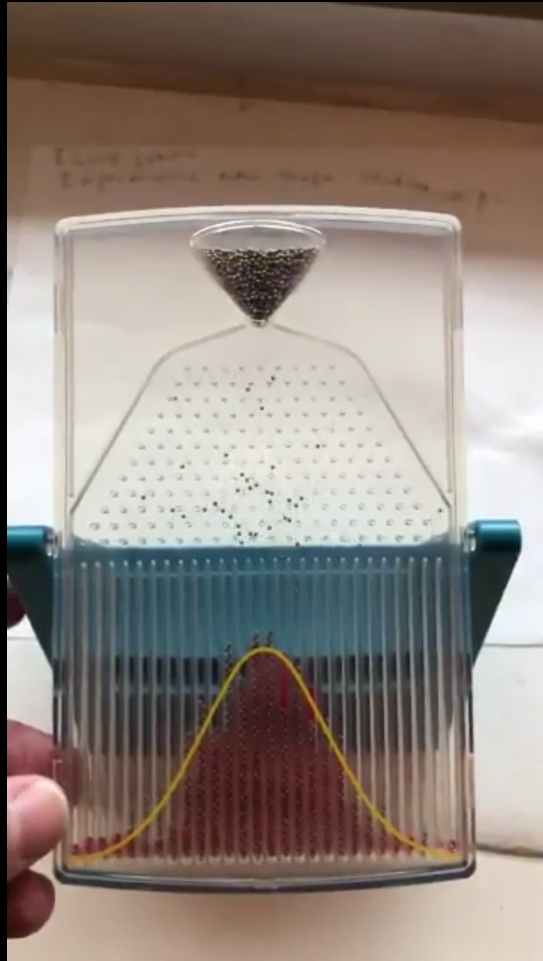



# Likelihood-free inference in Physical Sciences

Machine Learning in High Energy Physics Summer School 2019  
July 4, DESY

Gilles Louppe  
[g.louppe@uliege.be](mailto:g.louppe@uliege.be)





The probability of ending in bin  $x$  corresponds to the total probability of all the paths  $z$  from start to  $x$ .


$$p(x|\theta) = \int p(x, z|\theta) dz = \binom{n}{x} \theta^x (1 - \theta)^{n-x}$$

What if we shift or remove some of the pins?


The Galton board is a **metaphore** of simulation-based science:

Galton board device	→	Computer simulation
Parameters $\theta$	→	Model parameters $\theta$
Buckets $x$	→	Observables $x$
Random paths $z$	→	Latent variables $z$ (stochastic execution traces through simulator)

Inference in this context requires **likelihood-free algorithms**.




- Prediction:
- Well-understood mechanistic model
  - Simulator can generate samples






- Prediction:
- Well-understood mechanistic model
  - Simulator can generate samples

- Inference:
- Likelihood function  $p(x|\theta)$  is intractable
  - Inference based on estimator  $\hat{p}(x|\theta)$


