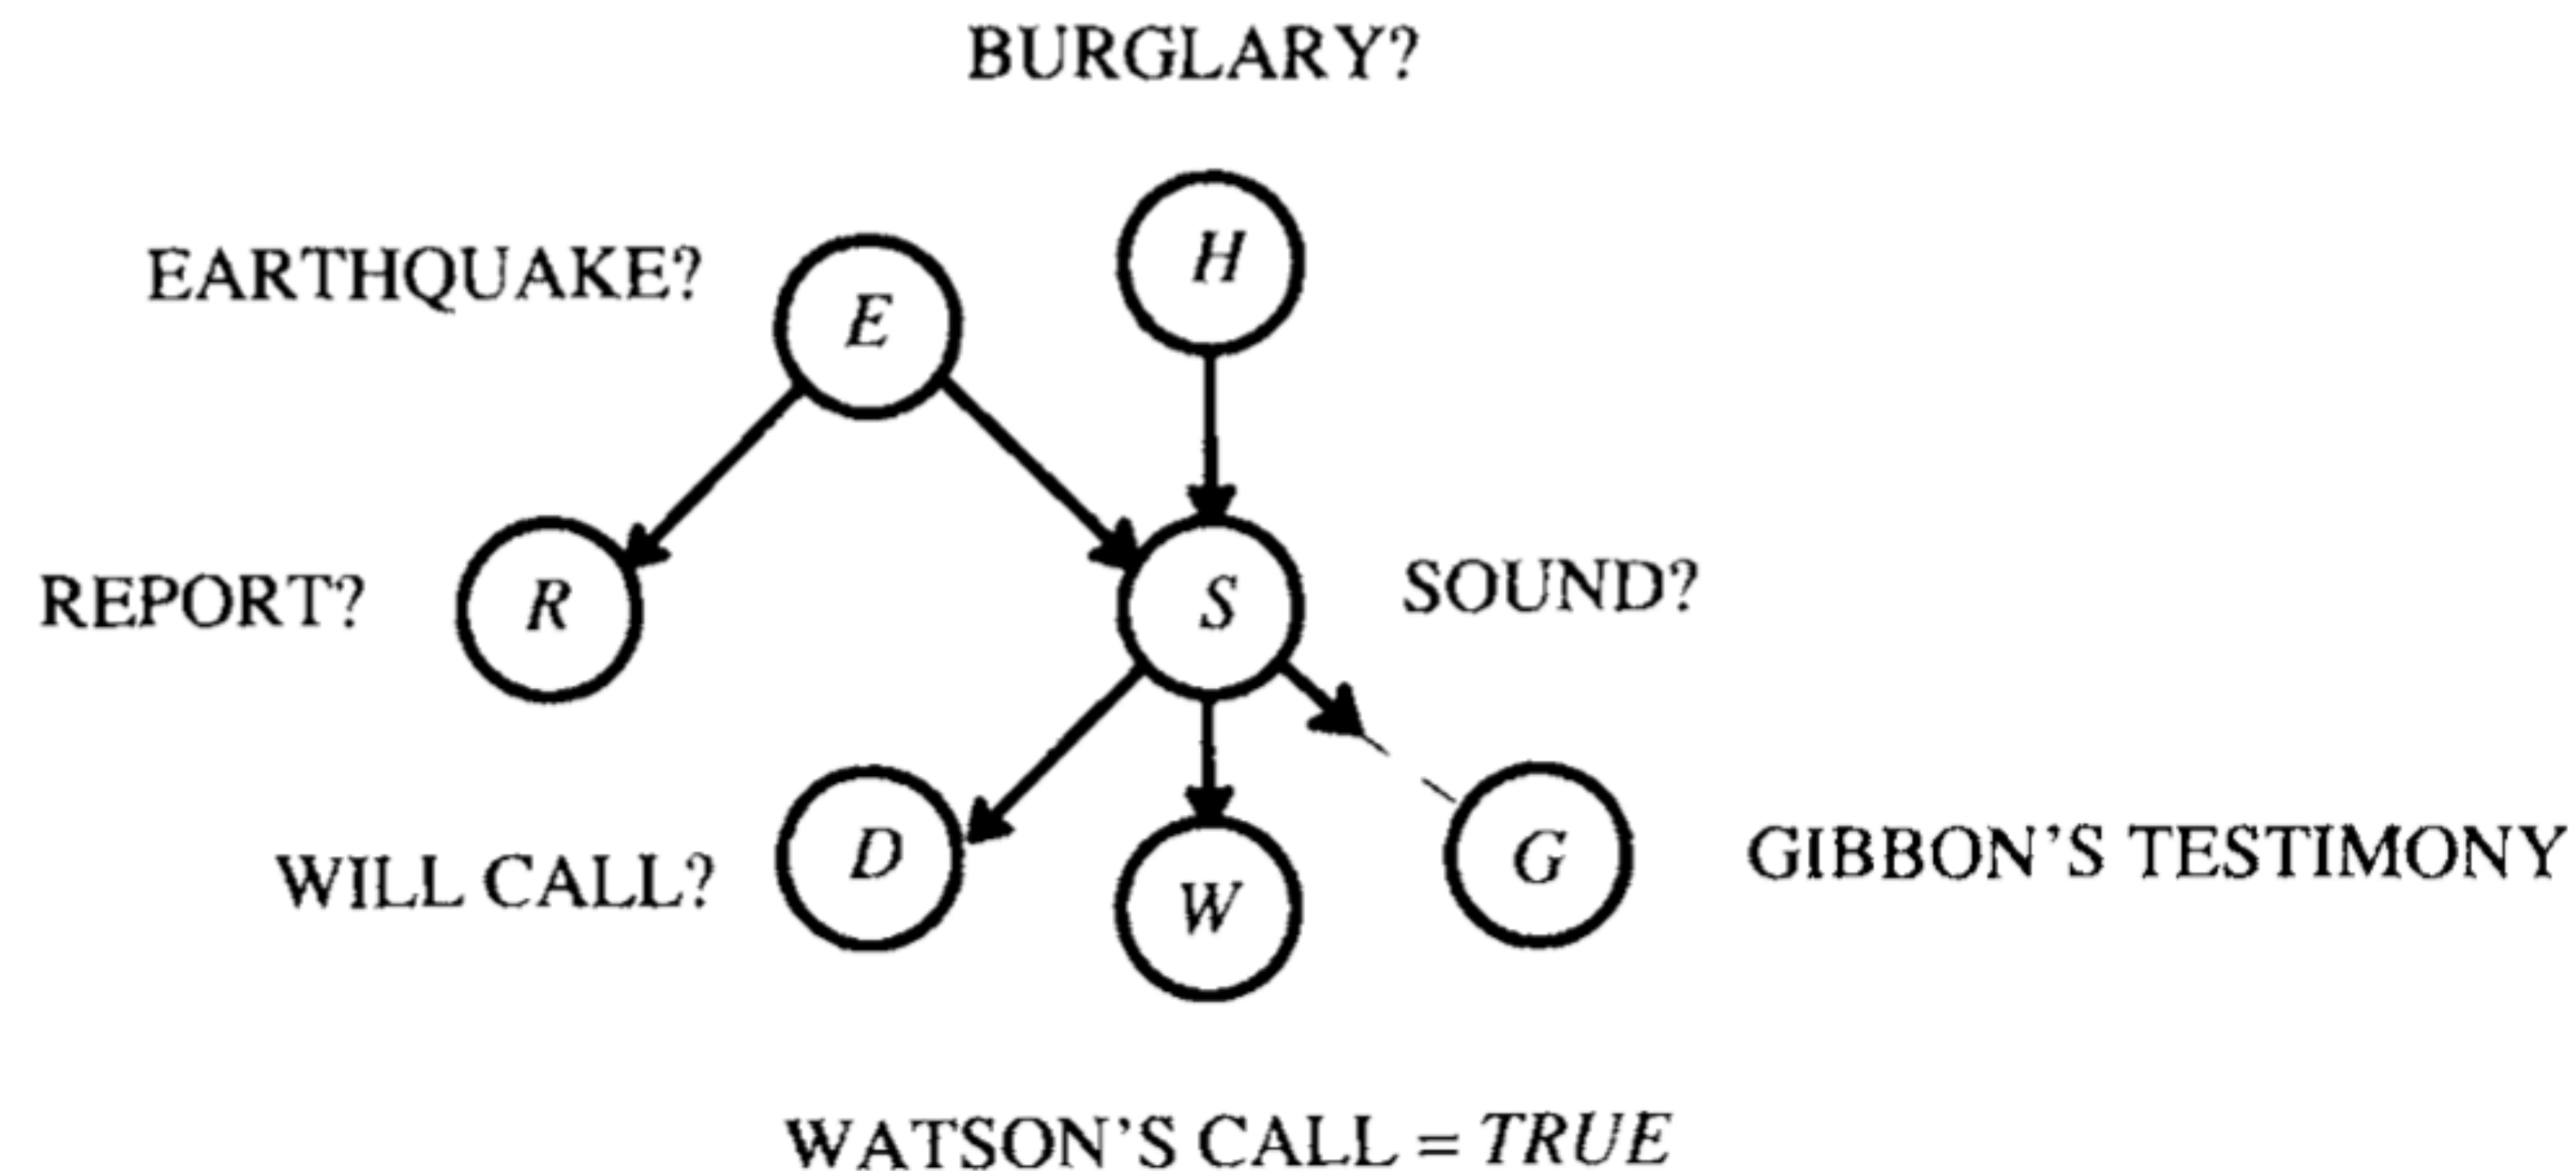



# Graphical Models and Simulation-Based Inference

Graphical Models: Discrete Inference and Learning

# Introduction to DAG and their relationship with Probability Functions (Pearl)

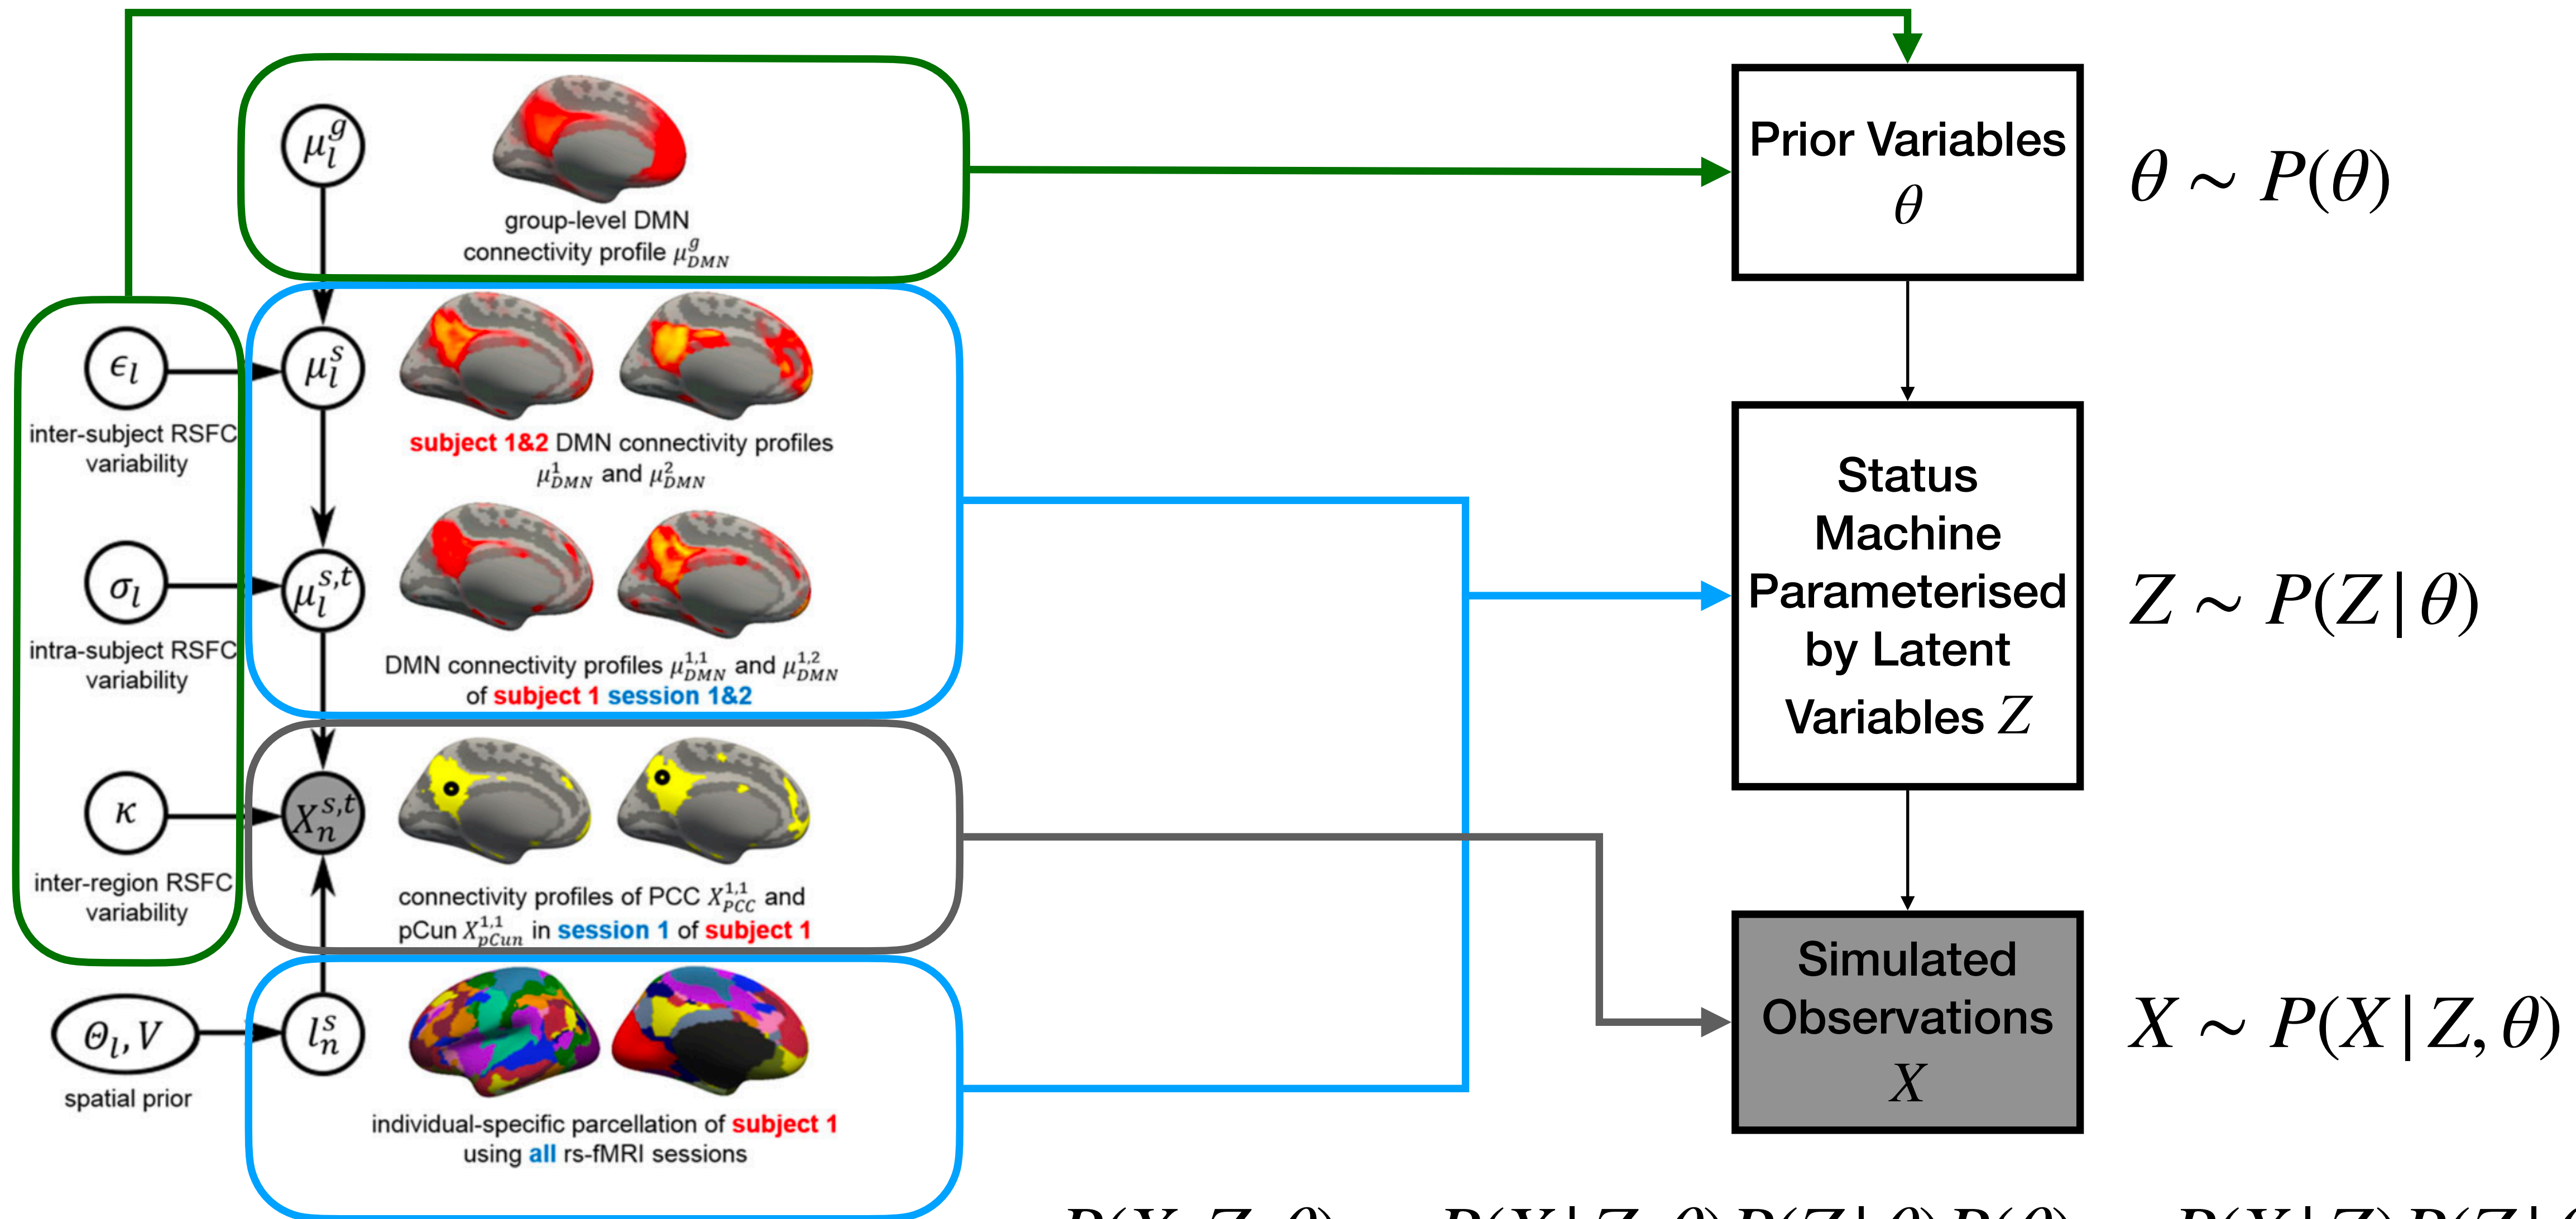


[Pearl 1987]



[Kong et al 2019]

# Graphical Models and Simulation Systems



$$P(X, Z, \theta) \stackrel{=}{=} P(X | Z, \theta)P(Z | \theta)P(\theta) = P(X | Z)P(Z | \theta)P(\theta)$$

# General Inference Notation

$\theta$ : parameters     $X$ : observations

$$P(\theta | X) = \frac{\overset{\text{Likelihood}}{P(X | \theta)} \overset{\text{Prior}}{P(\theta)}}{\underset{\text{Evidence}}{P(X)}}$$


Posterior

$Z$ : latent random variables

$$P(\theta | X) = \frac{\mathbb{E}_Z [P(X | Z, \theta)] P(\theta)}{P(X)}$$

$\eta$ : nuisance random variables

$$P(\theta | X) = \frac{\mathbb{E}_\eta [P(X | \theta, \eta)] P(\theta)}{P(X)}$$



Intractable in general:  
full likelihood impossible to evaluated or computation cost is extremely high

