how to do integrals the brute force way with dan foreman-mackey

how to do integrals using differentiation with dan foreman-mackey

[who am I?] Dan Foreman-Mackey

[where am I?] SF / FI / CCA

[what do I do?] exoplanets?

[wth am I doing here?]

[what do I do?] write **software** for astronomy

[what do I do?] focus on implementation

[what do I do?] try to **learn things** & **share them**

[today's takeaways: #1]
data analysis in astrophysics
 is getting ambitious

[today's takeaways: #1] data analysis in grav waves is getting ambitious

[today's takeaways: #2]
even if you're not doing machine
 learning, there are
 useful open source tools

[today's takeaways: #3] using them might be easy or a complete pain

[today's takeaways: #4] it might be worth it

[1] inference

want: data ⇒ physics

have: physics ⇒ data

integral of the form f(physics) p(physics data) dphysics



[one option]
(markov chain) monte carlo
physics ~ p(physics | data)

so, "all" you need is:

[a] a "good" sampler

[b] fast p(data | physics)

[a] a "good" sampler cost per effective sample

[a] a "good" sampler depends on # of parameters and "geometry" of the problem

[b] fast p(data physics) should be interpretable

[b] fast p(data|physics)
 depends on size of data
 and simplifying assumptions

in astrophysics, we want to
[] do rigorous inference
[] with huge datasets, and
[] physics-based models

in astrophysics, we want to [] do rigorous inference [] with huge datasets, and [] physics-based models [] finish in finite time

so. what do we do?



Theme by the Executable Book

Project

A more complete example is available in the Quickstart tutorial.

Ē ☆ 53 \mathbf{c}

∷ Contents

Basic Usage How to Use This Guide License & Attribution Changelog

8

emcee is an MIT licensed pure-Python implementation of Goodman & Weare's Affine Invariant Markov chain Monte Carlo (MCMC) Ensemble sampler and these pages will show you how to use it.

This documentation won't teach you too much about MCMC but there are a lot of resources available for that (try this one). We also published a paper explaining the emcee algorithm and implementation in detail.

emcee has been used in quite a few projects in the astrophysical literature and it is being actively

If you wanted to draw samples from a 5 dimensional Gaussian, you would do something like:





emcee isn't a very "good" sampler





ref: personal experience

tenish not outrageously many number of parameters





encee tenish not outrageously many number of parameters

ref: personal experience





ref: personal experience



[X] huge datasets [?] finite time

[X] rigorous inference [X] physics-based models

this is a solved problem*





[2] gradients
dp(data physics) / dphysics



log p(data physics)

physics



U(**q**) П p(data physics) -log



U(**q**) П p(data physics) -log



П p(data physics) -log

physics = q

cool. so we're done?

maybe!



let's not get ahead of ourselves



but, probably not.

[X] huge datasets [X] finite time

[X] rigorous inference [?] physics-based models

dp(data physics) / dphysics

automatic differentiation

your model is just code

apply the chain rule

apply the **chain rule** over and over again...



it's not! (mostly)

what about things like: M = E - e sin(E)

ref: Kepler (1609)



custom "ops" [a] C/C++/Fortran/CUDA/... code [b] derivative rules

[3] my (current) recommendations

[caveat] I mostly work in Python

[caveat] this is **a moving target**

[caveat] this is not a comprehensive list





Stan state-of-the-art inference algorithms

TensorFlow + TF Probability fast; powerful but poorly documented inference algorithms

+ Numpyro/TF Probability the above?

my current recommendation



import numpy as np \Rightarrow

import jax.numpy as jnp

[4] a case study

e exoplanet – exoplanet × + $\leftarrow \rightarrow$ ○ △ https://docs.exoplanet.codes/en/latest/ \mathbf{C} ÷ e exoplanet

exoplanet is a toolkit for probabilistic modeling of time series data in astronomy with a focus on observations of exoplanets, using PyMC3. *PyMC3* is a flexible and high-performance model building language and inference engine that scales well to problems with a large number of parameters. *exoplanet* extends *PyMC3*'s language to support many of the custom functions and distributions required when fitting exoplanet datasets. These features include:

- A fast and robust solver for Kepler's equation.
- Scalable Gaussian Processes using celerite.
- Fast and accurate limb darkened light curves using starry.
- And many others!