# Applications




Particle physics




Cosmology



Epidemiology



Climatology




Computational topography




Astronomy

# Particle physics



$$\begin{aligned}
 C_{SM} = & -\frac{1}{2}\partial_\mu g_\nu^2 \partial_\mu g_\nu^2 - g_\nu^2 f^{abc} \partial_\mu g_\nu^c \partial_\mu g_\nu^c - \frac{1}{2}g_\nu^2 f^{abc} f^{def} g_\nu^c g_\nu^d g_\nu^e g_\nu^f - \partial_\mu W_\nu^+ \partial_\mu W_\nu^- - \\
 & M^2 W_\nu^+ W_\nu^- - \frac{1}{2}\partial_\mu Z_\nu^0 \partial_\mu Z_\nu^0 - \frac{1}{2}M^2 Z_\nu^0 Z_\nu^0 - \frac{1}{2}\partial_\mu A_\nu \partial_\mu A_\nu - i g_{\nu\alpha} (\partial_\mu Z_\nu^0 (W_\nu^+ W_\nu^- - \\
 & W_\nu^- W_\nu^+) - 2Z_\nu^0 \partial_\mu W_\nu^+ W_\nu^- - W_\nu^- \partial_\mu W_\nu^+) + 2Z_\nu^0 (W_\nu^+ \partial_\mu W_\nu^- - W_\nu^- \partial_\mu W_\nu^+) - \\
 & i g_{\nu\alpha} (\partial_\mu A_\nu (W_\nu^+ W_\nu^- - W_\nu^- W_\nu^+) - A_\nu (W_\nu^+ \partial_\mu W_\nu^- - W_\nu^- \partial_\mu W_\nu^+) + A_\nu (W_\nu^+ \partial_\mu W_\nu^- - \\
 & W_\nu^- \partial_\mu W_\nu^+) - \frac{1}{2}g^2 W_\nu^+ W_\nu^- W_\nu^+ W_\nu^- + \frac{1}{2}g^2 W_\nu^+ W_\nu^- W_\nu^- W_\nu^+) + g^2 \tilde{c}_2^2 (Z_\nu^0 W_\nu^+ Z_\nu^0 W_\nu^- - \\
 & Z_\nu^0 Z_\nu^0 W_\nu^+ W_\nu^-) + g^2 \tilde{c}_3^2 (A_\nu W_\nu^+ A_\nu W_\nu^- - A_\nu A_\nu W_\nu^+ W_\nu^-) + g^2 s_{\nu\alpha} (A_\nu Z_\nu^0 (W_\nu^+ W_\nu^- - \\
 & W_\nu^- W_\nu^+) - 2A_\nu Z_\nu^0 W_\nu^+ W_\nu^-) - \frac{1}{2}\partial_\mu H \partial_\mu H - 2M^2 \alpha H^2 - \partial_\mu \phi^+ \partial_\mu \phi^- - \frac{1}{2}\partial_\mu \phi^0 \partial_\mu \phi^0 - \\
 & \partial_\mu \left( \frac{2M^2}{\alpha} H + \frac{1}{2}(H^2 + \phi^0 \phi^0 + 2\phi^+ \phi^-) \right) + \frac{2M^2}{\alpha} \alpha \alpha - \\
 & \frac{1}{2}g^2 \alpha \alpha (H^4 + (\phi^0)^4 + 4(\phi^+ \phi^-)^2 + 4(\phi^0)^2 \phi^+ \phi^- + 4H^2 \phi^+ \phi^- + 2(\phi^0)^2 H^2) - \\
 & g M W_\nu^+ W_\nu^- H - \frac{1}{2}g \frac{M^2}{\alpha} Z_\nu^0 Z_\nu^0 H - \\
 & \frac{1}{2}i g (W_\nu^+ (\phi^0 \partial_\mu \phi^- - \phi^- \partial_\mu \phi^0) - W_\nu^- (\phi^0 \partial_\mu \phi^+ - \phi^+ \partial_\mu \phi^0)) + \\
 & \frac{1}{2}g (W_\nu^+ (H \partial_\mu \phi^- - \phi^- \partial_\mu H) + W_\nu^- (H \partial_\mu \phi^+ - \phi^+ \partial_\mu H)) + \frac{1}{2}g \frac{1}{\alpha} (Z_\nu^0 (H \partial_\mu \phi^0 - \phi^0 \partial_\mu H) + \\
 & M (\frac{1}{\alpha} Z_\nu^0 \partial_\mu \phi^0 + W_\nu^+ \partial_\mu \phi^- + W_\nu^- \partial_\mu \phi^+) - i g \frac{M^2}{\alpha} Z_\nu^0 (W_\nu^+ \phi^- - W_\nu^- \phi^+) + i g_{\nu\alpha} M A_\nu (W_\nu^+ \phi^- - \\
 & W_\nu^- \phi^+) - i g \frac{1}{\alpha} \frac{2M^2}{\alpha} Z_\nu^0 (\phi^+ \partial_\mu \phi^- - \phi^- \partial_\mu \phi^+) + i g_{\nu\alpha} A_\nu (\phi^+ \partial_\mu \phi^- - \phi^- \partial_\mu \phi^+) - \\
 & \frac{1}{2}g^2 W_\nu^+ W_\nu^- (H^2 + (\phi^0)^2 + 2\phi^+ \phi^-) - \frac{1}{2}g^2 \frac{1}{\alpha} Z_\nu^0 Z_\nu^0 (H^2 + (\phi^0)^2 + 2(\phi^\pm)^2 - 1)^2 \phi^+ \phi^- - \\
 & \frac{1}{2}g^2 \frac{1}{\alpha} Z_\nu^0 \partial_\mu (W_\nu^+ \phi^- + W_\nu^- \phi^+) - \frac{1}{2}i g^2 \frac{1}{\alpha} Z_\nu^0 H (W_\nu^+ \phi^- - W_\nu^- \phi^+) + \frac{1}{2}g^2 s_{\nu\alpha} A_\nu (W_\nu^+ \phi^- + \\
 & W_\nu^- \phi^+) + \frac{1}{2}i g^2 s_{\nu\alpha} H (W_\nu^+ \phi^- - W_\nu^- \phi^+) - g^2 \frac{1}{\alpha} (2\tilde{c}_2^2 - 1) Z_\nu^0 A_\nu \phi^+ \phi^- - \\
 & g^2 \tilde{c}_3^2 A_\nu A_\nu \phi^+ \phi^- + \frac{1}{2}i g_\nu \lambda_\nu^2 (g_\nu^+ \phi^+ g_\nu^- \phi^- - \phi^+ (\gamma \partial + m_\nu^2) \phi^- - \phi^- (\gamma \partial + m_\nu^2) \phi^+) - \\
 & \frac{1}{2}(g_\nu^+ \gamma \partial + m_\nu^2) g_\nu^- + i g_{\nu\alpha} A_\nu (-e^{\gamma \lambda} \nu e^\lambda) + \frac{1}{2}(g_\nu^+ \gamma \partial g_\nu^-) - \frac{1}{2}(g_\nu^- \gamma \partial g_\nu^+) + \\
 & \frac{i g_\nu}{2} Z_\nu^0 ((\nu^+ \gamma \partial (1 + \gamma^2) \nu^+) + (\nu^+ \nu^+ (4\tilde{c}_2^2 - 1 - \gamma^2) \nu^+) + (g_\nu^+ \gamma (4\tilde{c}_2^2 - 1 - \gamma^2) \nu^+) + \\
 & (g_\nu^+)^2 (1 - \frac{1}{2}\tilde{c}_2^2 + \gamma^2) \nu^+)) + \frac{i g_\nu}{2} W_\nu^+ ((\nu^+ \nu^+ (1 + \gamma^2) \nu^+) + (g_\nu^+ \gamma (4\tilde{c}_2^2 - 1 - \gamma^2) \nu^+) + \\
 & \frac{i g_\nu}{2} W_\nu^- ((\nu^- \nu^- \nu^- (1 + \gamma^2) \nu^-) + (g_\nu^- \gamma (4\tilde{c}_2^2 - 1 - \gamma^2) \nu^-) + \\
 & \frac{i g_\nu}{2} W_\nu^0 (-m_\nu^2 (\nu^0 \nu^0 \nu^0 (1 - \gamma^2) \nu^0) + m_\nu^2 (\nu^0 \nu^0 \nu^0 (1 + \gamma^2) \nu^0) + \\
 & \frac{i g_\nu}{2} W_\nu^0 (-m_\nu^2 (\nu^0 \nu^0 \nu^0 (1 + \gamma^2) \nu^0) - m_\nu^2 (\nu^0 \nu^0 \nu^0 (1 - \gamma^2) \nu^0) - \frac{1}{2} \frac{M^2}{\alpha} H (\nu^0 \nu^0) - \\
 & \frac{1}{2} \frac{M^2}{\alpha} H (e^{\lambda} \nu^0) + \frac{1}{2} \frac{M^2}{\alpha} \partial^\mu (\nu^0 \lambda^0 \nu^0) - \frac{1}{2} \frac{M^2}{\alpha} \partial^\mu (\nu^0 \lambda^0 \nu^0) - \frac{1}{2} \nu_\alpha M_\mu^\alpha (1 - \gamma_\alpha) \nu_\alpha - \\
 & \frac{1}{4} \nu_\alpha M_\mu^\alpha (1 - \gamma_\alpha) \nu_\alpha + \frac{1}{2} \frac{M^2}{\alpha} \partial^\mu (-m_\nu^2 (g_\nu^+ C_{\nu\alpha} (1 - \gamma^2) \nu^+) + m_\nu^2 (g_\nu^- C_{\nu\alpha} (1 + \gamma^2) \nu^-) + \\
 & \frac{i g_\nu}{2} W_\nu^0 (-m_\nu^2 (g_\nu^+ C_{\nu\alpha}^+ (1 + \gamma^2) \nu^+) - m_\nu^2 (g_\nu^- C_{\nu\alpha}^- (1 - \gamma^2) \nu^-) - \frac{1}{2} \frac{M^2}{\alpha} H (g_\nu^+ \nu^+) - \\
 & \frac{1}{2} \frac{M^2}{\alpha} H (g_\nu^- \nu^-) + \frac{1}{2} \frac{M^2}{\alpha} \partial^\mu (g_\nu^+ \gamma^2 \nu^+) - \frac{1}{2} \frac{M^2}{\alpha} \partial^\mu (g_\nu^- \gamma^2 \nu^-) + G^+ \partial^\mu G^+ + g_\nu f^{abc} \partial_\mu G^+ G^+ G_\mu^+ + \\
 & X^+ (\partial^\mu - M^2) X^+ + X^+ (\partial^\mu - M^2) X^+ + X^0 (\partial^\mu - \frac{M^2}{\alpha}) X^0 + Y^+ \partial^\mu Y^+ + i g_{\nu\alpha} W_\nu^+ (\partial_\mu X^+ X^+ - \\
 & \partial_\mu X^+ X^+) + i g_{\nu\alpha} W_\nu^- (\partial_\mu X^- X^- - \partial_\mu X^- X^-) + i g_{\nu\alpha} W_\nu^0 (\partial_\mu X^0 X^0 - \\
 & \partial_\mu X^0 X^0) + i g_{\nu\alpha} W_\nu^0 (\partial_\mu X^+ X^- - \partial_\mu X^- X^+) + i g_{\nu\alpha} Z_\nu^0 (\partial_\mu X^+ X^+ - \\
 & \partial_\mu X^- X^-) + i g_{\nu\alpha} A_\nu (\partial_\mu X^+ X^+ - \partial_\mu X^- X^-) + \\
 & \partial_\mu X^+ X^- - \frac{1}{2} i g M (X^+ X^+ H + X^- X^- H + \frac{1}{2} X^0 X^0 H) + \frac{1}{2} \frac{M^2}{\alpha} i g M (X^+ X^0 \phi^+ - X^- X^0 \phi^-) + \\
 & \frac{1}{2} i g M (X^0 X^+ \phi^- - X^0 X^+ \phi^+) + i g M s_{\nu\alpha} (X^0 X^+ \phi^- - X^0 X^+ \phi^+) + \\
 & \frac{1}{2} i g M (X^+ X^+ \phi^- - X^- X^- \phi^+) .
 \end{aligned}$$