# Likelihood computation is hard: Enter Mechanistic, Example Models Galton Board

$\theta$ : parameters     $X$ : observations

$$P(\theta | X) = \frac{\overset{\text{Likelihood}}{P(X | \theta)} \overset{\text{Prior}}{P(\theta)}}{\underset{\text{Evidence}}{P(X)}}$$

Posterior

$Z$ : latent random variables


$$P(\theta | X) = \frac{\mathbb{E}_Z [P(X | \mathbf{Z}, \theta)] P(\theta)}{P(X)}$$

$\eta$ : nuisance random variables

$$P(\theta | X) = \frac{\mathbb{E}_\eta [P(X | \theta, \eta)] P(\theta)}{P(X)}$$



# Simulation-Based Inference



- Inference is defined as finding the  $\theta$  that could be at the origin of an observation  $X$ . Specifically computing  $P(\theta | X) = \mathbb{E}_Z[P(\theta, Z | X)]$
- For this, we use Bayes 
$$P(\theta, Z | X) = \frac{P(X | Z, \theta)P(Z, \theta)}{P(X)}$$
, nonetheless the likelihood  $P(X | Z, \theta)$  is often unknown or intractable.
- Hence simulation-based inference either approximates or eliminates the need for an explicit likelihood by simulating observations.