All of these functions and distributions include methods for efficiently calculating their gradients so that they can be used with gradient-based inference methods like Hamiltonian Monte Carlo, No U-Turns Sampling, and variational inference. These methods tend to be more robust than the methods more commonly used in astronomy (like ensemble samplers and nested sampling) especially when the model has more than a few parameters. For many exoplanet applications, *exoplanet* (the code) can improve the typical performance by orders of magnitude.

exoplanet is being actively developed in a public repository on GitHub so if you have any trouble, open an issue there.

- **?** Where to find what you need

For more in depth examples of *exoplanet* used for more realistic problems, go to the Case studies page.

USER GUIDE

Installation

Citing exoplanet & its dependencies

exoplanet

Q Search the docs ...

Theano vs. Aesara

Multiprocessing

API documentation

Developer documentation

Changelog

TUTORIALS

About these tutorials Automatic differentation & gradientbased inference A quick intro to PyMC3 Data & models Light travel time delay Reparameterization Case studies Z celerite2 🗹

≣ ☆

[]

 \mathbf{O}

∷ Contents

Contents License & attribution 80

• Common reparameterizations for exoplanet-specific parameters like limb darkening and eccentricity.

For general installation and basic usage, continue scrolling to the table of contents below.





celerite \se.le.вi.te\ *noun, archaic literary* A scalable method for Gaussian Process regression in one dimension. From French célérité.

celerite2

celerite2

User Guide Installation Upgrading from celerite Citing celerite2 Tutorials Getting started API Details Python interface Theano interface C++ interface

celerite2 ¶

celerite is an algorithm for fast and scalable Gaussian Process (GP) Regression in one dimension and this library, *celerite2* is a re-write of the original celerite project to improve numerical stability and integration with various machine learning frameworks. This implementation includes interfaces in Python and C++, with full support for Theano/PyMC3 and JAX.

This documentation won't teach you the fundamentals of GP modeling but the best resource for learning about this is available for free online: Rasmussen & Williams (2006). Similarly, the *celerite* algorithm is restricted to a specific class of covariance functions (see the original paper for more information and a recent generalization for extensions to structured two-dimensional data). If you need scalable GPs with more general covariance functions, GPyTorch might be a good choice.

celerite2 is being actively developed in a public repository on GitHub so if you have any trouble, open an issue there.

User Guide

Installation

Using pip

$\boxed{\bigcirc}$ $\boxed{\bigcirc}$ $\boxed{\bigcirc}$ $\boxed{\bigcirc}$ $\boxed{\bigcirc}$ $\boxed{\bigcirc}$ $\underbrace{\bigcirc}$ $\underbrace{\bigcirc}$ Celerite2 40 Stars \cdot 4 Forks

Contents

celerite2 License & attribution Changelog

🗐 v: latest 🔻

[gaussian process] a likelihood function for correlated noise

[gaussian process] in my case, caused by stochastic stellar variability

[gaussian process]
in your case, caused by
systematics and/or GW background



[gaussian process] the problem: O(N³) scaling

[gaussian process] in special cases: O(NlogN) scaling

[gaussian process] with celerite: O(N) scaling


•••	\bigcirc celerite2/forward.hpp at main $\cdot e \times +$
$\leftarrow \ \rightarrow \ {\rm G}$	https://github.com/exoplanet-dev/celerite2/blob/ma
	<pre>celerite2/forward.hpp at main .e × +</pre>
	<pre>79) { 80 ASSERT_ROW_MAJOR(Work); 81 82 typedef typename Diag::Scalar Scalar; 83 typedef typename Eigen::internal::plain_row_type<lowrank>::type R 84 typedef typename Eigen::internal::plain_col_type<coeffs>::type Co 85 86 Eigen::Index N = U.rows(), J = U.cols(); 87 CAST_VEC(DiagOut, d, N); 88 CAST_MAT(LowRankOut, W, N, J); 89 CAST_BASE(Work, S); 90 if (update_workspace) { 91 S.derived().resize(N, J * J); 92 S.row(0).setZero(); 93 } </coeffs></lowrank></pre>