Latent variables

Parameters of interest

Parton-level momenta

Theory parameters

$$z_p \longleftarrow \theta$$



Latent variables


Parameters  
of interest

Shower  
splittings

Parton-level  
momenta

Theory  
parameters

$z_s$  ←  $z_p$  ←  $\theta$



# Latent variables

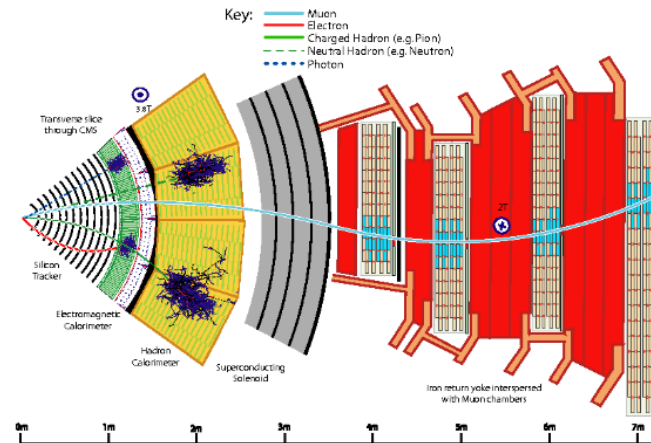

Parameters of interest



Detector interactions

Shower splittings

Parton-level momenta

Theory parameters



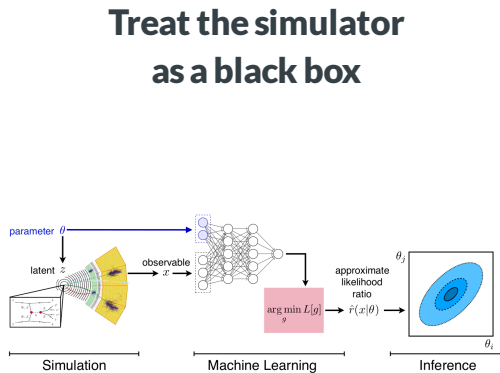


[Image source: M. Cacciari, G. Salam, G. Soyez 0802.1189]

$$p(x|\theta) = \underbrace{\iiint}_{\text{intractable}} p(z_p|\theta)p(z_s|z_p)p(z_d|z_s)p(x|z_d)dz_pdz_sdz_d$$


# Likelihood-free inference algorithms

## Learn a proxy for inference




Histograms of observables  
Neural density (ratio) estimation

Adversarial variational optimization




## Make use of the inner structure

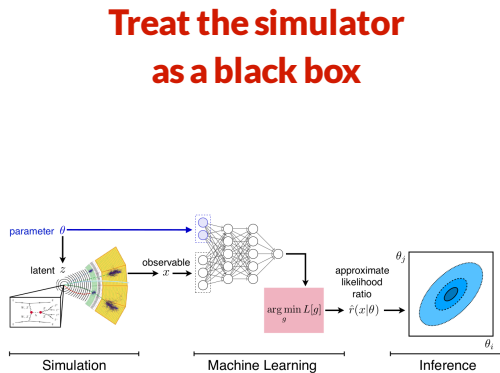


Mining gold from implicit models


Probabilistic programming



## Learn a proxy for inference




Histograms of observables  
Neural density (ratio) estimation




Adversarial variational optimization

## Make use of the inner structure



Mining gold from implicit models

## Learn to control the simulator



Probabilistic programming



# Likelihood ratio

The likelihood ratio

$$r(x|\theta_0, \theta_1) = \frac{p(x|\theta_0)}{p(x|\theta_1)}$$

is the quantity that is **central** to many **statistical inference** procedures.