# Simulation-Based Inference: Neural Network Approximations

Prior Variables  
 $\theta$

$$\theta \sim P(\theta)$$

Status  
Machine  
Parameterised  
by Latent  
Variables  $Z$

$$Z \sim P(Z | \theta)$$

Simulated  
Observations  
 $X$

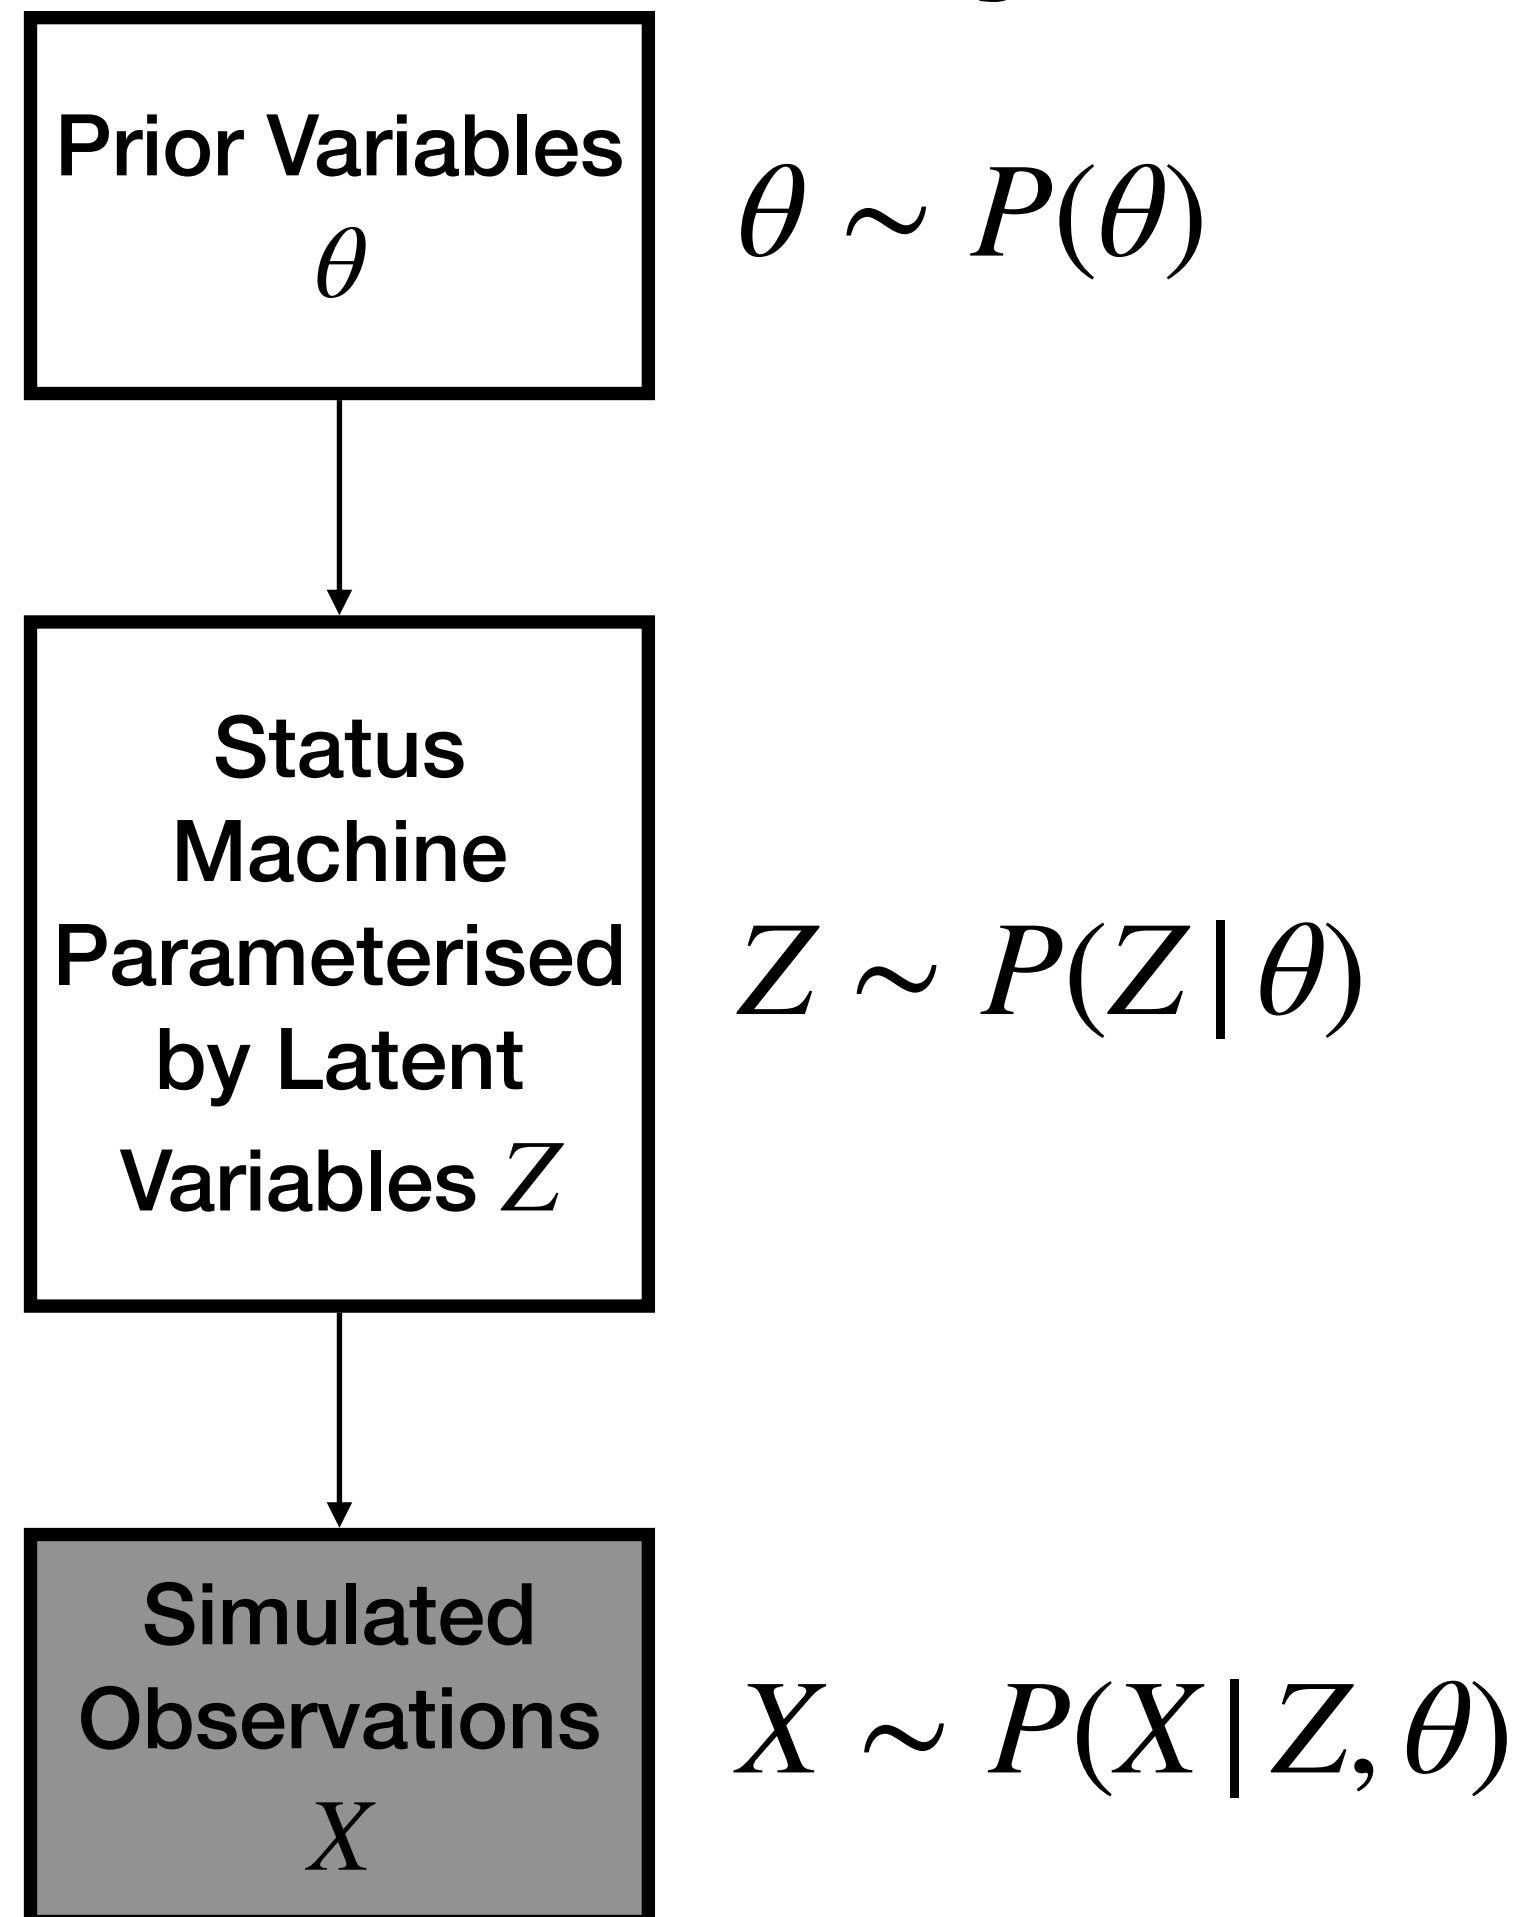
$$X \sim P(X | Z, \theta)$$

$$P(\theta | X) = \frac{\overset{\text{Likelihood}}{P(X | \theta)} \overset{\text{Prior}}{P(\theta)}}{\underset{\text{Evidence}}{P(X)}}$$

$\theta$ : parameters    $X$ : observations

- $P(\theta | X)$  approximated through “Neural Posterior” estimators
- $P(X | \theta)$  approximated through “Neural Likelihood” estimators
- $\frac{P(X | \theta)}{P(X)}$  approximated through the “Neural ratio” estimators

# Simulation-Based Inference: Why now it works (Cranmer et al 2019)



$$P(\theta | X) = \frac{P(X | \theta)P(\theta)}{P(X)}$$

$\theta$ : parameters     $X$ : observations

Likelihood    Prior


Posterior

Evidence

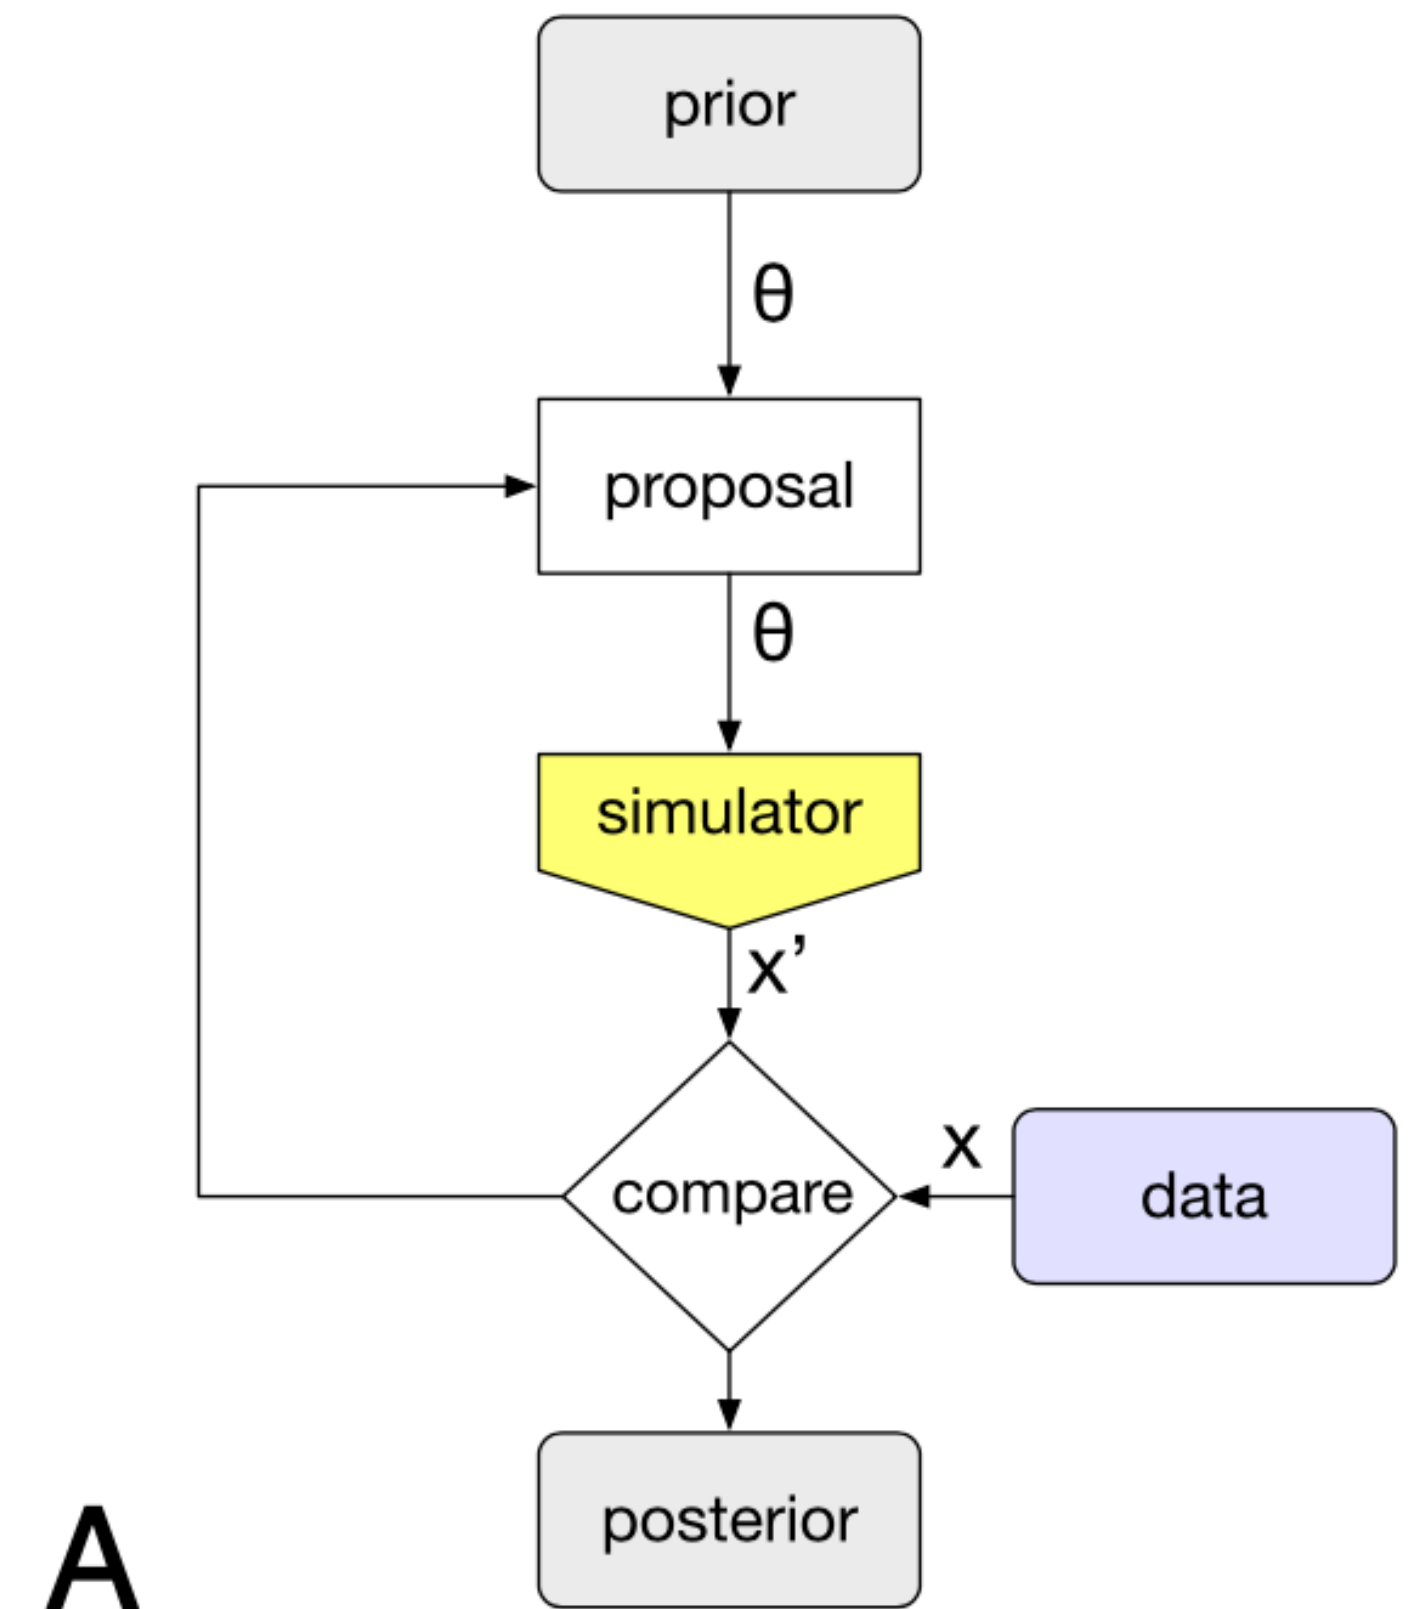
- Novel ML-based approaches allow us to massively generate simulated observations
- Autodifferentiation and neural network approaches are great non-linear function estimators
- Active learning can help improving sampling efficiency much better than Markov Chains