			~
nain/c%2B%2B/include/celerite2/forward.hpp	80% 🟠		≡
nd `W_out` can			
case, the but this			
esky factor			
n step			
oeffs, <mark>typename</mark> Diag, typename LowRank, typename DiagOut, typename LowRankOut,			
′ (N,) ′ (J,)			
(N,)			
(N, J)			
(N,) (N, J)			
(N, J*J)			
RowVector:			
coeffVector;			

https://celerite2.readthedocs.io/en/latest/tutoria

celerite2

Also

 $\leftarrow \ \ \rightarrow \ \ \mathbf{G}$

User Guide

Installation

Upgrading from celerite

Getting started

Citing celerite2

Tutorials

Getting started

API Details

Python interface Theano interface

C++ interface

Posterior inference using emcee

Now, to get a sense for the uncertainties on our numerically estimate the posterior expectations to run our MCMC. Our likelihood function is the choose a wide normal prior on each of our para

```
[9]: import emcee
    prior_sigma = 2.0
    def log_prob(params, gp):
        gp = set_params(params, gp)
        return (
             gp.log_likelihood(y) - 0.5
            gp.kernel.get_psd(omega),
    np.random.seed(5693854)
    coords = soln.x + 1e-5 * np.random.
    sampler = emcee.EnsembleSampler(
        coords.shape[0], coords.shape[1
    state = sampler.run_mcmc(coords, 200
    sampler.reset()
    state = sampler.run_mcmc(state, 5000
                     2000/2000 [00:32<00
    100%|
                   | 5000/5000 [01:20<00
    100%|
```

After running our MCMC, we can plot the predic chain. This gives a qualitative sense of the unce

ls/first/#Posterior-inference-using-emcee	② 目 公 =
Q Search	celerite2 40 Stars · 4 Forks
e r model, let's use Markov chain Monte Carlo (MCMC) to s of the model. In this first example, we'll use the emcee package same as the one we used in the previous section, but we'll also ameters.	Contents Getting started Maximum likelihood Posterior inference using emcee Posterior inference using PvMC3
* np.sum((params / prior_sigma) ** 2),	Posterior inference using numpyro Comparison
randn(32, len(soln.x))	
], log_prob, args=(gp,) 00, progress=True) 0. progress=True)	
0:00, 62.37it/s] 0:00, 62.34it/s]	
ctions that the model makes for a handful of samples from the ertainty in the predictions.	■ v: latest ▼

🛑 😑 🔵 🥢 Getting started — celerite2	× +
$\leftarrow \rightarrow C$ \bigcirc \triangle htt	ps://celerite2.readthedocs.io/en/latest/tutorials/firs
Getting started	
	<pre>import numpyro.distributions as dist</pre>
celerite2	<pre>from numpyro.infer import MCMC, NUTS</pre>
User Guide	<pre>import celerite2.jax</pre>
Installation	<pre>from celerite2.jax import terms as jax_t</pre>
Upgrading from celerite	
Citing celerite2	<pre>def numpyro_model(t, yerr, y=None):</pre>
Tutorials	<pre>mean = numpyro.sample("mean", dist.N</pre>
Getting started	<pre>log_jitter = numpyro.sample("log_jit</pre>
API Details	<pre>log_sigma1 = numpyro.sample("log_sig log_sho1 = numpyro.sample("log_sig</pre>

```
Python interface
Theano interface
```

```
C++ interface
```

```
term2 = jax_terms.OverdampedSHOTerm(
```

```
kernel = term1 + term2
gp = celerite2.jax.GaussianProcess(kernel, mean=mean)
```

```
numpyro.sample("obs", gp.numpyro_dist(), obs=y)
numpyro.deterministic("psd", kernel.get_psd(omega))
```

```
nuts_kernel = NUTS(numpyro_model, dense_mass=True)
rng_key = random.PRNGKey(34923)
%time mcmc.run(rng_key, t, yerr, y=y)
```

```
CPU times: user 16.8 s, sys: 122 ms, total: 16.9 s
Wall time: 16.9 s
```