## Examples

- Frequentist hypothesis testing
- Supervised learning
- Bayesian posterior sampling with MCMC
- Bayesian posterior inference through Variational Inference
- Generative adversarial networks
- Empirical Bayes with Adversarial Variational Optimization
- Optimal compression

*When solving a problem of interest, do not solve a more general problem as an intermediate step. – Vladimir Vapnik*



$$\frac{p_{\mathbf{x}}(\mathbf{x}|\theta_0)}{p_{\mathbf{x}}(\mathbf{x}|\theta_1)} = r(\mathbf{x}; \theta_0, \theta_1)$$

The equation is annotated with a blue curved arrow pointing from the left side to the right side, and a red curved arrow pointing from the right side to the left side, with a red diagonal slash through it.

Direct likelihood ratio estimation is simpler than density estimation.


(This is fortunate, we are in the likelihood-free scenario!)

# The frequentist physicist's way

The Neyman-Pearson lemma states that the likelihood ratio

$$r(x|\theta_0, \theta_1) = \frac{p(x|\theta_0)}{p(x|\theta_1)}$$

is the **most powerful test statistic** to discriminate between a null hypothesis  $\theta_0$  and an alternative  $\theta_1$ .



## IX. *On the Problem of the most Efficient Tests of Statistical Hypotheses.*


By J. NEYMAN, *Nencki Institute, Soc. Sci. Lit. Varsoviensis, and Lecturer at the Central College of Agriculture, Warsaw,* and E. S. PEARSON, *Department of Applied Statistics, University College, London.*

(Communicated by K. PEARSON, F.R.S.)

(Received August 31, 1932.—Read November 10, 1932.)

### CONTENTS.

	PAGE.
I. Introductory . . . . .	289
II. Outline of General Theory . . . . .	293
III. Simple Hypotheses . . . . .	




Define a projection function  $s : \mathcal{X} \rightarrow \mathbb{R}$  mapping observables  $x$  to a summary statistics  $x' = s(x)$ .

Then, **approximate** the likelihood  $p(x|\theta)$  as

$$p(x|\theta) \approx \hat{p}(x|\theta) = p(x'|\theta).$$

From this it comes

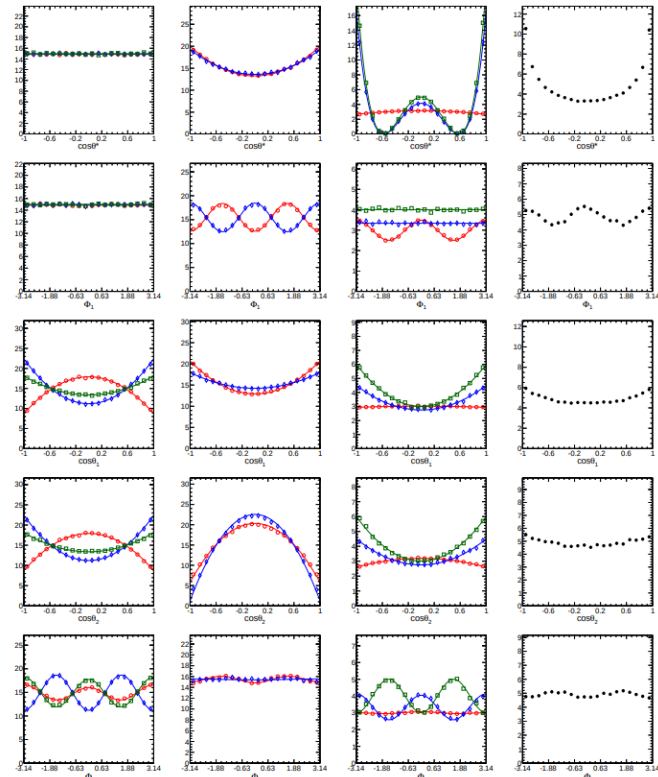
$$\frac{p(x|\theta_0)}{p(x|\theta_1)} \approx \frac{\hat{p}(x|\theta_0)}{\hat{p}(x|\theta_1)} = \hat{r}(x|\theta_0, \theta_1).$$



This methodology has worked great for physicists for the last 20-30 years, but ...

- Choosing the projection  $s$  is difficult and problem-dependent.
- Often there is no single good variable: compressing to any  $x'$  loses information.
- Ideally: analyse high-dimensional  $x'$ , including all correlations.


Unfortunately, filling high-dimensional histograms is **not tractable**.



# Bayesian inference

Bayesian inference usually consists in computing the posterior

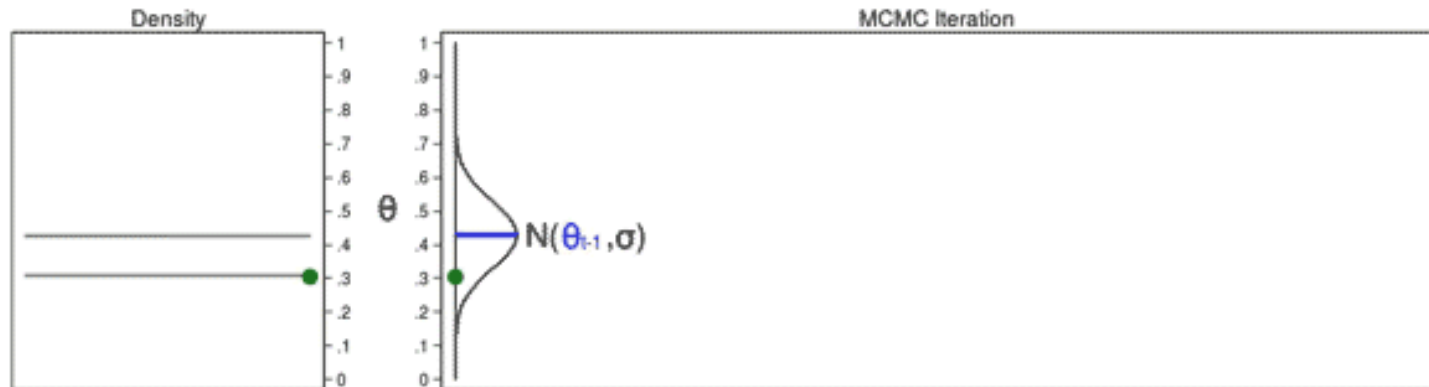
$$p(\theta|x) = \frac{p(x|\theta)p(\theta)}{p(x)}.$$



Doubly **intractable** in the likelihood-free scenario:

- Cannot evaluate the evidence  $p(x) = \int p(x|\theta)p(\theta)d\theta$ .
- Cannot evaluate the likelihood  $p(x|\theta) = \int p(x, z|\theta)dz$ .