# Simulation-Based Inference: Approximate Bayesian MC




Approximate Bayesian Computation  
with Monte Carlo sampling



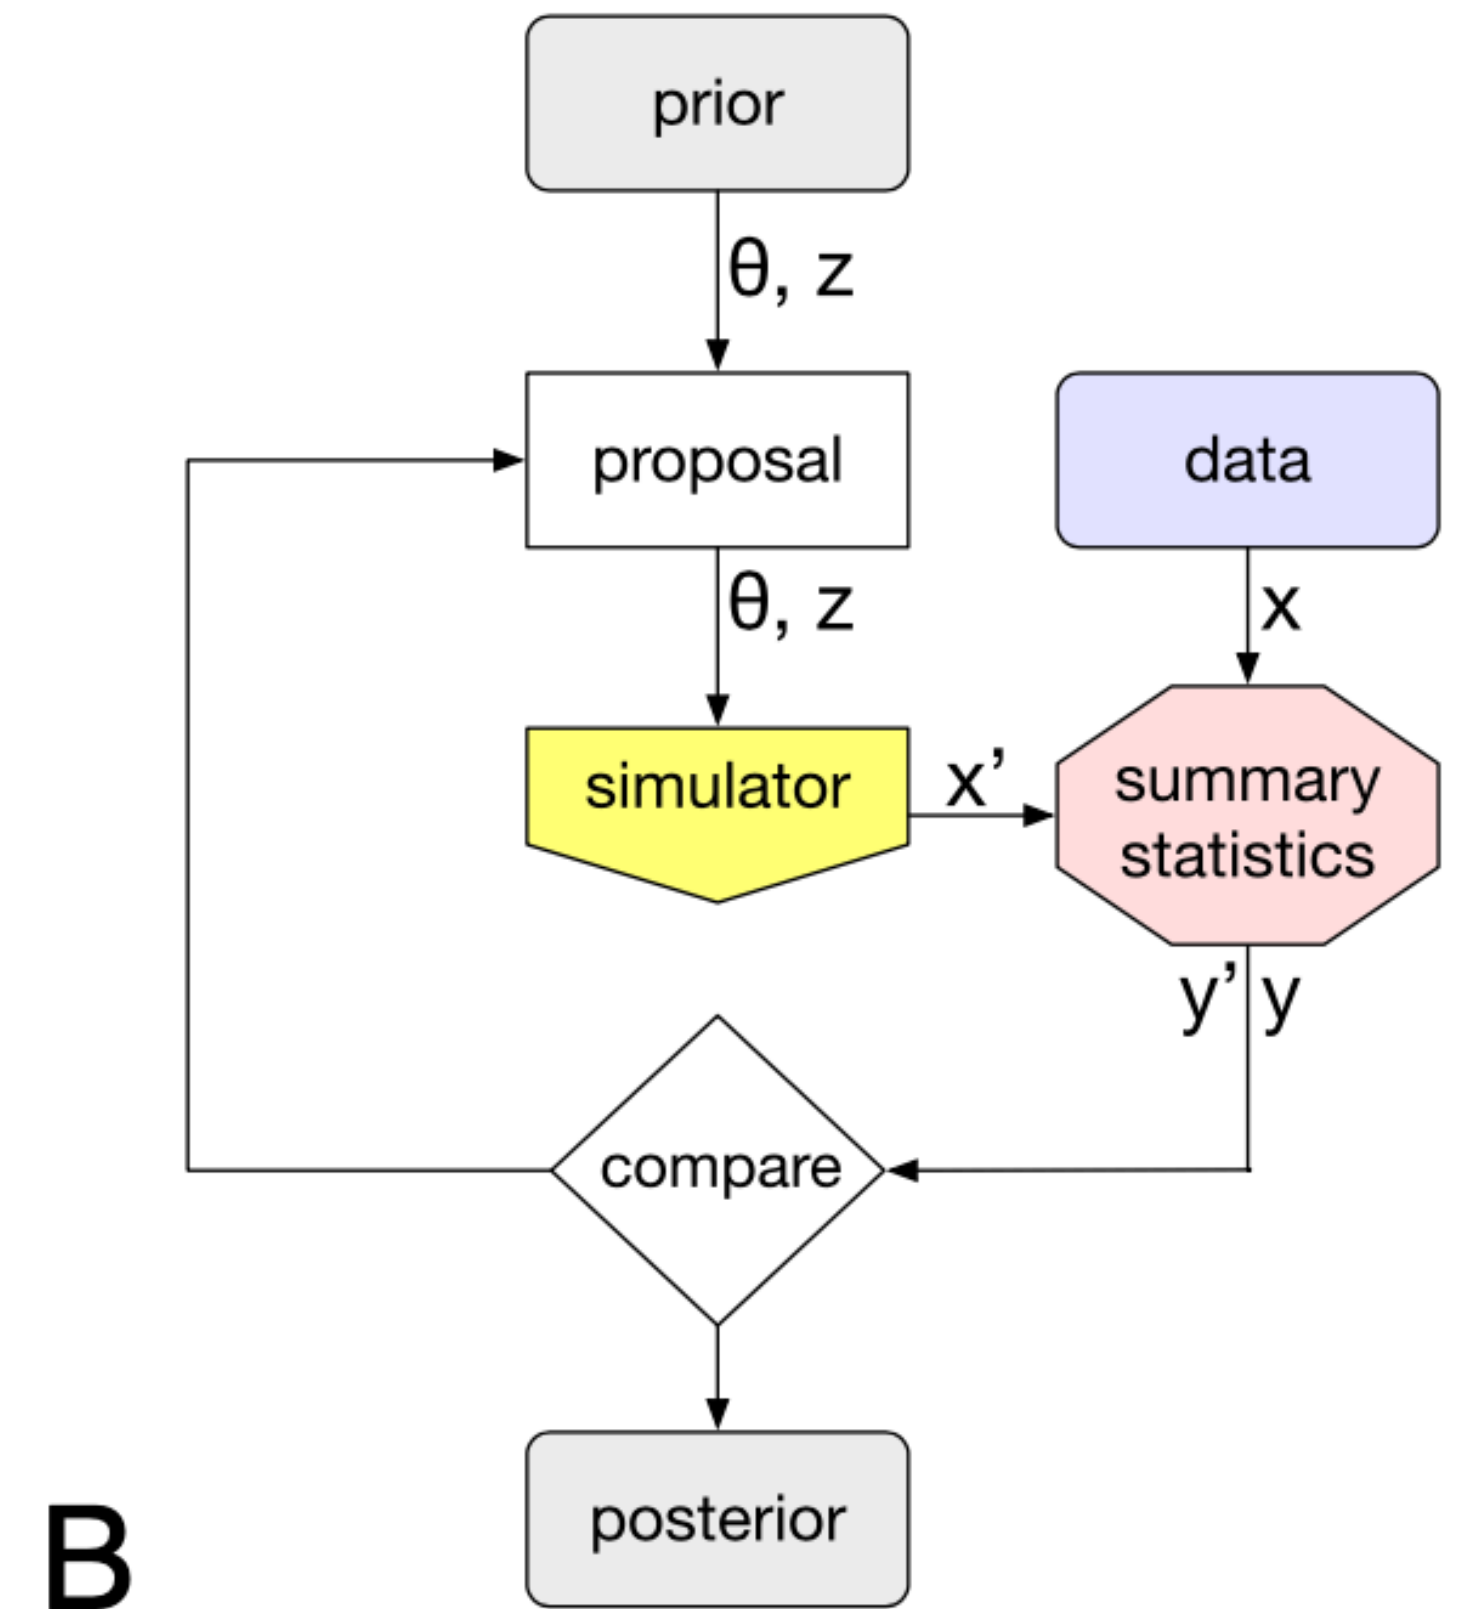
A

(Cranmer et al 2019)

# Simulation-Based Inference: Approximate Bayesian MC




Approximate Bayesian Computation  
with learned summary statistics

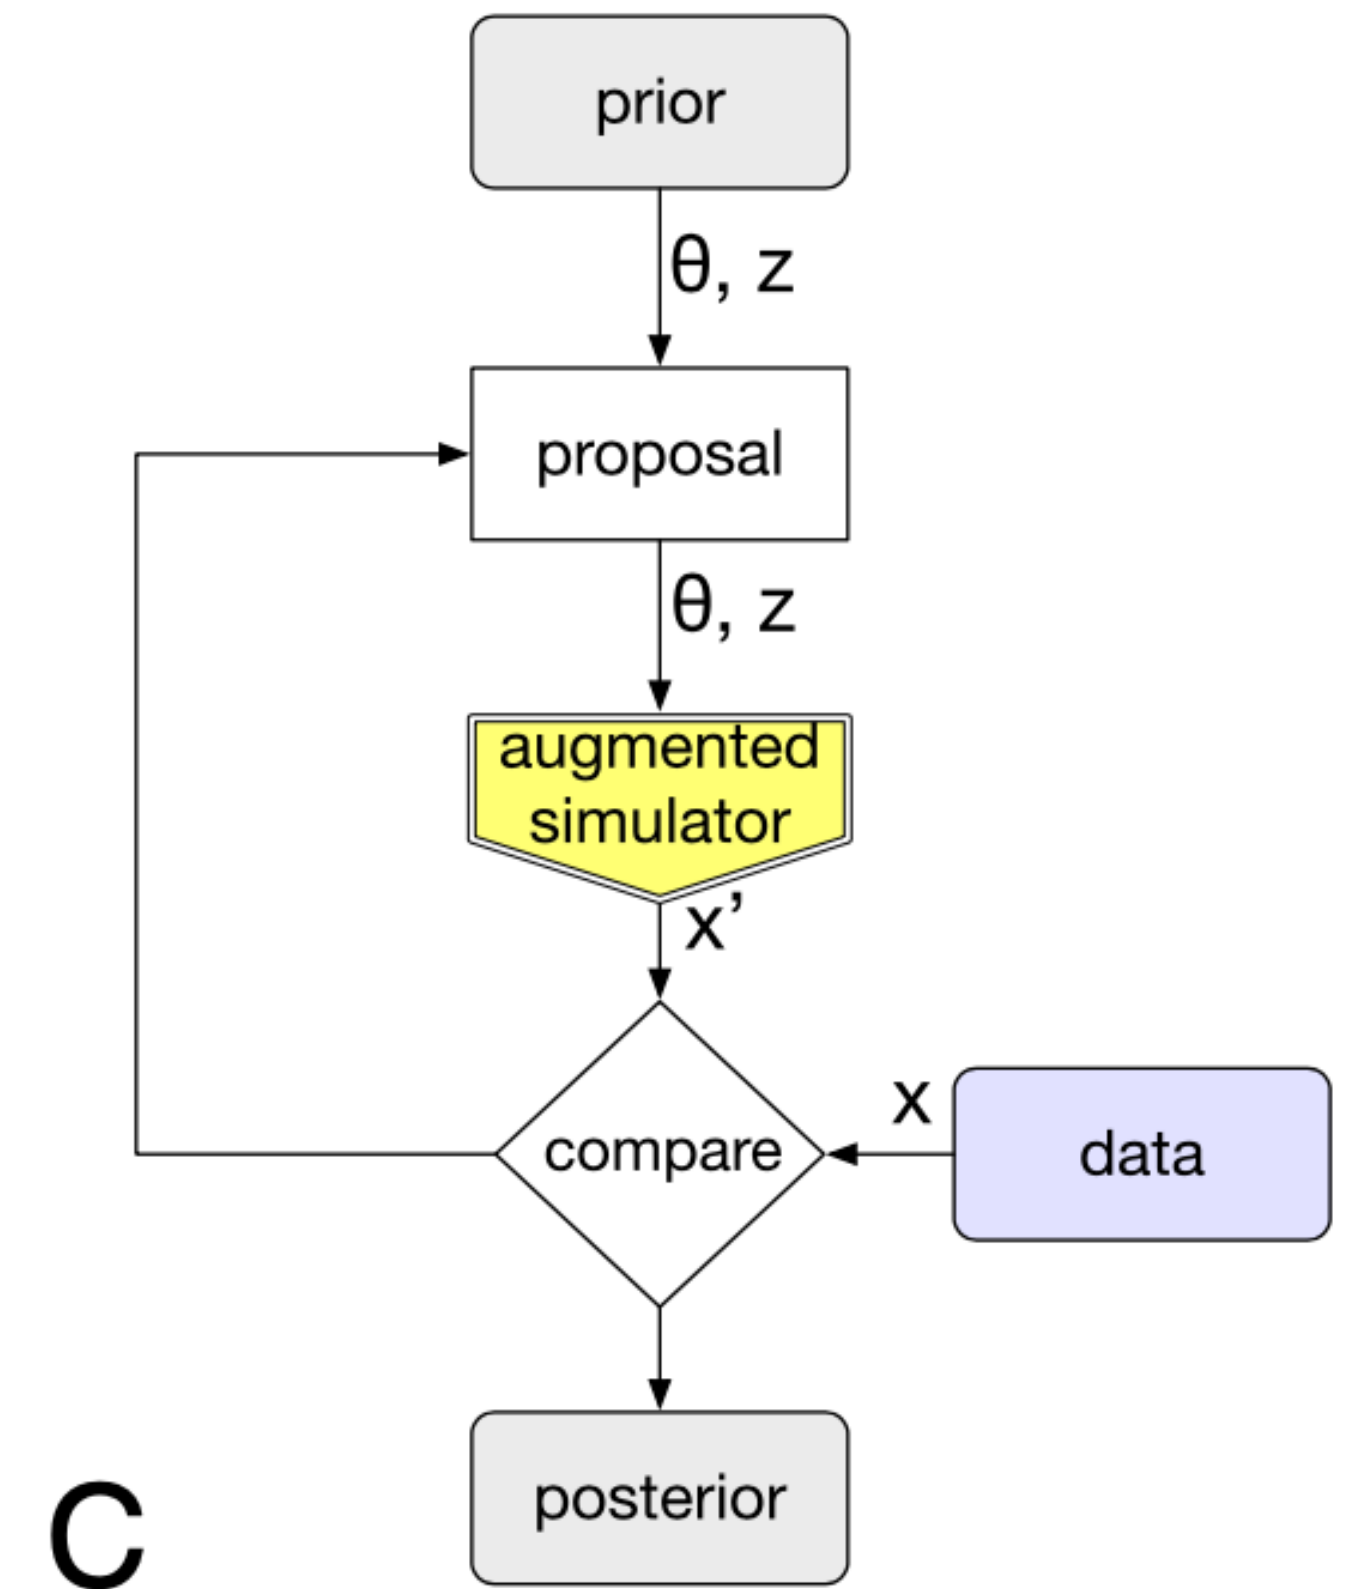


(Cranmer et al 2019)

