Probabilistic Programming and Inference

Large-Scale Data Processing and Optimisation 13 Nov 2024

Hong Ge hg344@cam.ac.uk

Machine Learning Group, Department of Engineering, University of Cambridge

slides courtesy of Tor Fjelde

Bayesian inference: probability view

Given some "real world" data ${\mathfrak D}$

Want to determine parameters θ given "real world" data \mathcal{D} , i.e.

 $p(\theta \mid D)$

 $\begin{aligned} \theta &\sim p(\theta) \\ \mathcal{D} &\sim p(\mathcal{D} \mid \theta) \end{aligned}$

Bayes' rule tells us that $p(\theta \mid D) \propto p(D \mid \theta)p(\theta)$ Which we can write as generative model

Bayesian inference: simulator view

Have simulator which, given parameters theta, can generate new data $\tilde{\mathcal{D}}$

Have data ${\mathfrak D}$ from "real world"

Have prior knowledge regarding values of theta are probable, giving rise to a prior distribution $p(\theta)$

Aim: determine parameters θ such that the generated data $\tilde{\mathcal{D}}$ is close to true data \mathcal{D} , i.e. sample from $p(\theta \mid \mathcal{D})$.

 $\theta \sim p(\theta)$ $\mathcal{D} \sim p(\mathcal{D} \mid \theta)$

Bayesian inference: biased coin

Data \mathcal{D} of coin flips from a **biased** coin $\mathcal{D} = (H, H, H, T, \dots)$ **Initial belief:** it's a fair coin

Update our belief using the observations

More data => more certain!

$$\theta \sim \text{Beta}(3,3)$$

 $x_i \sim \text{Bernoulli}(\theta) \text{ for } i = 1, ..., n$



Bayesian inference: biased coin

Data \bigcirc of coin flips from a **biased** coin WARNING! Initial Updating the posterior with Update our b incoming data is, generally, More not that as easy as these animations make it seem! $\theta \sim \text{Beta}$



 $x_i \sim \text{Bernoulli}(\theta)$ for i = 1, ..., n

Bayesian inference: linear regression

$$\begin{split} \mathcal{D} &= \left((x_1, y_1), \dots, (x_n, y_n)\right) \\ \alpha &\sim \mathcal{N}(0, 1) \\ \beta &\sim \mathcal{N}(0, I) \\ \sigma &\sim \mathcal{N}_+(0, 1) \\ \mu_i &= \beta x_i + \alpha \\ y_i &\sim \mathcal{N}(\mu_i, \sigma^2) \end{split}$$

Bayesian inference: linear regression

$$\mathcal{D} = ((\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{y}_n))$$

 $\alpha \sim \mathcal{N}(0,1)$ $\beta \sim \mathcal{N}(0,I)$ $\sigma \sim \mathcal{N}_{+}(0,1)$ $\mu_{i} = \beta x_{i} + \alpha$ $y_{i} \sim \mathcal{N}(\mu_{i},\sigma^{2})$





Bayesian inference: in practice

Posterior $p(\theta \mid D)$ is generally **not available in closed-form**

Requires various approximate methods to inspect $p(\theta \mid \mathcal{D})$

E.g. Markov Chain Monte Carlo (MCMC), variational inference (VI)

Annoying & difficult to implement!

Bayesian inference: summary

Principled way to incorporate prior knowledge through p(theta)

Allows quantification of uncertainty

imes Non-trivial for interesting problems

Both computationally expensive and requires knowledge of inference methods

 $\begin{aligned} \theta &\sim p(\theta) \\ \mathcal{D} &\sim p(\mathcal{D} \mid \theta) \end{aligned}$

Aim: make Bayesian inference approachable.

User should only have to write down their model, and the rest should be "automatic"

Key ideas:

Decouple modelling and inference
 General-purpose inference





PYMC BUGS



Bean Machine

Probabilistic programming language (PPL)

Suite of tools allowing easy specification of models + inference method

Many different approaches to PPLs



Stan

Anglican PYRO

Many different approaches to PPLs

Stan

Has its own DSL

Implemented in C++

Provides convenient interfaces in R, Python, Julia, etc.