# Posterior sampling



$$\text{Step 1: } r(\theta_{new}, \theta_{t-1}) = \frac{\text{Posterior}(\theta_{new})}{\text{Posterior}(\theta_{t-1})} = \frac{\text{Beta}(1,1,0.306) \times \text{Binomial}(10,4,0.306)}{\text{Beta}(1,1,0.429) \times \text{Binomial}(10,4,0.429)} = 0.834$$

$$\text{Step 2: } \text{Acceptance probability } \alpha(\theta_{new}, \theta_{t-1}) = \min\{r(\theta_{new}, \theta_{t-1}), 1\} = \min\{0.834, 1\} = 0.834$$

$$\text{Step 3: } \text{Draw } u \sim \text{Uniform}(0,1) = 0.617$$

$$\text{Step 4: } \text{If } u < \alpha(\theta_{new}, \theta_{t-1}) \rightarrow \text{If } 0.617 < 0.834 \quad \text{Then } \theta_t = \theta_{new} = 0.306$$

$$\text{Otherwise } \theta_t = \theta_{t-1} = 0.429$$

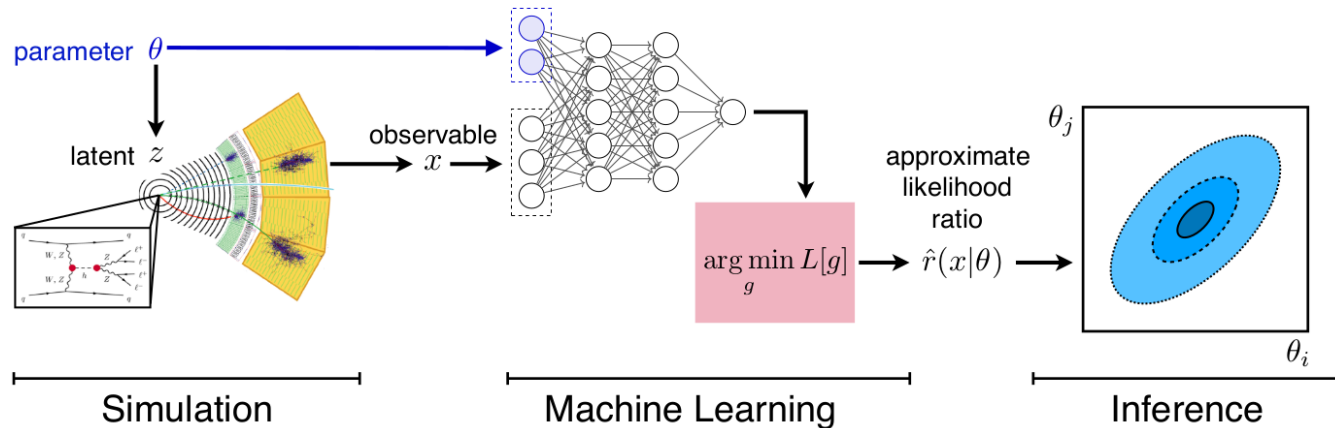
MCMC algorithms can be made likelihood-free by plugging in the likelihood ratio.

# CARL

Supervised learning provides a way to **automatically** construct  $s$ :

- Let us consider a binary classifier  $\hat{s}$  (e.g., a neural network) trained to distinguish  $x \sim p(x|\theta_0)$  from  $x \sim p(x|\theta_1)$ .
- $\hat{s}$  is trained by minimizing the cross-entropy loss

$$L_{XE}[\hat{s}] = -\mathbb{E}_{p(x|\theta)\pi(\theta)} [1(\theta = \theta_0) \log \hat{s}(x) + 1(\theta = \theta_1) \log(1 - \hat{s}(x))]$$





The solution  $\hat{s}$  found after training approximates the optimal classifier

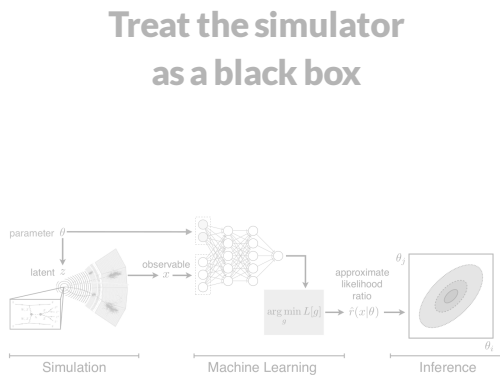
$$\hat{s}(x) \approx s^*(x) = \frac{p(x|\theta_1)}{p(x|\theta_0) + p(x|\theta_1)}.$$

Therefore,


$$r(x|\theta_0, \theta_1) \approx \hat{r}(x|\theta_0, \theta_1) = \frac{1 - \hat{s}(x)}{\hat{s}(x)}$$

That is, supervised classification is equivalent to likelihood ratio estimation.

## Learn a proxy for inference




Histograms of observables  
Neural density (ratio) estimation




Adversarial variational optimization

## Make use of the inner structure




Mining gold from implicit models



Probabilistic programming

# Mining gold from simulators



$p(x|\theta)$  is usually intractable.

What about  $p(x, z|\theta)$ ?

As the trajectory  $z_1, \dots, z_T$  and the observable  $x$  are emitted, it is often possible:

- to calculate the **joint likelihood**  $p(x, z|\theta)$ ;
- to calculate the **joint likelihood ratio**  $r(x, z|\theta_0, \theta_1)$ ;
- to calculate the **joint score**  $t(x, z|\theta_0) = \nabla_{\theta} \log p(x, z|\theta)|_{\theta_0}$ .


We call this process **mining gold** from your simulator!

Observe that the joint likelihood ratios

$$r(x, z|\theta_0, \theta_1) = \frac{p(x, z|\theta_0)}{p(x, z|\theta_1)}$$

are scattered around  $r(x|\theta_0, \theta_1)$ .

Can we use them to approximate  $r(x|\theta_0, \theta_1)$ ?