# Simulation-Based Inference: Approximate Bayesian MC



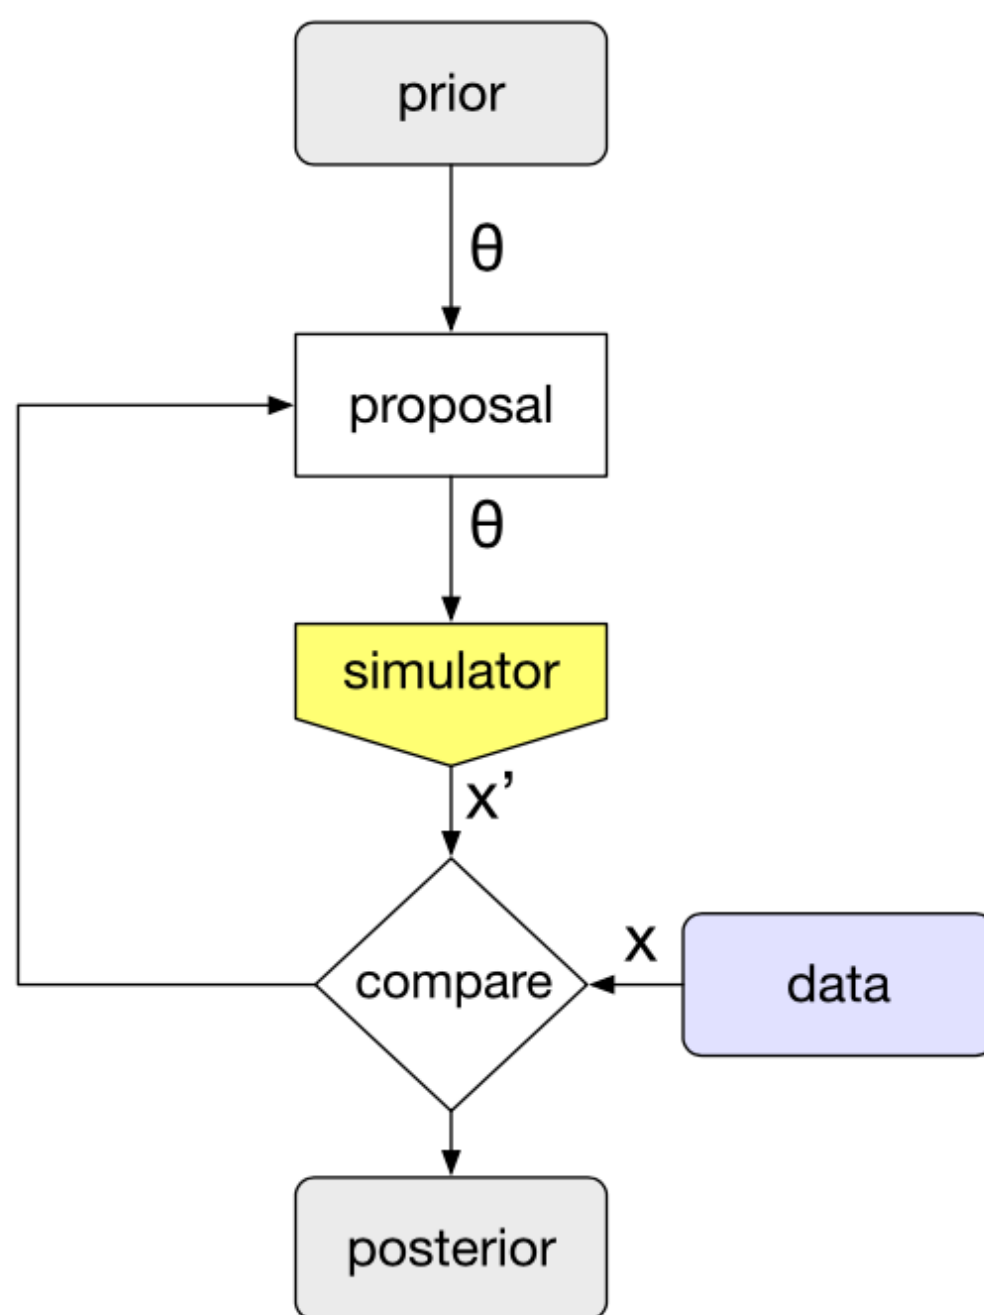
Probabilistic Programming  
with Monte Carlo sampling



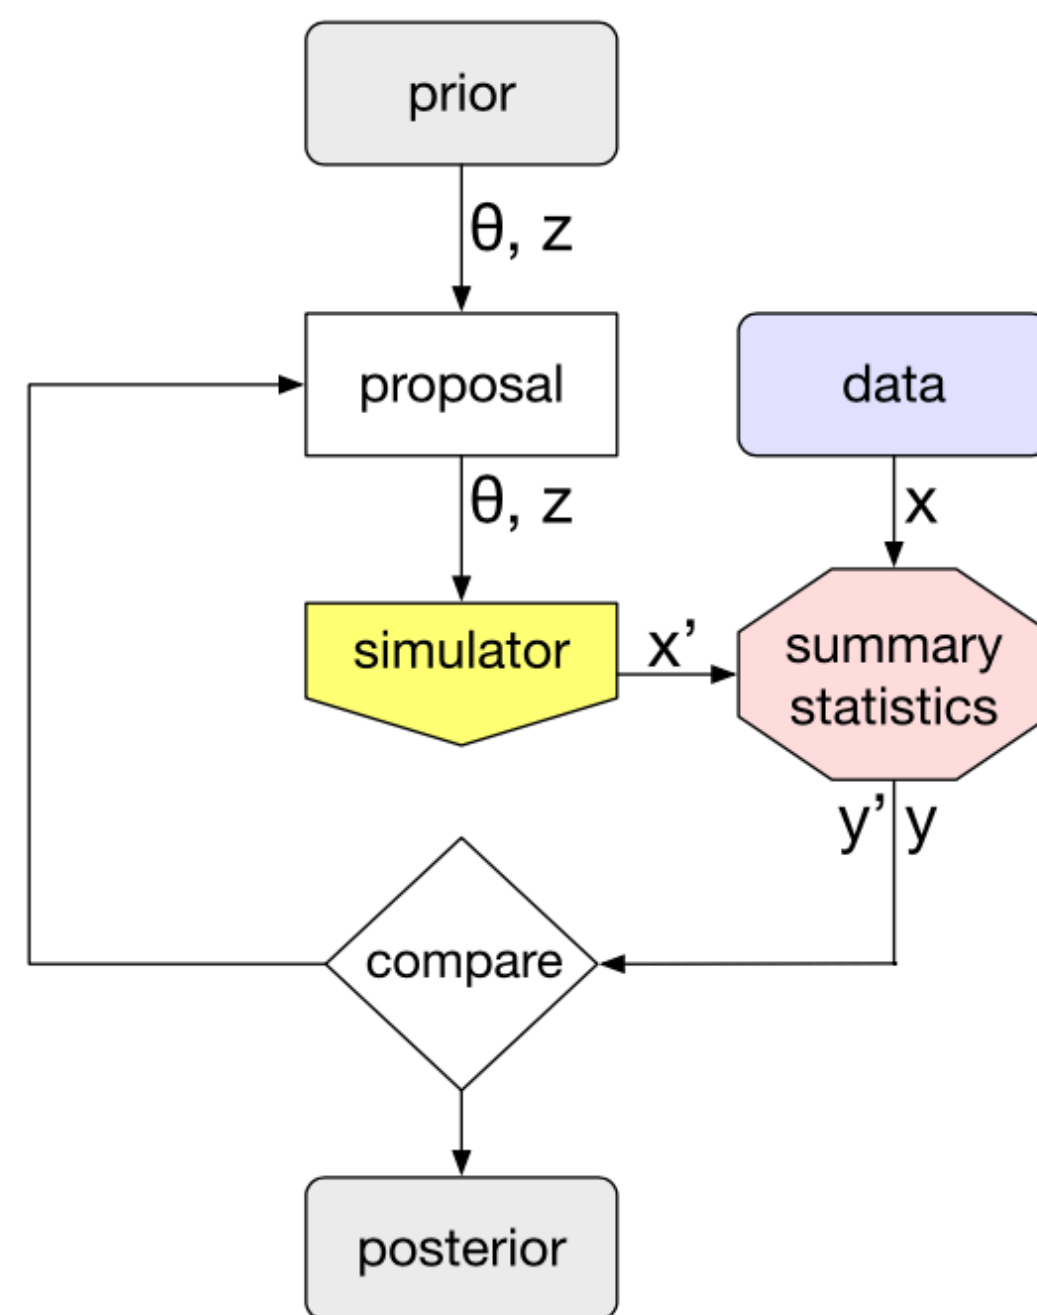
(Cranmer et al 2019)