Very well-established with great track record (gold standard)









Many different approaches to PPLs

BUGS

Has its own DSL

Implemented in Component Pascal

Provides convenient interface in R

Old now, but was one of the very first PPLs





TuringLang/Turing.jl Bayesian inference with probabilistic programming



Many different approaches to PPLs

PyMC, **Pyro** and **Bean**

Embedded in Python

Relies on underlying "graph" computational frameworks, e.g. JAX, PyTorch

But have to work with a subset of Python compatible with framework



Many different approaches to PPLs

Model and inference are closely coupled

Model and inference are decoupled, but modelling language is closely coupled with inference

e.g. BUGS: Gibbs, Stan: HMC, Infer.NET: EP/VMP

Modelling language is decoupled with inference details, but closely coupled with a computational backend

e.g. Pyro: JAX, PyTorch; PyMC: Theano, BlackJAX: JAX





Anglican



Stan





Turing.jl: what is it?

 \Rightarrow More general purpose

Inference-agnostic: models are compatible with a wide range of MCMC and VI algorithms

Interoperability with other modelling libraries (GPs, DiffEqs): requires all modelling libraries using the same computational backend

Good tools for profiling/debugging/interactively_run
models and inference; High performance

Modular / Compositional inference among others





Stan





Turing.jl: what is it?

TuringLang/**Turing.jl**

Bavesian inference with probabilistic programming



Fully implemented in the **julia** programming language

Approachable for users familiar with Matlab, Python and R

Support BUGS syntax via JuliaBUGS

Growing in popularity

\sim	88	U	69	5	23	\mathcal{M}	ZK	ጉ	213	
	Contributors		Issues		Discussions		Stars		Forks	

```
1 using Turing
2
3 # Define the model.
4 @model function demo()
5  θ ~ Prior
6  Ø ~ Likelihood(θ)
7 end
8
9 # Instantiate the model.
10 model = demo()
11 # Condition `model` on the data.
12 model_cond = model | (Ø = data,)
13 # Sample from the posterior.
14 chain = sample(model cond, NUTS(), 1000)
```

Turing.jl: what is it?

TuringLang/**Turing.jl**

Bayesian inference with probabilistic programming



Fully implemented in the unit programming language

Approachable for users familiar with Matlab, Python and R

Support BUGS syntax via JuliaBUGS

Growing in popularity

X 88	0 69	V-U 23	א ∠K	7 213	
Contributors	Issues	Discussions	Stars	Forks	

Turing.jl: Bayesian linear regression

using Turing

```
# Define model.
@model function linear_regression(x)
    # Priors.
    α ~ Normal(0, 1)
    β ~ Normal(0, 1)
    σ<sup>2</sup> ~ InverseGamma(2, 3)
    # Likelihood.
    y ~ MvNormal(β * x .+ α, σ<sup>2</sup>)
end
```

```
# Condition model on some data.
model = linear_regression(x_data) | (y = y_data, )
```

Perfurm inference!
sample(model, NUTS(), 1000)

$$\begin{split} \alpha &\sim \mathcal{N}(0,1) \\ \beta &\sim \mathcal{N}(0,I) \\ \sigma &\sim \mathcal{N}_{+}(0,1) \\ \mu_{i} &= \beta x_{i} + \alpha \\ y_{i} &\sim \mathcal{N}(\mu_{i},\sigma^{2}) \end{split}$$

JuliaBUGS: Rats model

```
model_def = @bugs("""model{
                                                             \beta_i \sim \operatorname{Normal}\left(\beta_c, \tau_\beta\right)
    for( i in 1 : N ) {
         for( j in 1 : T ) {
         Y[i , j] ~ dnorm(mu[i , j],tau.c)
         mu[i, j] <- alpha[i] + beta[i] * (x[j] - xbar)
         alpha[i] ~ dnorm(alpha.c,alpha.tau)
         beta[i] ~ dnorm(beta.c,beta.tau)
    tau.c ~ dgamma(0.001,0.001)
    sigma <- 1 / sqrt(tau.c)</pre>
    alpha.c \sim dnorm(0.0, 1.0E-6)
    alpha.tau ~ dgamma(0.001,0.001)
    beta.c ~ dnorm(0.0, 1.0E-6)
    beta.tau ~ dgamma(0.001,0.001)
    alpha0 <- alpha.c - xbar * beta.c
}""", false) # `false` means R-style variable names are kept
```

 $Y_{ij} \sim \operatorname{Normal}\left(lpha_i + eta_i \left(x_j - ar{x}
ight), au_c
ight)$

 $\alpha_i \sim \operatorname{Normal}\left(\alpha_c, \tau_{\alpha}\right)$

Nothing is different from the original BUGS

Bayesian workflow is important

Bayesian workflow^{*} Andrew Gelman[†] Aki Vehtari[‡] Daniel Simpson[§] Charles C. Margossian[†] Bob Carpenter[¶] Yuling Yao[†] Lauren Kennedy[∥] Jonah Gabry[†] Paul-Christian Bürkner^{***} Martin Modrák^{††}

2 Nov 2020

Should be "simple" to perform

Standard Bayesian workflow

Define model "in maths"

Define model in Turing.jl

Perform prior checks

Perform inference

Evaluate, e.g. predict

Rinse and repeat!