[recently]
generalized to
multivariate data

celerite2.readthedocs.io

```
\bigcirc celeriac/ops.py at main \cdot dfm/cel	imes +
\leftarrow \rightarrow \mathbf{C}
                             ○ A https://github.com/dfm/celeriac/blob/main/src/cel
                 if P.ndim == 1:
                     if transpose:
                         return other * P[None, :]
                     return P[:, None] * other
         82
                 if transpose:
                     return other @ P
                 return P.T @ other
         88 Carry = Tuple[Array, Array, Array]
         89 Data = Tuple[Array, Array, Array, Array]
         90 MatmulData = Tuple[Array, Array, Array]
    •••• 93 def _factor_impl(
                state: Carry, data: Data
             ) -> Tuple[Carry, Tuple[Array, Array]]:
                Sp, dp, Wp = state
                an, Un, Vn, Pn = data
                 Sn = _pdot(Pn, _pdot(Pn, Sp + dp * jnp.outer(Wp, Wp), transpose=
                 tmp = Sn @ Un
                 dn = an – tmp @ <mark>Un</mark>
                 Wn = (Vn - tmp) / dn
                 return (Sn, dn, Wn), (dn, Wn)
        105 def _solve_impl(state: Carry, data: Data) -> Tuple[Carry, Array]:
        106 Fp, Wp, Zp = state
                Un, Wn, Pn, Yn = data
                Fn = _pdot(Pn, Fp + jnp.outer(Wp, Zp))
                 Zn = Yn - Un @ Fn
                 return (Fn, Wn, Zn), Zn
        113 def _matmul_impl(state: Carry, data: MatmulData) -> Tuple[Carry, Arr
                 (Fp, Vp, Yp) = state
        115 Vn, Pn, Yn = data
                 Fn = _pdot(Pn, Fp + jnp.outer(Vp, Yp))
                 return (Fn, Vn, Yn), Fn
        118
```

		~
eriac/ops.py#L93-L102	☆	⊌ =
True))		
ay]:		

$\leftarrow \rightarrow C \qquad \bigcirc A \text{ https://github.com/exoplanet-dev/celerite2/blob/main/c%2B%2B/include/celerite2/forward.hpp#L69-L135} 50\% \textcircled{2}$	Celerite2/forward.hpp at main · e × +			~
b4 * @param V (N, J): Ine second low rank matrix 65 * @param W out (N,): The diagonal component of the Cholesky factor 66 * @param W out (N, I): The second low rank component of the Cholesky factor	\leftarrow \rightarrow C \bigcirc A https://github.co	m/exoplanet-dev/celerite2/blob/main/c%2B%2B/include/celerite2/forward.hpp#L69-L135	50% 公	
<pre> Auto Contain Contain</pre>		<pre>strumt =</pre>		

[5] summary & remaining issues

[X] huge datasets [X] finite time

[X] rigorous inference [X] physics-based models





documentation & tutorials?



```
[7]: [obs]
```

```
[8]: model.logp({"mu": 0})
```

```
[8]: array(-136.56820547)
```

It's worth highlighting the design choice we made with logp. As you can see above. logp is being called with arguments, so it's a method of the model instance. More precisely, it puts

				₽ •••		lii/	•	0	▣
Developer Guide	About PyMC3	Search	Q	Ģ					
					-				
Warning)									
om variables (RVs) and	computes the model logp a	nd its gradients. Usually, yo	u would inst	antiate it as					
-lass									
class.									
.random.randn(10	0))								



									~
					⊠ ☆	III\ 🗊	0	•	≡
Developer Guide	About PyMC3	Search	Q	0					
Warning)									
		-	71						
	-	hat							
om taules V and	l coputes the moder log	p and its gradients. Usually		ntiate it as					
		, and the second s							
-									
1955									
1055.									
random randn/10									
.random.randn(10	(0))								

reparameterization



integration with legacy code
 remains Hard™

ambitious data analysis calls for powerful tools that may (or may not) work out of the box but we have work to do

I want you to keep doing awesome science and learning about new methods then share what you learn and contribute back

get in touch! dfm.io github.com/dfm twitter.com/exoplaneteer