# Simulation-Based Inference: Approximate Bayesian MC


Approximate Bayesian Computation  
with Monte Carlo sampling




Approximate Bayesian Computation  
with learned summary statistics



Probabilistic Programming  
with Monte Carlo sampling



Probabilistic Programming  
with Inference Compilation



A

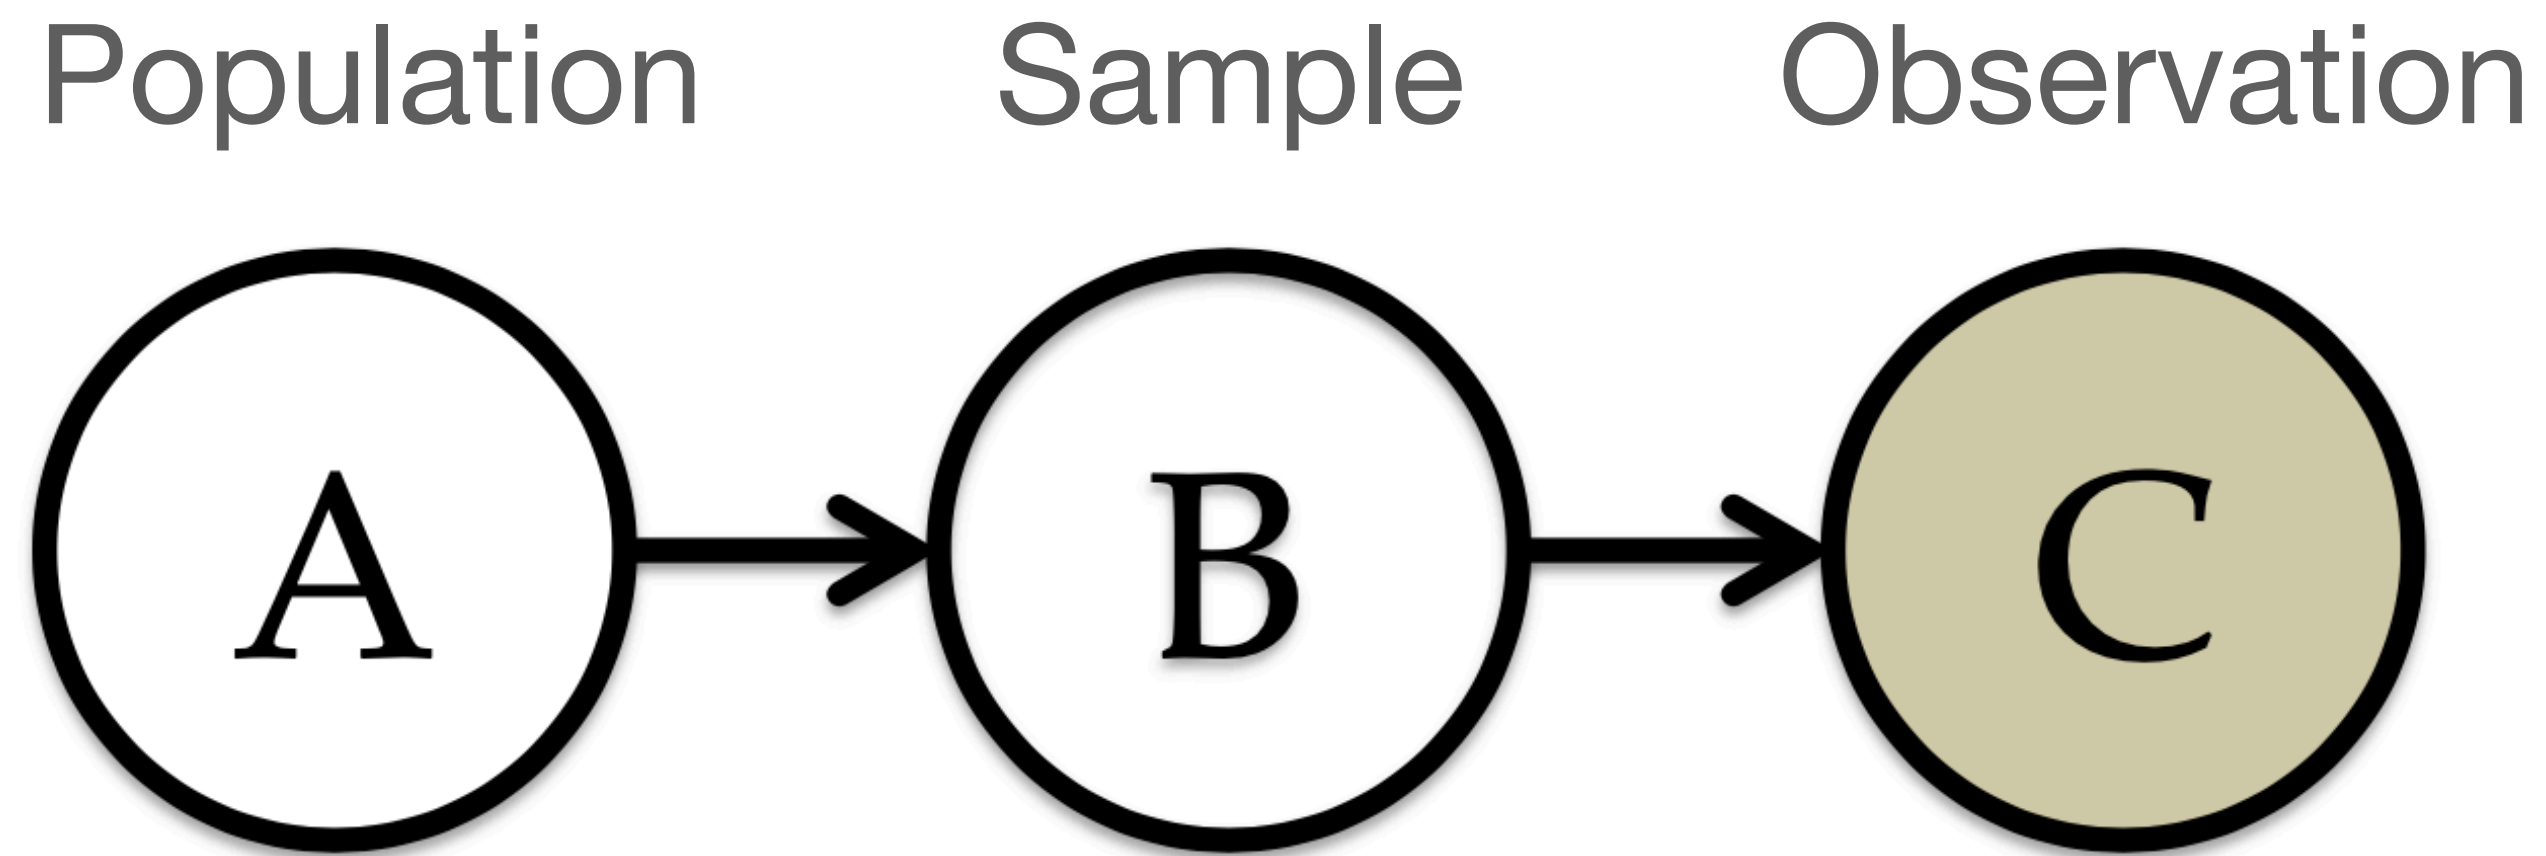
B

C

D

(Cranmer et al 2019)

# Simulation-Based Inference: Amortization

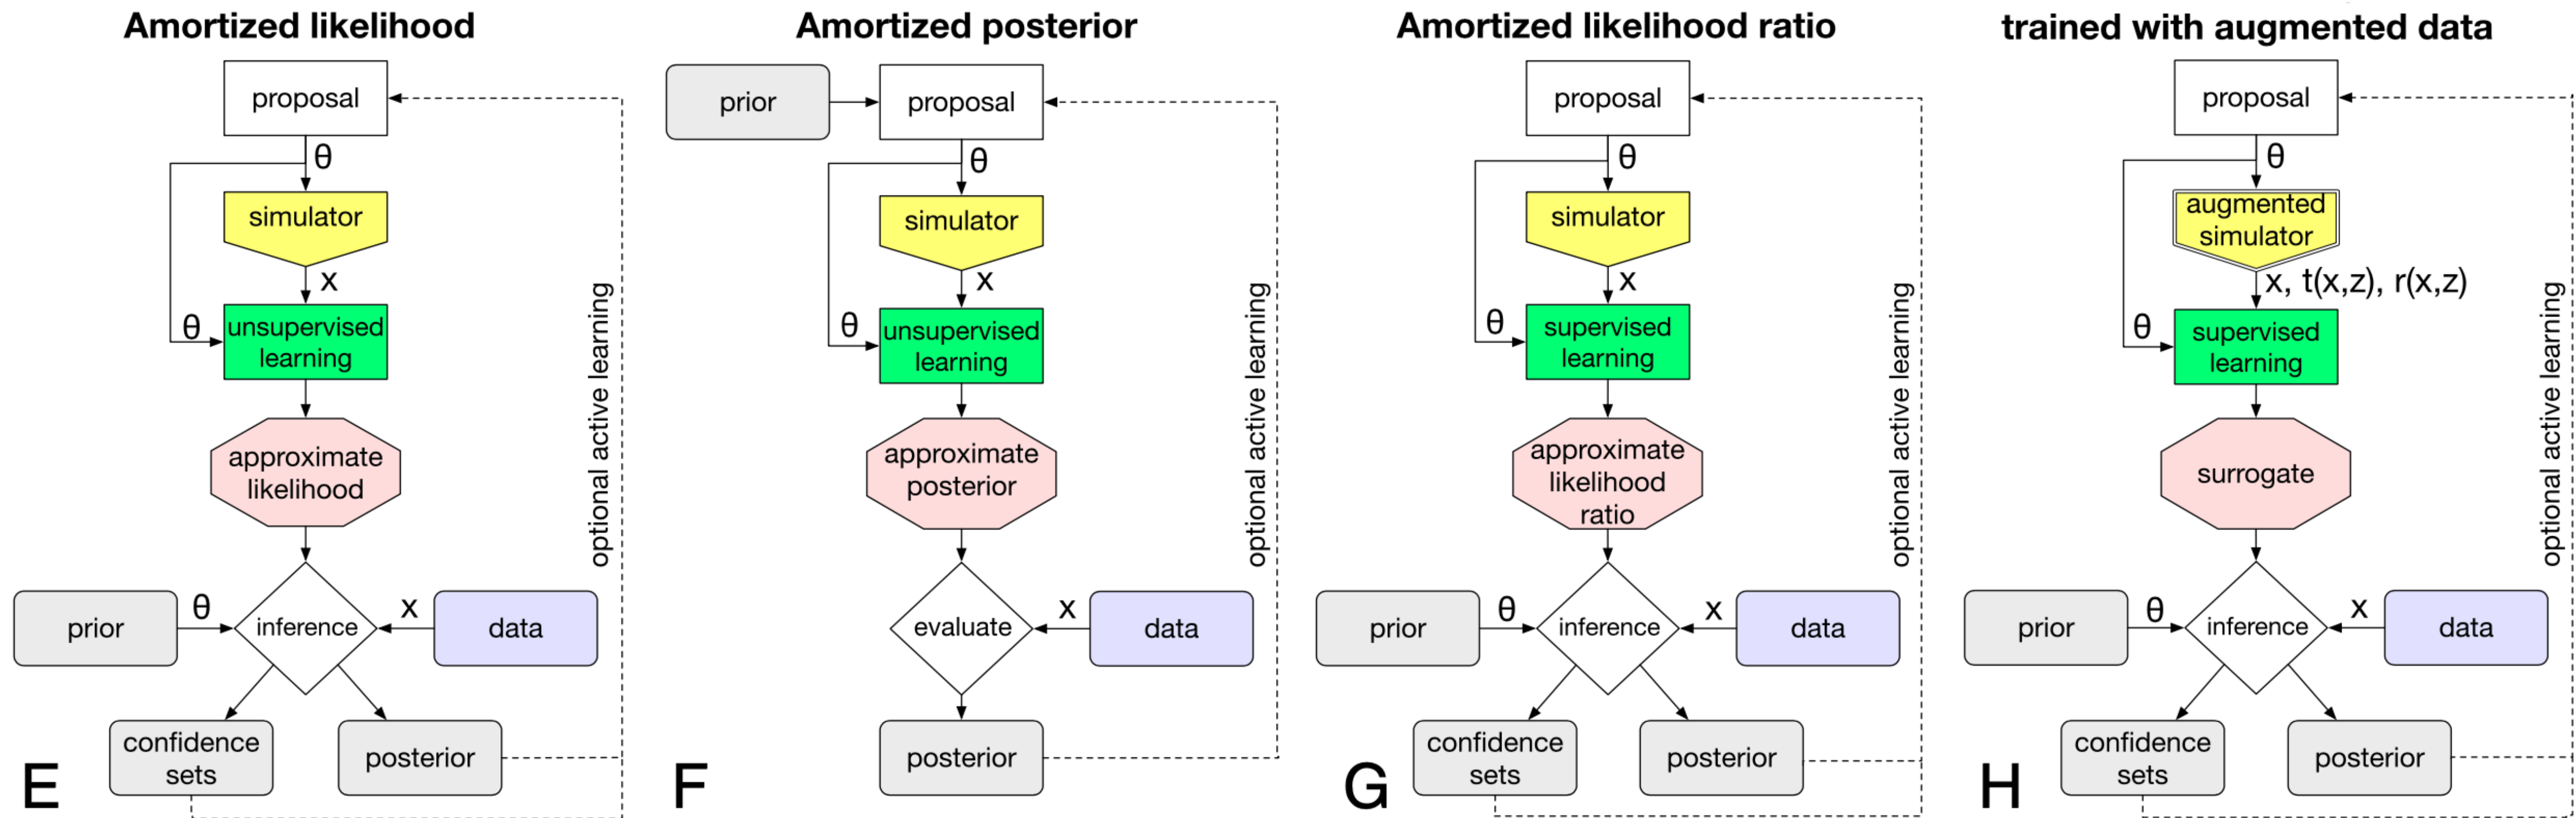


**Query 1:**  $P(B|C) = P(C|B)P(B)/P(C)$

**Query 2:**  $P(A|C) = \sum_B P(A|B)P(B|C)$

(Gershman et al 2014)

# Simulation-Based Inference: Amortisation Techniques



(Cranmer et al 2019)

# Simulation-Based Inference: Neural Network Approximations

Prior Variables  
 $\theta$

$$\theta \sim P(\theta)$$

Status  
Machine  
Parameterised  
by Latent  
Variables  $Z$

$$Z \sim P(Z | \theta)$$

Simulated  
Observations  
 $X$

$$X \sim P(X | Z, \theta)$$

$$P(\theta | X) = \frac{\overset{\text{Likelihood}}{P(X | \theta)} \overset{\text{Prior}}{P(\theta)}}{\underset{\text{Evidence}}{P(X)}}$$

$\theta$ : parameters    $X$ : observations

Posterior

- $P(\theta | X)$  approximated through “Neural Posterior” estimators
- $P(X | \theta)$  approximated through “Neural Likelihood” estimators
- $\frac{P(X | \theta)}{P(X)}$  approximated through the “Neural ratio” estimators

# Simulation-Based Inference: Neural Network Approximations Through Stochastic Flows

$\theta$ : parameters     $X$ : observations


$$P(\theta | X) = \frac{\overset{\text{Likelihood}}{P(X | \theta)} \overset{\text{Prior}}{P(\theta)}}{\underset{\text{Evidence}}{P(X)}}$$

Posterior

$$f(X, \theta) = N_{\mu, \Sigma}(\phi(X, \theta)) | J_{\phi}(X, \theta) |$$

$f$  the Neural estimator and  $\phi$  the stochastic flow

- $P(\theta | X)$  approximated through “Neural Posterior” estimators
- $P(X | \theta)$  approximated through “Neural Likelihood” estimators
- $\frac{P(X | \theta)}{P(X)}$  approximated through the “Neural ratio” estimators





# Simulation-Based Inference: Automatic Posterior Transformation (Greenberg et al 2019)

$$P(\theta | X) = \frac{\overset{\text{Likelihood}}{P(X | \theta)} \overset{\text{Prior}}{P(\theta)}}{\underset{\text{Evidence}}{P(X)}}$$

