

Automatic Reparameterisation in Probabilistic Programming

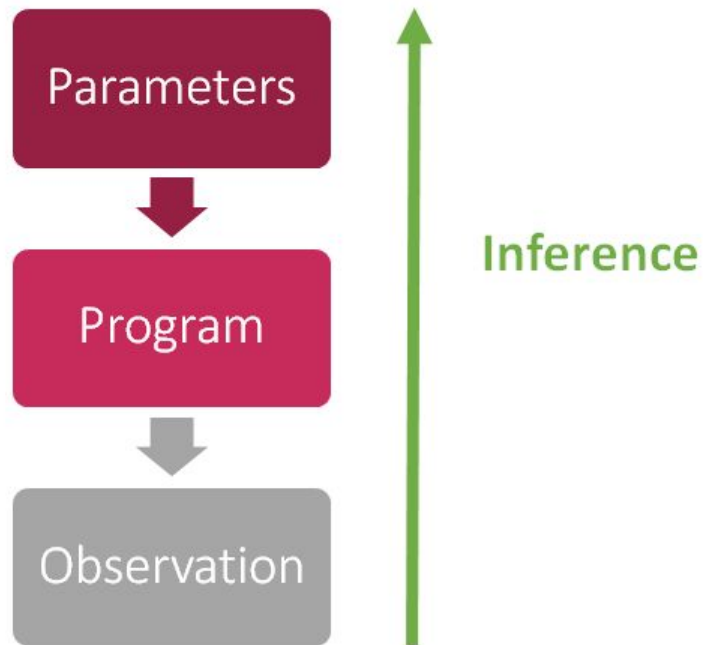
Maria I. Gorinova, Dave Moore, Matthew D. Hoffman

A probabilistic program: Schools

```
mu ~ normal(0, 5)
tau ~ halfCauchy(0, 5)



for(n in 1 .. 3)
  theta[n] ~ normal(mu, tau)
  y[n] ~ normal(theta, sigma[n])

observe y = [28, 8, -3]
observe sigma = [15, 10, 16]
```




Automatic reparameterisation overview

1. What is reparameterisation and why is it difficult?
2. Reparameterisation in probabilistic programming
3. Variationally inferred parameterisation (VIP)




Problem: The posterior geometry affects quality of inference

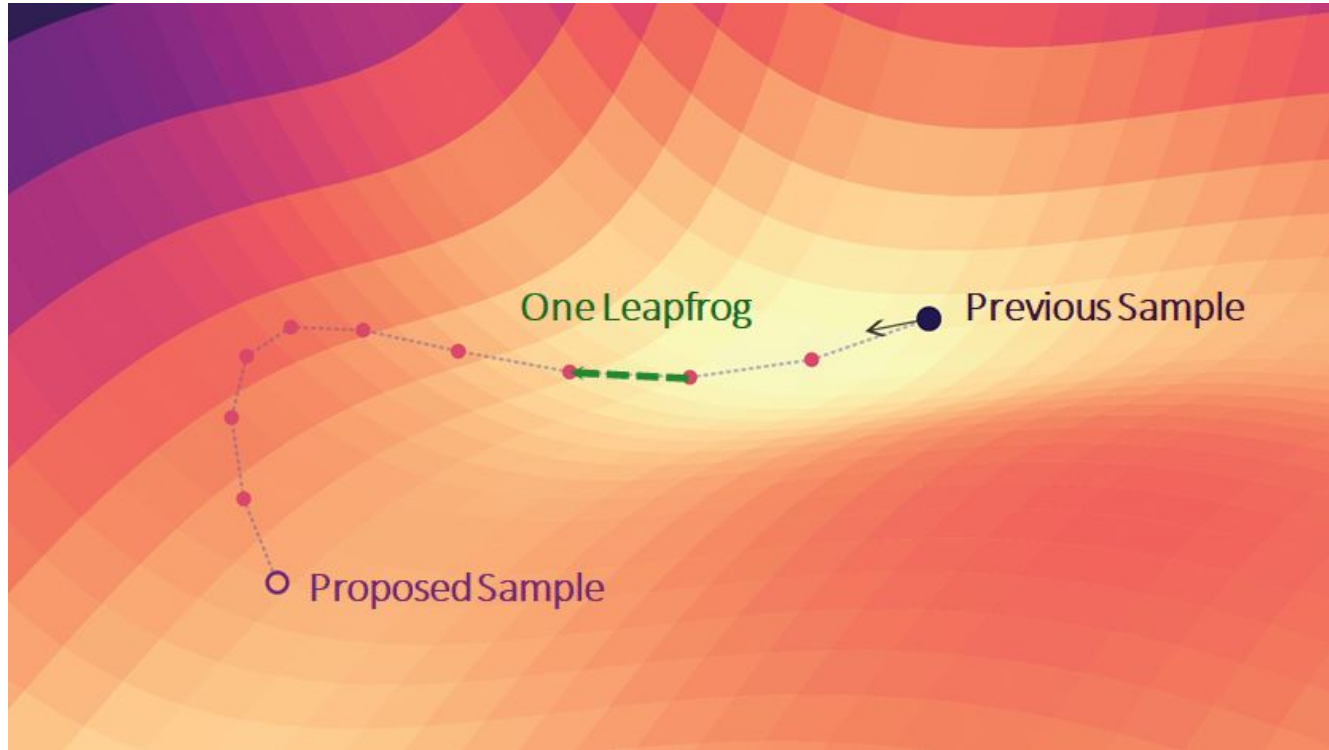
$$y \sim \mathcal{N}(0, 1)$$
$$x \sim \mathcal{N}(y, 0.1)$$