Standard Bayesian workflow

Define model "in maths"

Define model in Turing.jl

Perform prior checks

Perform inference

Evaluate, e.g. predict

Rinse and repeat!

$\sigma^{2} \sim \text{InverseGamma}(3, 4/10)$ $\gamma \sim \mathcal{N}(0, 10)$ $\beta \sim \mathcal{N}(0, \sigma^{2}I)$	<pre>model_conditioned = model (y = y_data,)</pre>
$\begin{array}{c c} \log y_i \sim \mathcal{N}(\beta \cdot x_i + \gamma, \sigma^*) & \text{for } i = 1, \dots, N \end{array}$ $\begin{array}{c c} \text{Turing.jl} & \text{Condition} \end{array}$	Perform inference
Consider a function linear_regression(x) d = size(x, 1) $\sigma^2 = InverseGamma(3, 4 / 10)$ $\gamma = Normal(0, 10)$	<pre>chain = sample(model_conditioned, NUTS(), 1000)</pre>
β ~ MvNormal(zeros(d), $\sigma^2 * I$) y ~ MvNormal(x' * β .+ γ , $\sigma^2 * I$) end model = linear_regression(x_data)	Post-inference analysis
Prior checks	$\begin{array}{cccccccccccccccccccccccccccccccccccc$
<pre>model_gen = fix(model, params_gen) (; y) = rand(model_gen) chain_gen = sample(model (y = y,), NUTS(), 1000)</pre>	
Identifiability	Predict
If necess	sary

Standard Bayesian workflow

Define model "in maths"

 $\sigma^{2} \sim \text{InverseGamma}(3, 4/10)$ $\gamma \sim \mathcal{N}(0, 10)$ $\beta \sim \mathcal{N}(0, \sigma^{2}I)$ $\log y_{i} \sim \mathcal{N}(\beta \cdot x_{i} + \gamma, \sigma^{2}) \text{ for } i = 1, \dots, N$



Standard Bayesian workflow

Define model "in maths"

Define model in Turing.jl

Perform prior checks

Perform inference

Evaluate, e.g. predict

Rinse and repeat!

$\begin{split} \sigma^2 &\sim \text{InverseGamma}(3, 4/10) \\ \gamma &\sim \mathcal{N}(0, 10) \\ \beta &\sim \mathcal{N}(0, \sigma^2 I) \\ \log y_i &\sim \mathcal{N}(\beta \cdot x_i + \gamma, \sigma^2) \text{for } i = 1, \dots, N \end{split}$	<pre>model_conditioned = model (y = y_data,)</pre>
Turing.jl Condition	Perform inference
Genodel function linear_regression(x) d = size(x, 1) σ^2 ⁻ InverseGamma(3, 4 / 10) γ ⁻ Normal(0, 10) β ⁻ MvNormal(zeros(d), $\sigma^2 * I$) y ⁻ MvNormal(x' * β .+ γ , $\sigma^2 * I$) end	chain = sample(model_conditioned, NUTS(), 1000)
<pre>model = linear_regression(x_data)</pre>	Post-inference analysis
	$\begin{array}{c c c c c c c c c c c c c c c c c c c $
Prior checks	
<pre>model_gen = fix(model, params_gen) (; y) = rand(model_gen) chain_gen = sample(model (y = y,), NUTS(), 1000)</pre>	
Identifiability	Predict
If necessa	

Standard Bayesian workflow

Define model in Turing.jl

```
\begin{array}{l} \texttt{Cmodel function linear_regression(x)} \\ \texttt{d} = \texttt{size}(\texttt{x}, 1) \\ \sigma^2 & \sim \texttt{InverseGamma}(\texttt{3}, \texttt{4} \neq \texttt{10}) \\ \gamma & \sim \texttt{Normal}(\texttt{0}, \texttt{10}) \\ \beta & \sim \texttt{Normal}(\texttt{o}, \texttt{10}) \\ \beta & \sim \texttt{MvNormal}(\texttt{zeros}(\texttt{d}), \sigma^2 * \texttt{I}) \\ \texttt{y} & \sim \texttt{MvNormal}(\texttt{x}^{+} * \beta . + \gamma, \sigma^2 * \texttt{I}) \\ \texttt{end} \\ \texttt{model} = \texttt{linear_regression}(\texttt{x\_data}) \end{array}
```



Standard Bayesian workflow

Define model "in maths"

Define model in Turing.jl

Perform prior checks

Perform inference

Evaluate, e.g. predict

Rinse and repeat!