## Key insights

Consider the squared error of a function  $\hat{g}(x)$  that only depends on  $x$ , but is trying to approximate a function  $g(x, z)$  that also depends on the latent  $z$ :

$$L_{MSE} = \mathbb{E}_{p(x, z | \theta)} [(g(x, z) - \hat{g}(x))^2].$$

Via calculus of variations, we find that the function  $g^*(x)$  that extremizes  $L_{MSE}[g]$  is given by


$$\begin{aligned} g^*(x) &= \frac{1}{p(x|\theta)} \int p(x, z|\theta) g(x, z) dz \\ &= \mathbb{E}_{p(z|x, \theta)} [g(x, z)] \end{aligned}$$

Therefore, by identifying the  $g(x, z)$  with the joint likelihood ratio  $r(x, z|\theta_0, \theta_1)$  and  $\theta$  with  $\theta_1$ , we define

$$L_r = \mathbb{E}_{p(x, z|\theta_1)} [(r(x, z|\theta_0, \theta_1) - \hat{r}(x))^2],$$

which is minimized by

$$\begin{aligned} r^*(x) &= \frac{1}{p(x|\theta_1)} \int p(x, z|\theta_1) \frac{p(x, z|\theta_0)}{p(x, z|\theta_1)} dz \\ &= \frac{p(x|\theta_0)}{p(x|\theta_1)} \\ &= r(x|\theta_0, \theta_1). \end{aligned}$$




$$r^*(x|\theta_0, \theta_1) = \arg \min_{\hat{r}} L_r[\hat{r}]$$



Similarly, we can mine the simulator to extract the joint score

$$t(x, z|\theta_0) = \nabla_{\theta} \log p(x, z|\theta)|_{\theta_0},$$

which indicates how much more or less likely  $x, z$  would be if one changed  $\theta_0$ .



Using the same trick, by identifying  $g(x, z)$  with the joint score  $t(x, z|\theta_0)$  and  $\theta$  with  $\theta_0$ , we define


$$L_t = \mathbb{E}_{p(x, z|\theta_0)} [(t(x, z|\theta_0) - \hat{t}(x))^2],$$

which is minimized by

$$\begin{aligned} t^*(x) &= \frac{1}{p(x|\theta_0)} \int p(x, z|\theta_0) (\nabla_{\theta} \log p(x, z|\theta)|_{\theta_0}) dz \\ &= \frac{1}{p(x|\theta_0)} \int p(x, z|\theta_0) \frac{\nabla_{\theta} p(x, z|\theta)|_{\theta_0}}{p(x, z|\theta_0)} dz \\ &= \frac{\nabla_{\theta} p(x|\theta)|_{\theta_0}}{p(x|\theta_0)} \\ &= t(x|\theta_0). \end{aligned}$$

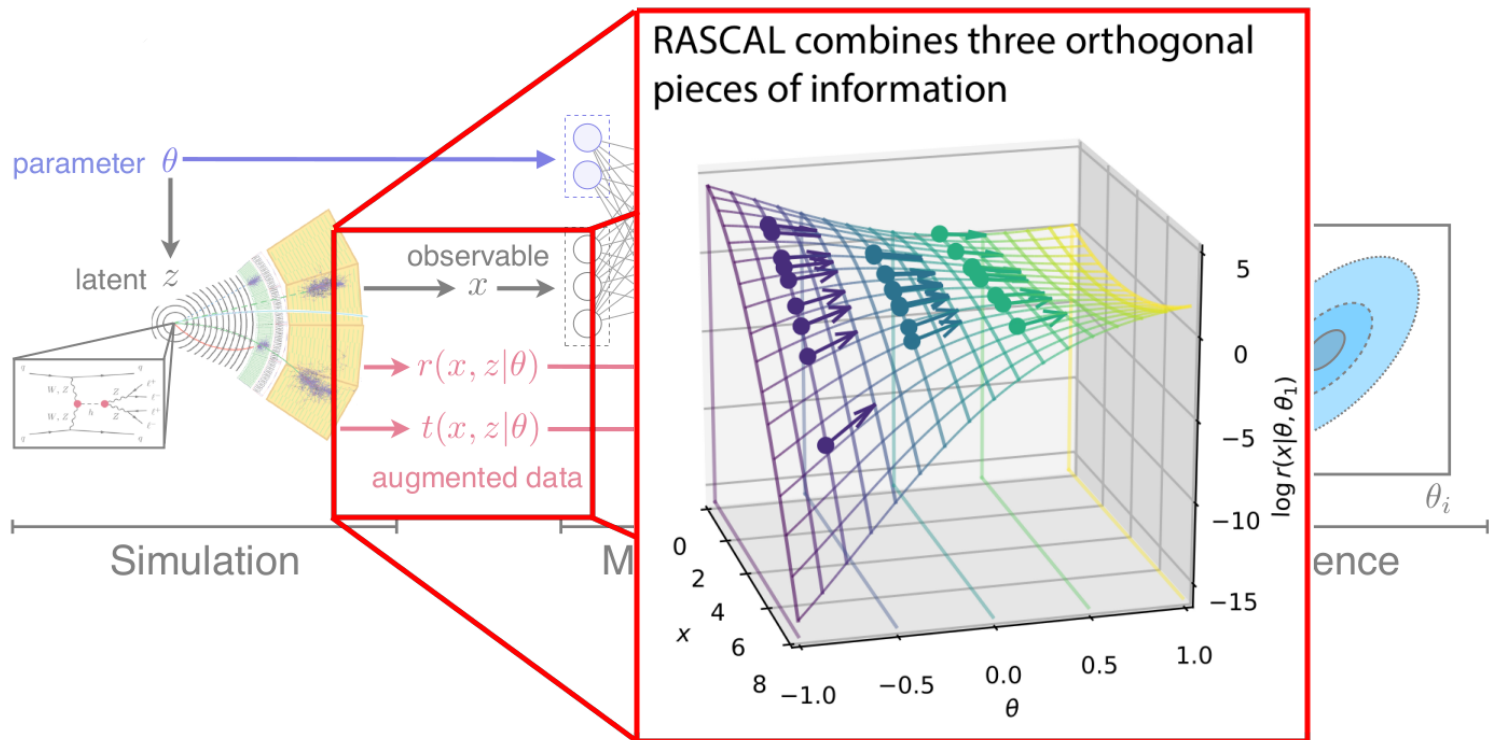
# RASCAL

$$L_{RASCAL} = L_r + L_t$$



# RASCAL

$$L_{RASCAL} = L_r + L_t$$



# SALLY (= optimal compression)

The likelihood ratio  $r$  relates to the **score**

$$t(x|\theta_{\text{ref}}) = \nabla_{\theta} \log p(x|\theta)|_{\theta_{\text{ref}}} = \nabla_{\theta} r(x|\theta, \theta_{\text{ref}})|_{\theta_{\text{ref}}}.$$

- It quantifies the relative change of the likelihood under infinitesimal changes.
- It can be seen as a **local equivalent of the likelihood ratio**.