$\theta$ : parameters    $X$ : observations

- $P(\theta | X)$  approximated through “Neural posterior” by a flow  $Q_{F(x_0, \phi)}(\theta)$

- Loss function:

$$\tilde{q}_{x, \phi}(\theta) = q_{F(x, \phi)}(\theta) \frac{\tilde{p}(\theta)}{p(\theta)} \frac{1}{Z(x, \phi)}, \quad (2)$$

- Where a proposal posterior is

$$\tilde{p}(\theta | x) = p(\theta | x) \frac{\tilde{p}(\theta) p(x)}{p(\theta) \tilde{p}(x)}$$

---

## Algorithm 1 APT with per-round proposal updates

---

**Input:** simulator with (implicit) density  $p(x|\theta)$ , data  $x_o$ , prior  $p(\theta)$ , density family  $q_\psi$ , neural network  $F(x, \phi)$ , simulations per round  $N$ , number of rounds  $R$ .

```

 $\tilde{p}_1(\theta) := p(\theta)$ 
for  $r = 1$  to  $R$  do
  for  $j = 1$  to  $N$  do
    Sample  $\theta_{r,j} \sim \tilde{p}_r(\theta)$ 
    Simulate  $x_{r,j} \sim p(x|\theta_{r,j})$ 
  end for
   $\phi \leftarrow \underset{\phi}{\operatorname{argmin}} \sum_{i=1}^r \sum_{j=1}^N -\log \tilde{q}_{x_{i,j}, \phi}(\theta_{i,j})$     using (2)
   $\tilde{p}_{r+1}(\theta) := q_{F(x_o, \phi)}(\theta)$ 
end for
return  $q_{F(x_o, \phi)}(\theta)$ 

```

---

# Simulation-Based Inference: Sequential Neural Likelihood (Papamakarios et al 2019)

$$P(\theta | X) = \frac{\overset{\text{Likelihood}}{P(X | \theta)} \overset{\text{Prior}}{P(\theta)}}{\underset{\text{Evidence}}{P(X)}}$$

$\theta$ : parameters    $X$ : observations

Posterior

- $P(X | \theta)$  approximated through “Neural likelihood” by a flow  $Q_\phi(X | \theta)$

---

## Algorithm 1: Sequential Neural Likelihood (SNL)

---

**Input** : observed data  $\mathbf{x}_o$ , estimator  $q_\phi(\mathbf{x} | \theta)$ ,  
number of rounds  $R$ , simulations per  
round  $N$

**Output**: approximate posterior  $\hat{p}(\theta | \mathbf{x}_o)$

set  $\hat{p}_0(\theta | \mathbf{x}_o) = p(\theta)$  and  $\mathcal{D} = \{\}$

**for**  $r = 1 : R$  **do**

**for**  $n = 1 : N$  **do**

        sample  $\theta_n \sim \hat{p}_{r-1}(\theta | \mathbf{x}_o)$  with MCMC

        simulate  $\mathbf{x}_n \sim p(\mathbf{x} | \theta_n)$

        add  $(\theta_n, \mathbf{x}_n)$  into  $\mathcal{D}$

    (re-)train  $q_\phi(\mathbf{x} | \theta)$  on  $\mathcal{D}$  and set

$\hat{p}_r(\theta | \mathbf{x}_o) \propto q_\phi(\mathbf{x}_o | \theta) p(\theta)$

**return**  $\hat{p}_R(\theta | \mathbf{x}_o)$

---

# Simulation-Based Inference: Neural Ratio (Hermans et al 2020)

$$P(\theta | X) = \frac{\overset{\text{Likelihood}}{P(X | \theta)} \overset{\text{Prior}}{P(\theta)}}{\underset{\text{Evidence}}{P(X)}}$$

$\theta$ : parameters    $X$ : observations

Posterior

- $P(X | \theta) / P(X)$  approximated through “Neural ratio” by a flow  $d_\phi(X | \theta)$

---

## Algorithm 1 Optimization of $d_\phi(\mathbf{x}, \theta)$ .

---

*Inputs:*

Criterion  $\ell$  (e.g., BCE)

Implicit generative model  $p(\mathbf{x} | \theta)$

Prior  $p(\theta)$

*Outputs:*

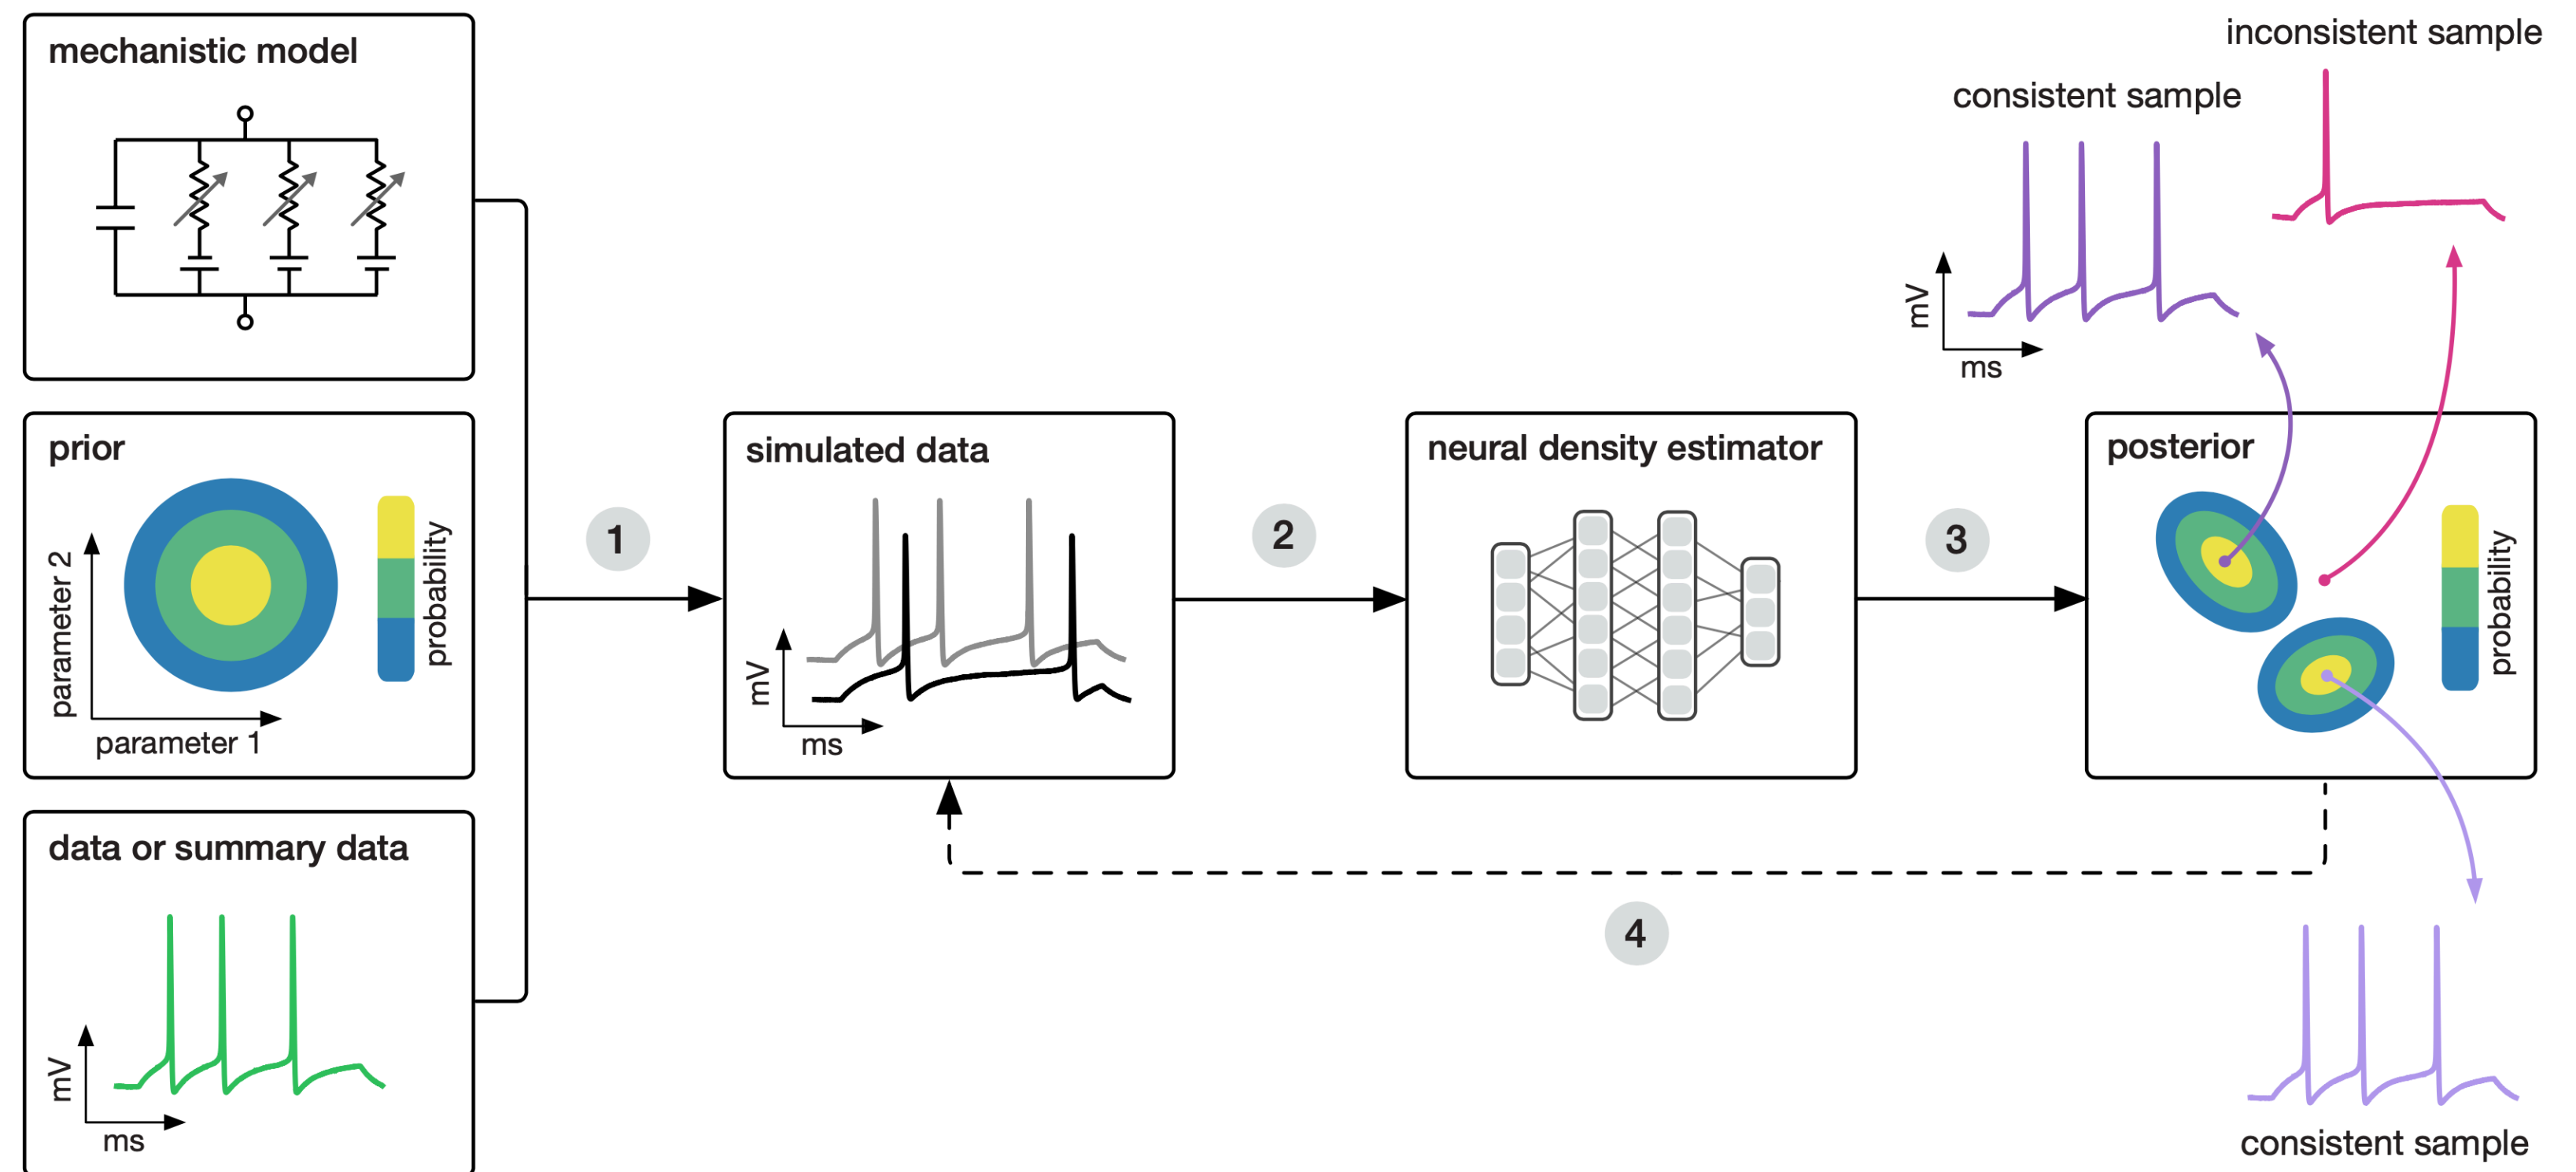
Parameterized classifier  $d_\phi(\mathbf{x}, \theta)$