Standard Bayesian workflow

Perform prior checks

model_gen = fix(model, params_gen)
(; y) = rand(model_gen)
chain_gen = sample(model | (y = y,), NUTS(), 1000)

Identifiability





Standard Bayesian workflow

Define model "in maths"

Define model in Turing.jl

Perform prior checks

Perform inference

Evaluate, e.g. predict

Rinse and repeat!

$\begin{split} \sigma^2 &\sim \text{InverseGamma}(3, 4/10) \\ \gamma &\sim \mathcal{N}(0, 10) \\ \beta &\sim \mathcal{N}(0, \sigma^2 I) \\ \log y_i &\sim \mathcal{N}(\beta \cdot x_i + \gamma, \sigma^2) \text{for } i = 1, \dots, N \end{split}$	model_conditioned = model (y = y_data,)
Turing.jl Condition	Perform inference
$ \begin{array}{c} \textbf{Gmodel function linear_regression(x)} \\ \textbf{d} = size(x, 1) \\ \sigma^2 & \neg \text{ InverseGamma}(3, 4 / 10) \\ \gamma & \neg \text{ Normal}(0, 10) \\ \beta & \neg \text{ NvNormal}(zeros(d), \sigma^2 * \textbf{I}) \\ \textbf{y} & \neg \text{ NvNormal}(x^* & \beta & + \gamma, \sigma^2 * \textbf{I}) \\ end \end{array} $	<pre>chain = sample(model_conditioned, NUTS(), 1000)</pre>
<pre>model = linear_regression(x_data)</pre>	Post-inference analysis
	Summary Statistics parameters mean std mcse Symbol Float64 Float64 Float64
	$\begin{array}{cccccccccccccccccccccccccccccccccccc$
\downarrow Prior checks	
<pre>model_gen = fix(model, params_gen) (; y) = rand(model_gen) chain_gen = sample(model (y = y,), NUTS(), 1000)</pre>	
Identifiability	Predict
1	
If necess	ary



If necessary

std

mcse

Float64

0.0000

0.0066

0.0022

0.0022

Standard Bayesian workflow

Define model "in maths"

Define model in Turing.jl

Perform prior checks

Perform inference

Evaluate, e.g. predict

Rinse and repeat!

$\sigma^2 \sim \text{InverseGamma}(3, 4/10)$ $\gamma \sim \mathcal{N}(0, 10)$ $\beta \sim \mathcal{N}(0, \sigma^2 I)$	model_conditioned = model (y = y_data,)
$\begin{tabular}{ c c c c c } \hline & & & & & & \\ \hline & & & & & & \\ \hline & & & &$	Perform inference
σ^{*} Cancel function linear_regression(x) $d = size(x, 1)$ $\sigma^{2} - InverseGamma(3, 4 / 10)$ $\gamma - Normal(0, 10)$ $\beta - MvNormal(zeros(d), \sigma^{2} * I)$ $y - MvNormal(z* * \beta + \gamma, \sigma^{2} * I)$	<pre> chain = sample(model_conditioned, NUTS(), 1000)</pre>
end model = linear_regression(x_data)	Post-inference analysis
Prior checks model_gen = fix(model, params_gen) (; y) = rand(model_gen) chain_gen = sample(model (y = y,), NUTS(), 1000)	
Identifiability	Predict
If necess	

Standard Bayesian workflow

Evaluate, e.g. predict

Summary Statis	stics			
parameters	mean	std	mcse	
Symbol	Float64	Float64	Float64	
σ^2	0.0543	0.0032	0.0000	
γ	7.7252	0.3688	0.0066	
β[1]	-0.2270	0.1247	0.0022	
β[2]	0.0133	0.1229	0.0022	





Standard Bayesian workflow

Define model "in maths"

Define model in Turing.jl

Perform prior checks

Perform inference

Evaluate, e.g. predict

Rinse and repeat!