In a small patch around  $\theta_{\text{ref}}$ , we have the approximation

$$p_{\text{local}}(x|\theta) = \frac{1}{Z(\theta)} p(t(x|\theta_{\text{ref}})|\theta_{\text{ref}}) \exp(t(x|\theta_{\text{ref}}) \cdot (\theta - \theta_{\text{ref}}))$$

where the score  $t(x|\theta_{\text{ref}})$  are its sufficient statistics. Therefore,

- in the local model the likelihood ratio between  $\theta$  and  $\theta_{\text{ref}}$  only depends on the product between the score and  $\theta - \theta_{\text{ref}}$ .
- That is,  $x$  can be compressed into a single scalar without loss of power.


# Results?



## Experimental setup

- Higgs production in weak boson fusion.
- Goal: constraints on two theory parameters.


$$\mathcal{L} = \mathcal{L}_{SM} + \underbrace{\frac{f_W}{\Lambda^2}} \frac{ig}{2} (D^\mu \phi)^\dagger \sigma^a D^\nu \phi W_{\mu\nu}^a - \underbrace{\frac{f_{WW}}{\Lambda^2}} \frac{g^2}{4} (\phi^\dagger \phi) W_{\mu\nu}^a W^{\mu\nu a}$$



New techniques  
require less data than  
generic ML method  
[CARL, see K. Cranmer, G. Louppe,  
J. Pavez 1506.02169]


Traditional Histogram






## Treat the simulator as a black box

Learn a proxy for inference




Histograms of observables  
Neural density (ratio) estimation




Adversarial variational optimization

## Make use of the inner structure



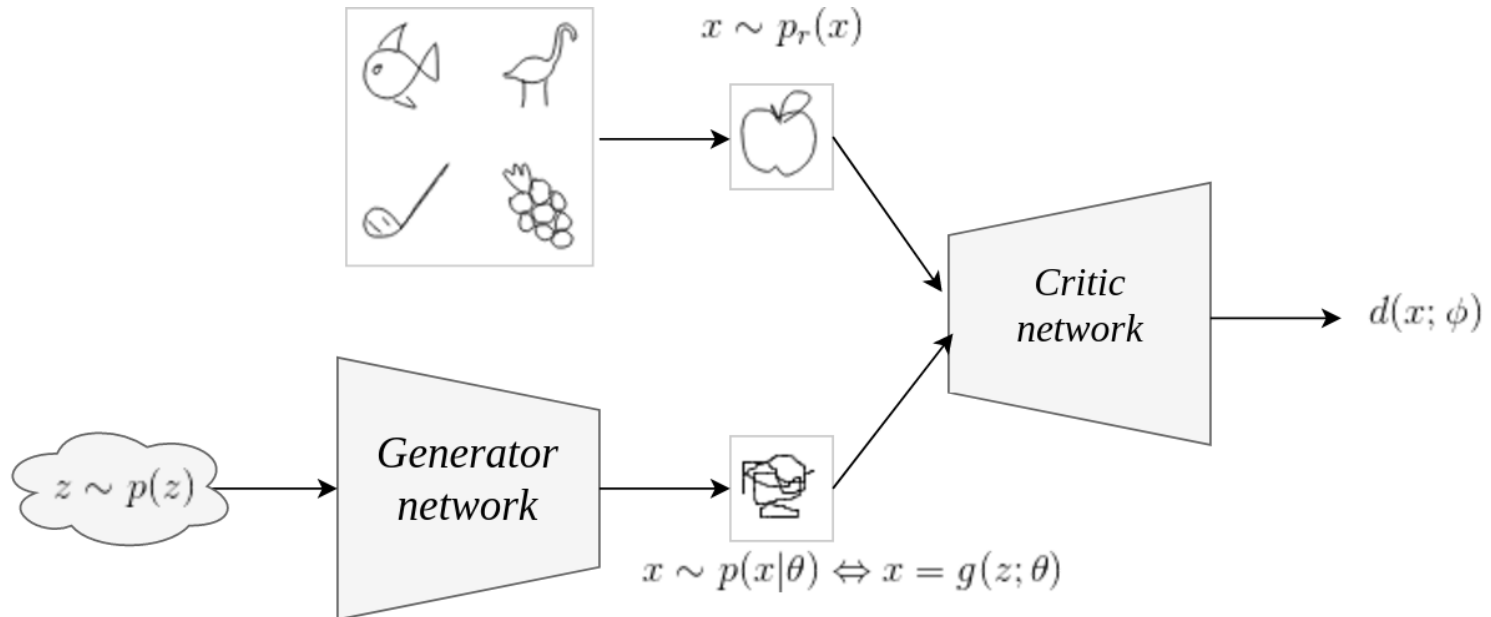
Mining gold from implicit models

Learn to control the simulator



Probabilistic programming

# Generative adversarial networks



$$\mathcal{L}_d(\phi) = \mathbb{E}_{\mathbf{x} \sim p_r(\mathbf{x})} [-\log(d(\mathbf{x}; \phi))] + \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} [-\log(1 - d(g(\mathbf{z}; \theta); \phi))]$$

$$\mathcal{L}_g(\theta) = \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} [\log(1 - d(g(\mathbf{z}; \theta); \phi))]$$