$$y \sim \mathcal{N}(0, 3)$$
$$x \sim \mathcal{N}(0, e^{y/2})$$




Hamiltonian Monte Carlo (HMC)




Problem: The posterior geometry affects quality of inference

$$y \sim \mathcal{N}(0, 1)$$
$$x \sim \mathcal{N}(y, 0.1)$$




$$y \sim \mathcal{N}(0, 3)$$
$$x \sim \mathcal{N}(0, e^{y/2})$$



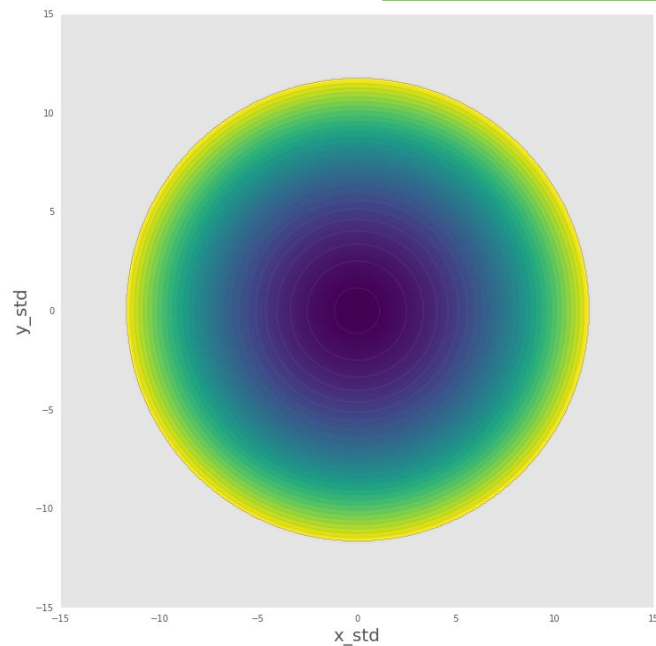
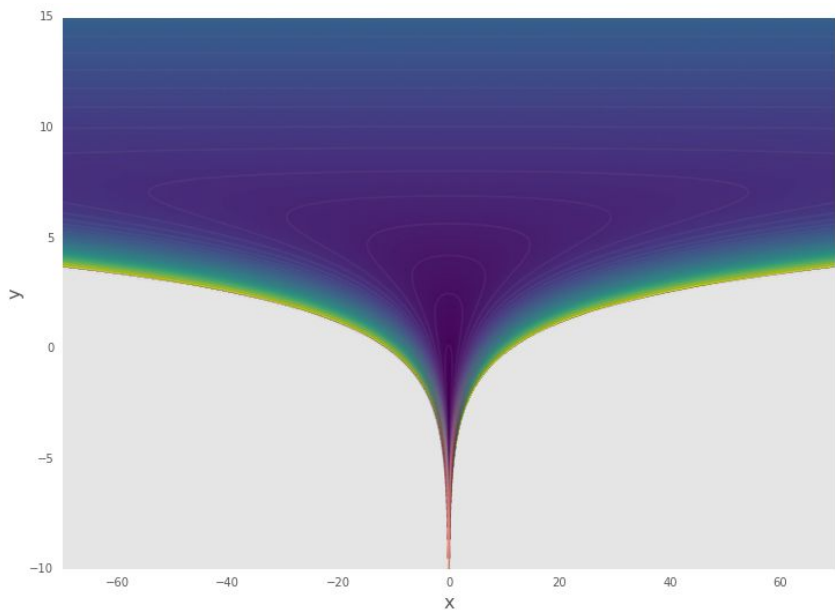
What is model reparameterisation?