Turing.jl: overview

- Convenience and performance from Julia Similar design philosophy as Julia: accept everything and optimize when possible
 - Models are Julia functions, distributions are Julia, *everything* is Julia!
 - Does mean that c**ertain operations are not possible**, e.g. automatic marginalization



Х

Naive implementation of model might not be the most performant

The Epimap model for COVID-19



[Nicholson et al, 2022]

Turing.jl: in the wild

Mixed IgG Fc immune complexes exhibit blended binding profiles and refine FcR affinity estimates

Zhixin Cyrillus Tan,¹ Anja Lux,² Markus Biburger,² Prabha Varghese,² Stephen Lees,³ Falk Nimmerjahn,² and Aaron S. Meyer^{1,3,4,5,*}

A Framework for Inference and Selection of Cell Signaling Pathway Dynamic Models

Marcelo Batista [⊠], <u>Fabio Montoni</u>, <u>Cristiano Campos</u>, <u>Ronaldo Nogueira</u>, <u>Hugo A. Armelin</u> & <u>Marcelo S.</u> <u>Reis</u> [⊠]

Saminfald Strinov 2022 MJ. 77, No. 2, 181–266 https://doi.org/10.1114/2.378584 The research was finded in which or in part, by URRI and/or Plan S sponsored finaling (details on grants are at the end of the paper). A CC BY 4.0 license is applied to this article arising from this submission, in

Interoperability of Statistical Models in Pandemic Preparedness: Principles and Reality

George Nicholson⁺, Marta Blanglardo⁺, Mark Briers⁺, Peter J. Diggle⁺, Tor Erlend Fjelde⁺, Hong Ge⁺, Robert J. B. Goudle, Radka Jersakova⁺, Ruairidh E. King⁺, Brieuc C. L. Lehmann⁺, Ann-Marie Mallon⁺, Tullia Padeillini⁺, Yee Whye Teh⁺, Chris Hoimes⁺ and Sylvia Richardson⁺

Hicks-Arrow Prices for US Federal Debt 1791-1930

George Hall; Jonathan Payne; Thomas J. Sargent; Bálint Szőke§

APPLIED RESEARCH

Bayesian Inference for Thermal Model of Synchronous Generator—Part I: Parameter Estimation

MADHUSUDHAN PANDEY AND BERNT LIE⁶⁰, (Member, IEEE) Telemark Modeling and Control Center (TMCC), 3901 Persgrunn, Norway Corresponding author: Bernt Lie (bernt.lie@usn.no)

Correlating tau pathology to brain atrophy using a physics-based Bayesian model

Amelie Schäfer¹[©] · Pavanjit Chaggar² · Alain Goriely² · Ellen Kuhl¹ · the Alzheimer's Disease Neuroimaging Initiative

Received: 1 December 2021 / Accepted: 4 April 2022 / Published online: 7 June 2022 © The Author(s), under exclusive licence to Springer-Verlag London Ltd., part of Springer Nature 2022

Digital twinning of cardiac electrophysiology for congenital heart disease

Matteo Salvador ¹ ² ³, Fanwei Kong ² ³, Mathias Peirlinck ⁴, David W Parker ⁵, Henry Chubb ³, Anne M Dubin ³, Alison Lesley Marsden ⁶ ¹ ² ³

Turing.jl: in the wild





Digital twinning of cardiac electrophysiology for congenital hear disease Matteo Sak-dar 1/2/3, Fanwer Kong Anne M Dubin⁽³⁾, Alison Lesley Marsder ⁶ (1/2/3)



David Widmann

PhD (Uppsala University) Now at Pumas AI



Hong Ge

Senior Research Fellow (University of Cambridge)



Cameron Pfiffer

PhD (University of Oregon) Now postdoc at Stanford



Kai Xu

PhD (University of Edinburgh) Now at MiT-IBM Watson AI Lab



Mohamed Tarek

PhD (University of Sidney) RA (University of Sideny) and Pumas AI



Tor Erlend Fjelde

PhD student (University of Cambridge)



Martin Trapp

Postdoc (Aalto University)



Qingliang Zhuo

Beijing Paoding Technology Co., LTD.



Seth Axen

ML Research Engineer (University of Tübingen)



Will Tebbutt

Postdoc (University of Cambridge & ATI)



Markus Hauru

Research Engineer (Alan Turing Institute)



Kyurae Kim

PhD student (University of Pennsylvania)

And many other contributors to Turing.jl and the Julia ecosystem!

Thank you for listening!