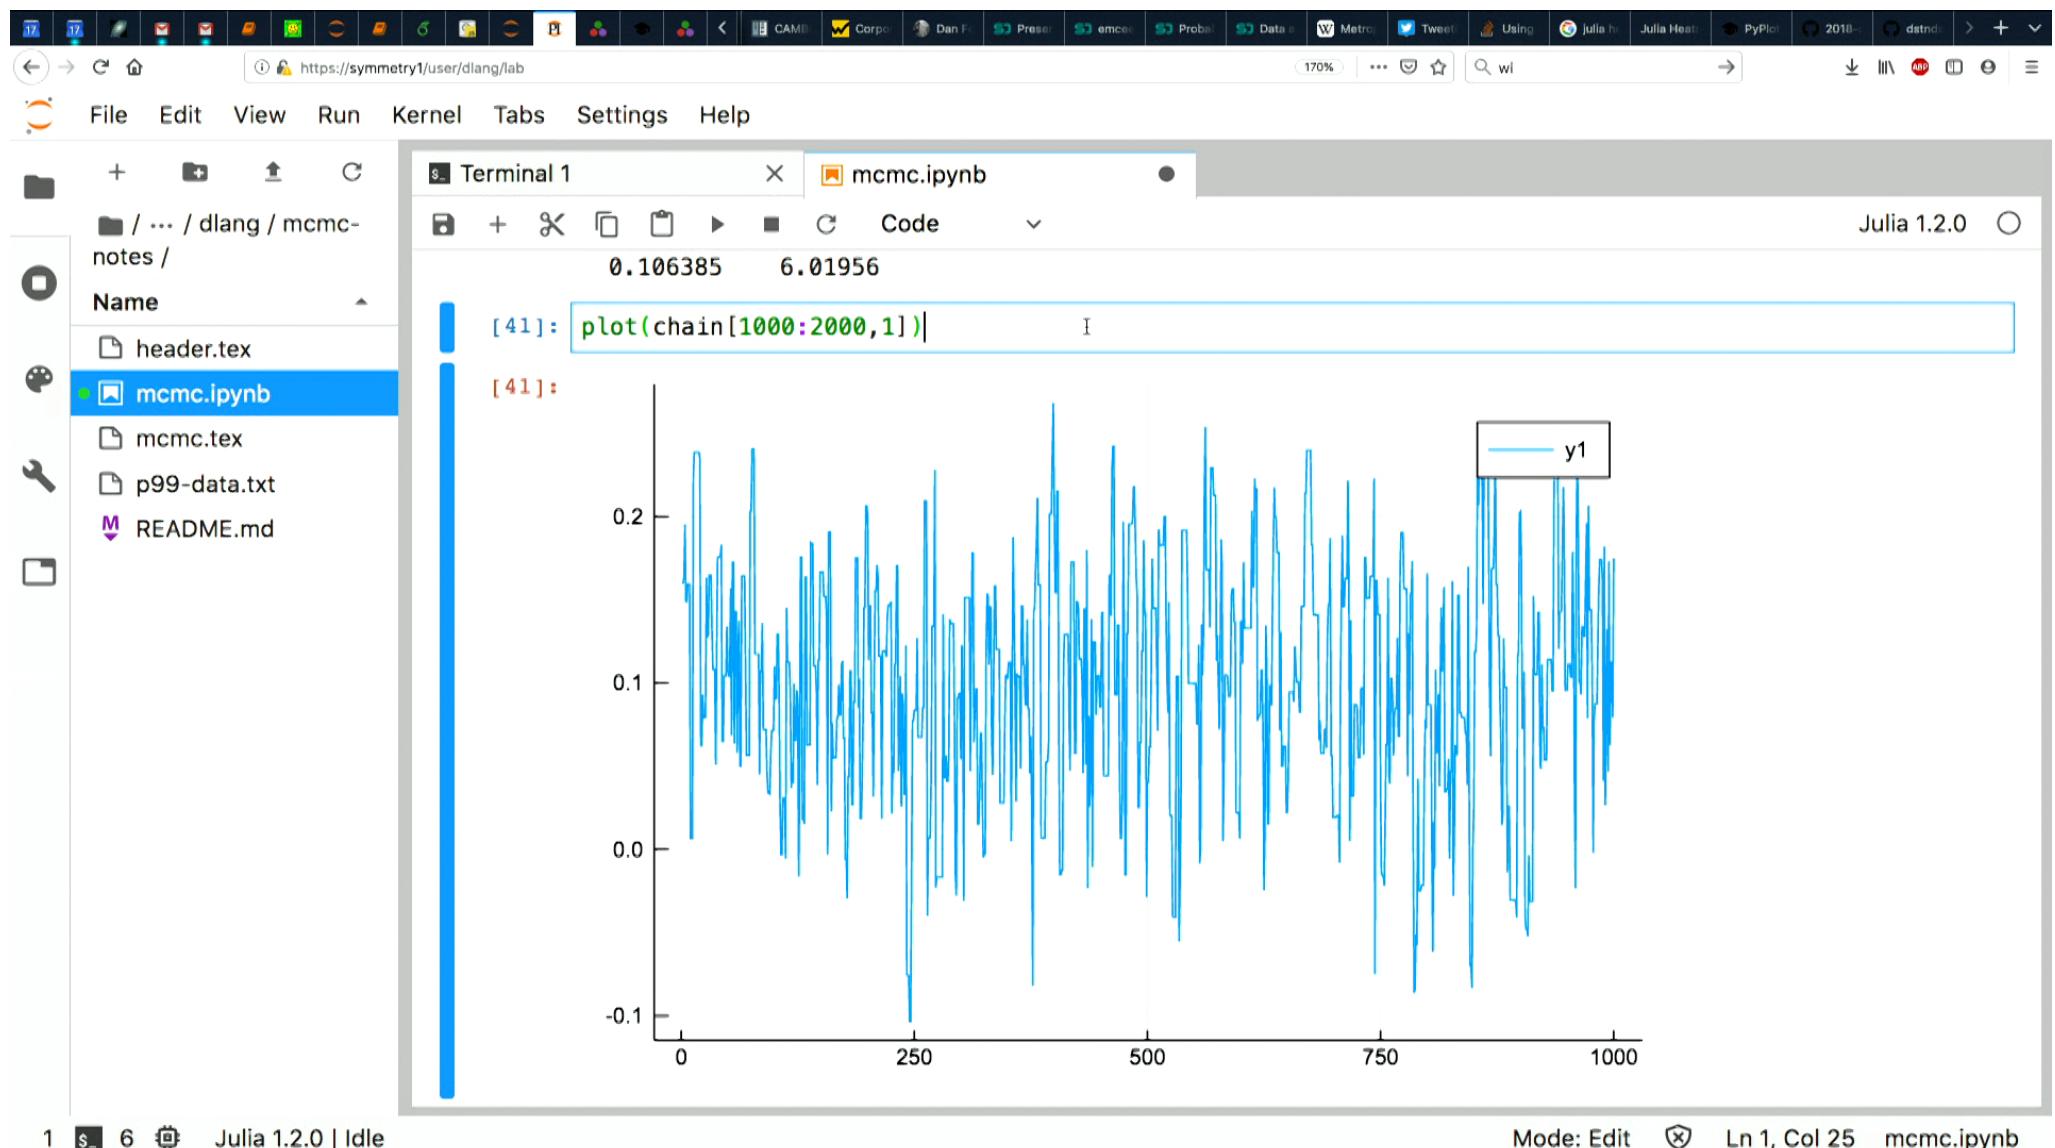
The terminal window shows the output of the code execution, including the acceptance rate (58.176%) and the resulting 1000x2 matrix `chain`. The bottom status bar indicates "Julia 1.2.0 | Idle", "Mode: Command", "Ln 1, Col 6", and the file name "mcmc.ipynb".