$$y \sim \mathcal{N}(0, 3)$$
$$x \sim \mathcal{N}(0, e^{y/2})$$




$$\tilde{y} \sim \mathcal{N}(0, 1)$$
$$\tilde{x} \sim \mathcal{N}(0, 1)$$

$$y = \tilde{y} \times 3$$
$$x = \tilde{x} \times e^{y/2}$$



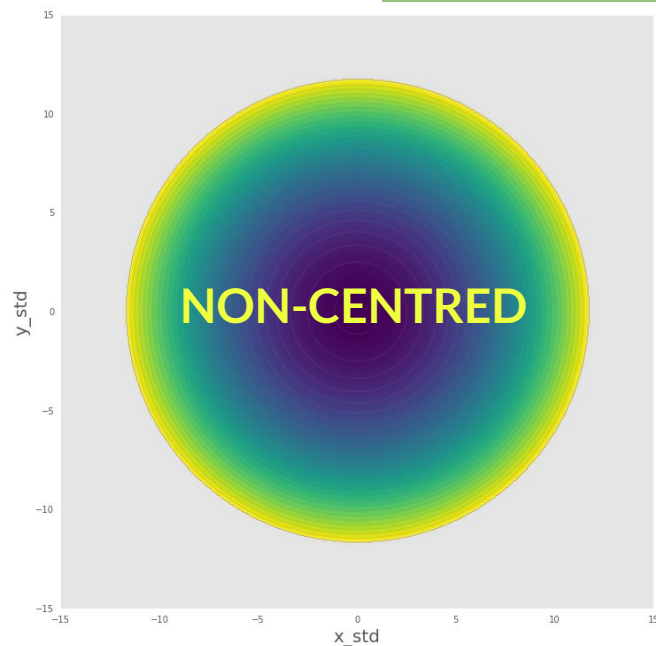
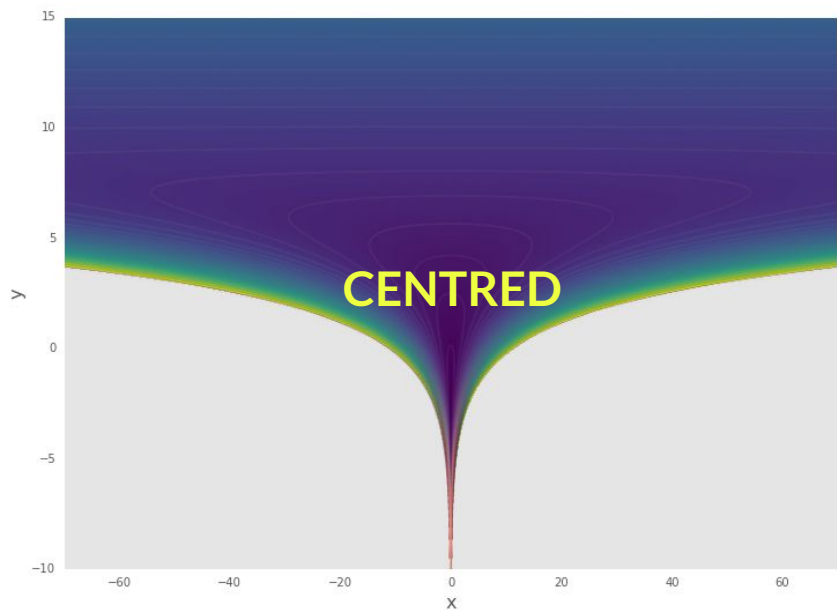
What is model reparameterisation?