*Hyperparameters:* Batch-size  $M$

- 1: **while not converged do**
  - 2:     **Sample**  $\theta \leftarrow \{\theta_m \sim p(\theta)\}_{m=1}^M$
  - 3:     **Sample**  $\theta' \leftarrow \{\theta'_m \sim p(\theta)\}_{m=1}^M$
  - 4:     **Simulate**  $\mathbf{x} \leftarrow \{\mathbf{x}_m \sim p(\mathbf{x} | \theta_m)\}_{m=1}^M$
  - 5:      $\mathcal{L} \leftarrow \ell(d_\phi(\mathbf{x}, \theta), 1) + \ell(d_\phi(\mathbf{x}, \theta'), 0)$
  - 6:      $\phi \leftarrow \text{OPTIMIZER}(\phi, \nabla_\phi \mathcal{L})$
  - 7: **end while**
  - 8: **return**  $d_\phi$
-


# Training deep neural density estimators to identify mechanistic models of neural dynamics

Pedro J Gonçalves<sup>1,2†\*</sup>, Jan-Matthis Lueckmann<sup>1,2†\*</sup>, Michael Deistler<sup>1,3†\*</sup>, Marcel Nonnenmacher<sup>1,2,4</sup>, Kaan Öcal<sup>2,5</sup>, Giacomo Bassetto<sup>1,2</sup>, Chaitanya Chintaluri<sup>6,7</sup>, William F Podlaski<sup>6</sup>, Sara A Haddad<sup>8</sup>, Tim P Vogels<sup>6,7</sup>, David S Greenberg<sup>1,4</sup>, Jakob H Macke<sup>1,2,3,9\*</sup>



# Training deep neural density estimators to identify mechanistic models of neural dynamics

Pedro J Gonçalves<sup>1,2†\*</sup>, Jan-Matthis Lueckmann<sup>1,2†\*</sup>, Michael Deistler<sup>1,3†\*</sup>, Marcel Nonnenmacher<sup>1,2,4</sup>, Kaan Öcal<sup>2,5</sup>, Giacomo Bassetto<sup>1,2</sup>, Chaitanya Chintaluri<sup>6,7</sup>, William F Podlaski<sup>6</sup>, Sara A Haddad<sup>8</sup>, Tim P Vogels<sup>6,7</sup>, David S Greenberg<sup>1,4</sup>, Jakob H Macke<sup>1,2,3,9\*</sup>



---



# SIMULATION-BASED BAYESIAN INFERENCE FOR MULTI-FINGERED ROBOTIC GRASPING

---

**Norman Marlier**  
University of Liège  
norman.marlier@uliege.be

**Olivier Bruls**  
University of Liège  
o.bruls@uliege.be

**Gilles Louppe**  
University of Liège  
g.louppe@uliege.be

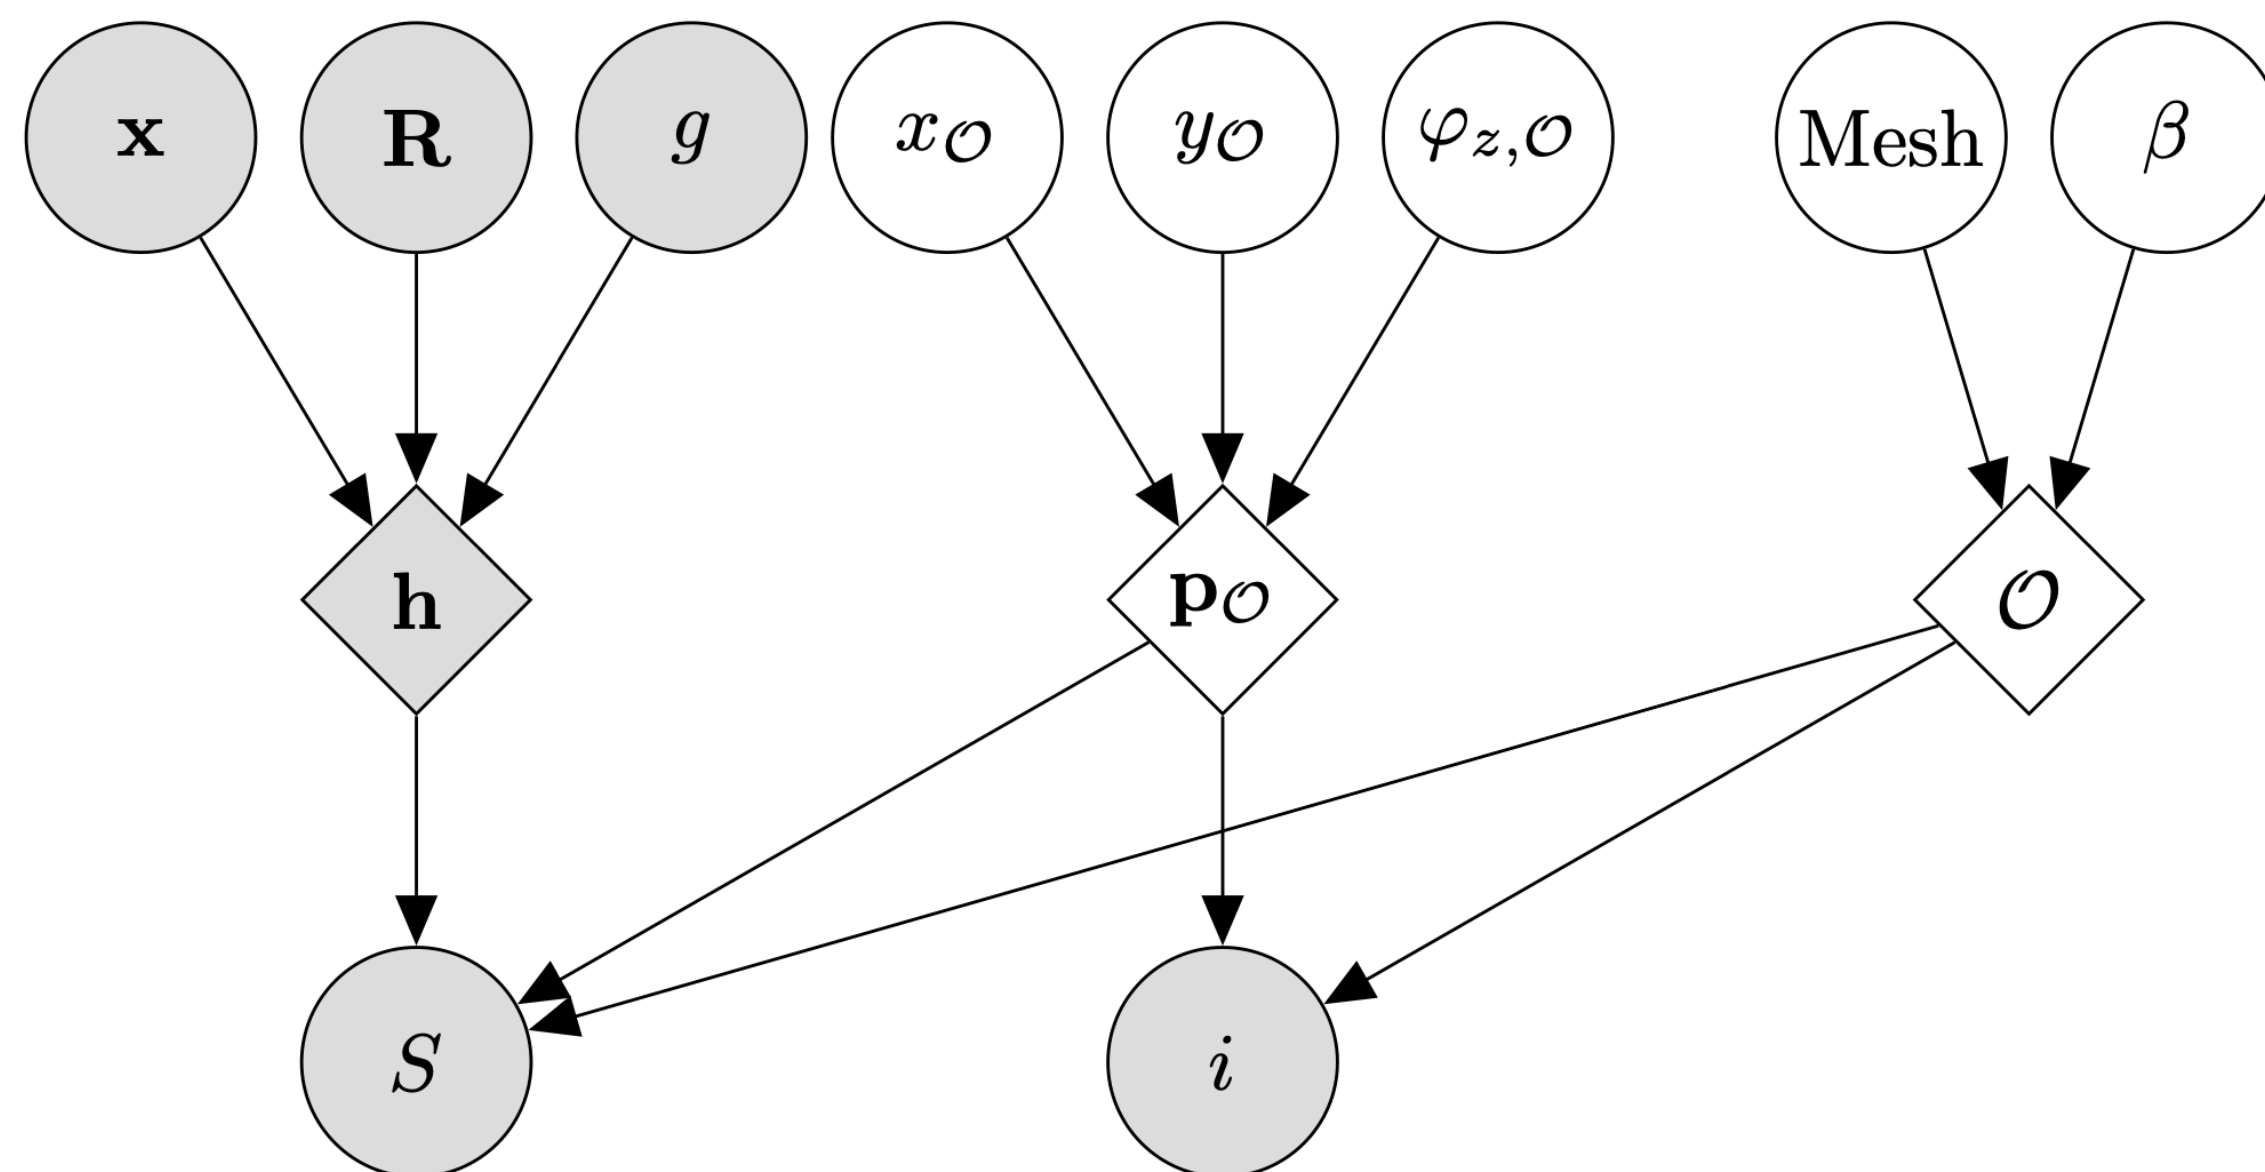


# SIMULATION-BASED BAYESIAN INFERENCE FOR MULTI-FINGERED ROBOTIC GRASPING


**Norman Marlier**  
University of Liège  
norman.marlier@uliege.be

**Olivier Bruls**  
University of Liège  
o.bruls@uliege.be

**Gilles Louppe**  
University of Liège  
g.louppe@uliege.be



(a)



Variable	Prior
$x$	$\text{uniform}(-0.15, 0.15)$
$y$	$\text{uniform}(-0.15, 0.15)$
$z$	$\text{uniform}(0.12, 0.34)$
$R$	$\text{mixture of power spherical}(\mu_i, \kappa)$
$g$	$\text{categorical}(\{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\})$
$x_O$	$\text{uniform}(-0.05, 0.05)$
$y_O$	$\text{uniform}(-0.05, 0.05)$
$\varphi_{z,O}$	$\text{uniform}(-\pi, \pi)$
$\text{Mesh}$	$\text{uniform in the set of objects}$
$\beta$	$\text{uniform}(0.9, 1.1)$

(b)

Figure 2: (a) Probabilistic graphical model of the environment. Gray nodes correspond to observed variables and white nodes to unobserved variables. (b) Prior distributions.

# SIMULATION-BASED BAYESIAN INFERENCE FOR MULTI-FINGERED ROBOTIC GRASPING

**Norman Marlier**  
University of Liège  
norman.marlier@uliege.be

**Olivier Brûls**  
University of Liège  
o.bruls@uliege.be

**Gilles Louppe**  
University of Liège  
g.louppe@uliege.be