The screenshot shows a Jupyter Notebook interface with a Julia 1.2.0 kernel. The left sidebar displays a file tree with files like header.tex, mcmc.ipynb (selected), mcmc.tex, p99-data.txt, and README.md. The main area has two tabs: Terminal 1 and mcmc.ipynb. The mcmc.ipynb tab shows the following Julia code and its output:

```
[39]: 58.176
[40]: chain
[40]: 100000×2 LinearAlgebra.Adjoint{Any,Array{Any,2}}:
      0.0      6.0
      -0.0221044  6.00859
      0.191279   6.005
      0.182272   5.87561
      0.182272   5.87561
      0.182272   5.87561
      0.16562    6.0168
      0.16562    6.0168
      0.16562    6.0168
      0.16562    6.0168
      0.175974   5.94926
      0.137782   5.99152
      0.137782   5.99152
      :
      0.284826   5.98298
      0.231389   5.94631
      0.272986   6.00874
      0.178822   6.19027
      0.00304063 6.12265
      0.00664297 6.15917
      -0.0450753  6.12275
      0.0422983   6.14672
      0.124702   6.02102
```

1 s_ 6 🏃 Julia 1.2.0 | Idle

Mode: Command ✘ Ln 1, Col 18 mcmc.ipynb



The screenshot shows a Jupyter Notebook interface running on a web browser. The top bar includes standard browser controls like back, forward, and search, along with tabs for various open documents. The main area has a file tree on the left containing files like `header.tex`, `mcmc.ipynb` (which is selected), `mcmc.tex`, `p99-data.txt`, and `README.md`. The central workspace contains two tabs: "Terminal 1" and "mcmc.ipynb". The "mcmc.ipynb" tab displays a Julia script for performing Metropolis-Hastings MCMC. The code defines parameters `nsteps` and `naccepts`, initializes parameters `param_phase` and `param_amp`, and then enters a loop. Inside the loop, it proposes new parameter values, computes the log-posterior at the new parameters, and checks if the proposal is accepted based on the log-probability difference. It then appends the new parameters to a chain and reshapes the chain into a matrix. The status bar at the bottom indicates "Julia 1.2.0 | Idle" and "Mode: Command".

```
# now many steps to take
nsteps = 100_000

# how many times we accepted a move
naccepts = 0

for i in 1:nsteps

    # propose new parameter values
    param_phase_new = param_phase + randn() * jump_phase
    param_amp_new   = param_amp   + randn() * jump_amp

    # compute log-posterior at new parameters
    logprob_new = sine_log_posterior(times, y, param_phase_new, param_amp_new)

    if (exp(logprob_new - logprob) >= rand(Float64))
        naccepts += 1
        logprob = logprob_new
        param_phase = param_phase_new
        param_amp = param_amp_new
    end
    append!(chain, (param_phase, param_amp))
end
# append! makes "chain" a 1-d vector; reshape to a matrix
chain = reshape(chain, (2,Int64(length(chain)/2)))';

[ 39 ]: 100. * naccepts/nsteps
```

MCMC: summary

Things MCMC can do for you

- ▶ produce constraints on model parameters, including covariances
- ▶ include (but ignore) *nuisance parameters* in your data analysis

Things MCMC doesn't do for you

- ▶ tell you if it has converged!
- ▶ tell you if your model is good
- ▶ (easily) allow comparisons between models
- ▶ handle multi-modal distributions (try *nested sampling*)
- ▶ Dan Foreman–Mackey has good notes on **advanced sampling methods**