$$y \sim \mathcal{N}(0, 3)$$
$$x \sim \mathcal{N}(0, e^{y/2})$$




$$\tilde{y} \sim \mathcal{N}(0, 1)$$
$$\tilde{x} \sim \mathcal{N}(0, 1)$$

$$y = \tilde{y} \times 3$$
$$x = \tilde{x} \times e^{y/2}$$




Understanding Reparameterisation Effects

Centred



$$\theta \sim \mathcal{N}(0, 1) \quad \mu \sim \mathcal{N}(\theta, \sigma_\mu)$$
$$y_n \sim \mathcal{N}(\mu, \sigma) \text{ for all } n \in 1 \dots N$$

Non-centred




$$\theta \sim \mathcal{N}(0, 1) \quad \epsilon \sim \mathcal{N}(0, 1) \quad \mu = \theta + \sigma_\mu \epsilon$$
$$y_n \sim \mathcal{N}(\theta + \sigma_\mu \epsilon, \sigma) \text{ for all } n \in 1 \dots N$$


Understanding Reparameterisation Effects

$$q = N/\sigma$$

Centred




Non-centred




Understanding Reparameterisation Effects

$$q = N/\sigma$$

Centred




Non-centred




Understanding Reparameterisation Effects

$$q = N/\sigma$$

Centred



Non-centred



Stan diagnostics example

Goal: Free modellers of the need to
choose model parameterisation

Reparameterisation in probabilistic programming

Centred

```
def model(N, sigma, sigma_mu):  
  
    theta = Normal(0., 3.)  
  
    mu = Normal(theta, sigma_mu)  
  
    y = Normal(mu, sigma)  
    return y
```

Non-centred

```
def model_ncp(N, sigma, sigma_mu):  
  
    theta_std = Normal(0., 1.)  
    theta = 3. * theta_std  
  
    mu_std = Normal(0., 1.)  
    mu = theta + mu_std * sigma_mu  
  
    y = Normal(mu, sigma)  
    return y
```

Algebraic effect handlers

```
handler h {  
  read(_) -> "I <3 PPLs"  
}
```

```
x = read(file1)  
y = read(file2)  
print(concatenate(x, y))
```

```
with h handle:  
  x = read(file1)  
  y = read(file2)  
  print(concatenate(x, y))
```

> "Contents of file1Contents of file2"

> "I <3 PPLsI <3 PPLs"

Algebraic effect handlers

```
handler h {  
  v ~ dist →  
    print("I <3 PPLs")  
    v = dist.sample()  
}
```

```
x ~ normal(0, 1)  
y ~ normal(0, 1)  
print(x * y)
```

with h **handle**:

```
x ~ normal(0, 1)  
y ~ normal(0, 1)  
print(x * y)
```

> -1.891424147896625

> "I <3 PPLs"

> "I <3 PPLs"

> -1.891424147896625

Reparameterisation in probabilistic programming

Effect Handler: A function that may change how and if a RV is constructed

```
def ncp(rv_constr, **args):  
    if is_location_scale(rv_constr):  
        std = rv_constr(0., 1.)  
        return args["scale"] * std + args["loc"]
```

```
with handler(ncp):  
    theta = Normal(0., 3.)  
    mu = Normal(theta, 1.)
```