$$G^{\mu\nu} = \partial_{\mu}\phi \sum_{n=1}^{\infty} \frac{1}{n!}$$

Ex X

$$P(x|t) = G$$
$$m(\theta) = \underbrace{A}_{\{A_i\}}$$
$$P(\{x_i\} | \theta)$$



$$G^{\mu\nu} = \partial_{\mu}\partial_{\nu} S^{00}/c^2$$



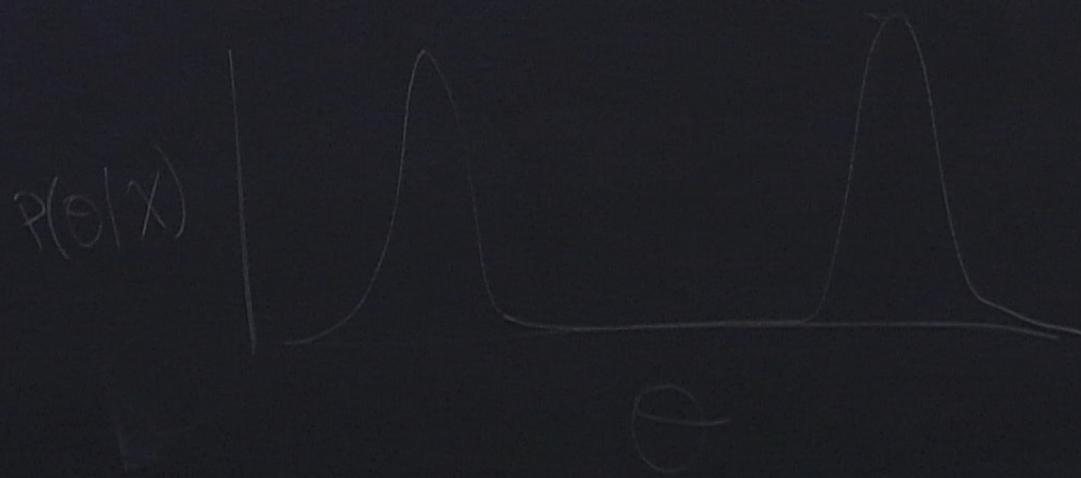
⊗ X

$$P(x|t) = \text{Gauss}$$

$$m(\theta) = \underbrace{A \sin(\theta)}_{\{A\}}$$

$$P(\{x\} | \theta) \propto \prod_i$$

$$G^{\mu\nu} = \partial^\mu \phi \partial^\nu \phi - \text{some terms}$$

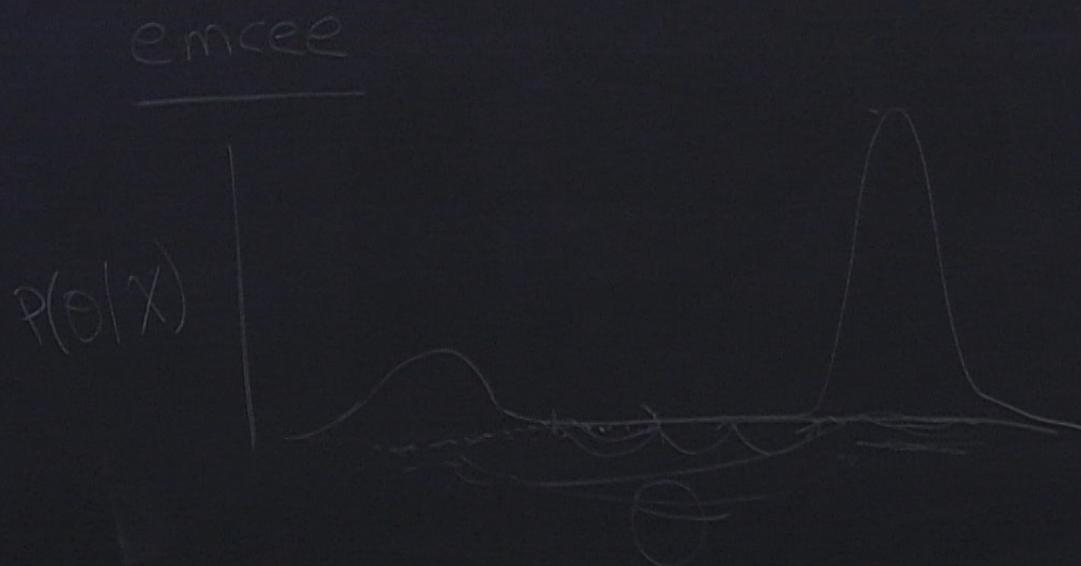


$$\text{Ansatz: } P(x|t) = \text{Gauss}$$

$$m(\theta) = \underbrace{A \sin(\theta)}_{\{A\}}$$

$$P(\{x\} | \theta) \propto \prod_i$$

$$G^{\mu\nu} = \partial_\mu S_{\alpha\beta}^{\mu\nu} \partial^\beta$$



fix X

$$P(x | t) = \text{Gauss}$$

$$m(\theta) = A \sin(\underbrace{\{A\}}_{\theta})$$

$$P(\{x\} | \theta) \propto$$