Reparameterisation in probabilistic programming

Effect Handler: A function that may change how and if a RV is constructed

```
def ncp(rv_constr, **args):  
    if is_location_scale(rv_constr):  
        std = rv_constr(0., 1.)  
        return args["scale"] * std + args["loc"]
```

```
with handler(ncp):  
    theta = Normal(0., 3.)  
    mu = Normal(theta, 1.)  
  
    theta_std = Normal(0., 1.)  
    theta = 3. * theta_std  
    mu_std = Normal(0., 1.)  
    mu = mu_std + theta
```

Reparameterisation in probabilistic programming

Centred



Non-centred

```
def model(N, sigma, sigma_mu):
```

```
    theta = Normal(0., 3.)
```

```
    mu = Normal(theta, sigma_mu)
```

```
    y = Normal(mu, sigma)
```

```
    return y
```

```
def model_ncp(N, sigma, sigma_mu):
```

```
    theta_std = Normal(0., 1.)
```

```
    theta = 3. * theta_std
```

```
    mu_std = Normal(0., 1.)
```

```
    mu = theta + mu_std * sigma_mu
```

```
    y = Normal(mu, sigma)
```

```
    return y
```

Variationally Inferred Parameterisation (VIP)

$$\tilde{z} \sim \mathcal{N}(\lambda\mu, \sigma^\lambda)$$

$$z = \mu + \sigma^{1-\lambda}(\tilde{z} - \lambda\mu)$$

Variationally Inferred Parameterisation (VIP)

$$\tilde{z} \sim \mathcal{N}(\lambda\mu, \sigma^\lambda)$$


$$z = \mu + \sigma^{1-\lambda}(\tilde{z} - \lambda\mu)$$

$$\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\lambda}) = \mathbb{E}_{q(\tilde{\mathbf{z}}; \boldsymbol{\theta})} (\log p(\mathbf{x}, \tilde{\mathbf{z}}; \boldsymbol{\lambda}) - \log q(\tilde{\mathbf{z}}; \boldsymbol{\theta}))$$

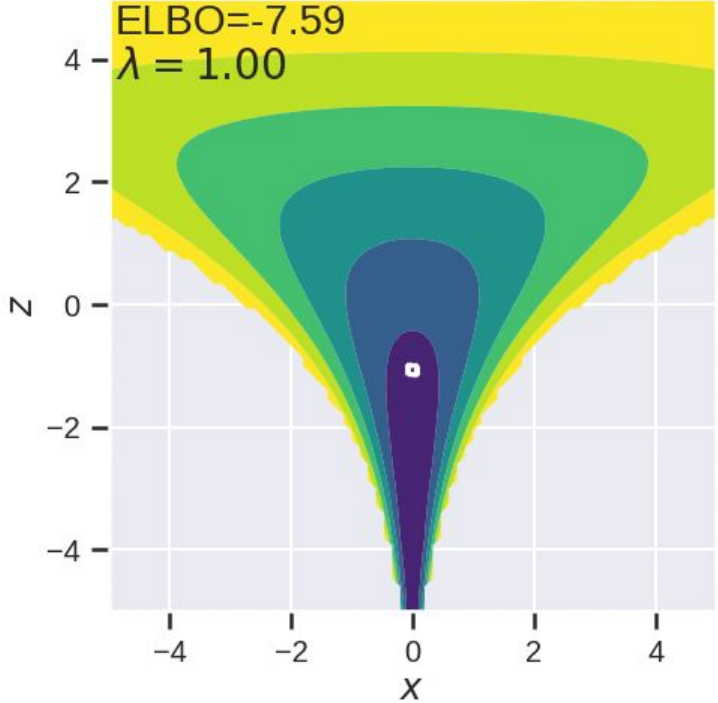
Example: Neal's funnel VIP



$$z \sim \mathcal{N}(0, 3) \quad x \sim \mathcal{N}(0, e^{z/2})$$



Reparameterized coordinates:





Original coordinates:










Automatic Reparameterisation in Probabilistic Programming

Maria I. Gorinova, Dave Moore, Matthew D. Hoffman



 <https://git.io/fpKeO>

M.Gorinova@gmail.com

 [@migorinova](https://twitter.com/migorinova)

Thank you!