Odena et al  
2016



Miyato et al  
2017

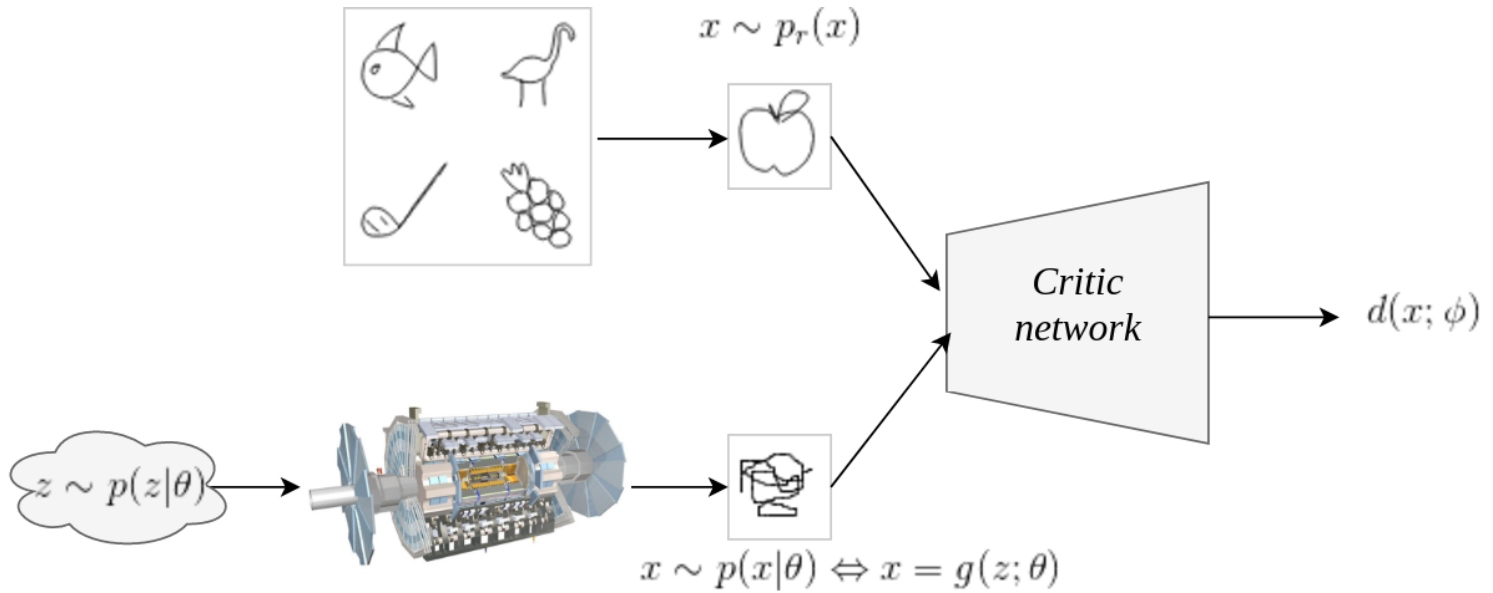


Zhang et al  
2018



Brock et al  
2018

# AVO



Replace  $g$  with an actual scientific simulator!



## Key insights

- Replace the generative network with a non-differentiable forward simulator  $g(\mathbf{z}; \theta)$ .
- Let the neural network critic figure out how to adjust the simulator parameters.
- Combine with variational optimization to bypass the non-differentiability by optimizing upper bounds of the adversarial objectives

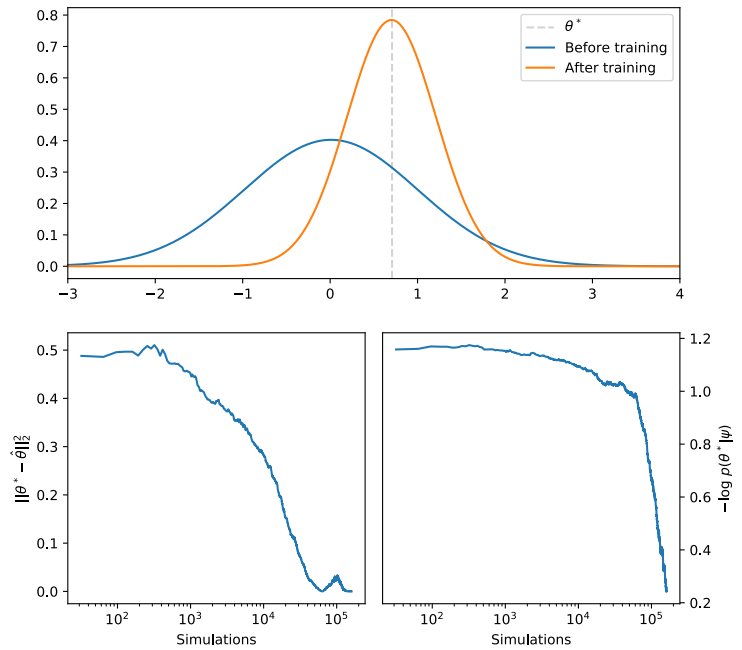
$$U_d(\phi) = \mathbb{E}_{\theta \sim q(\theta; \psi)} [\mathcal{L}_d(\phi)]$$

$$U_g(\psi) = \mathbb{E}_{\theta \sim q(\theta; \psi)} [\mathcal{L}_g(\theta)]$$

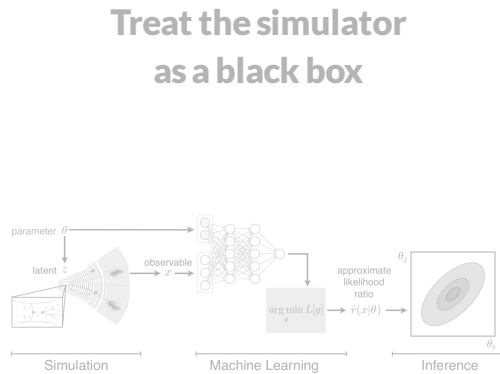
respectively over  $\phi$  and  $\psi$ .




Samples for  $\theta = 0$  (top) vs. samples for  $\theta = 0.81$  (bottom).



**Learn a proxy for inference**




Histograms of observables  
Neural density (ratio) estimation




Adversarial variational optimization

**Make use of the inner structure**



Mining gold from implicit models



Probabilistic programming




# Probabilistic programming

Probabilistic models define a set of **random variables** and their relationships.

- Observed variables
- Unobserved (hidden, latent) variables

Probabilistic **graphical models** use graphs to express conditional dependence.

- Bayesian networks
- Markov random fields



$$p(x, y, z) = p(x)p(y)p(z|x, y)$$

Probabilistic **programming** extends this to ordinary programming with two added constructs:

- Sampling from distributions
- Conditioning random variables by specifying observed values

## Example

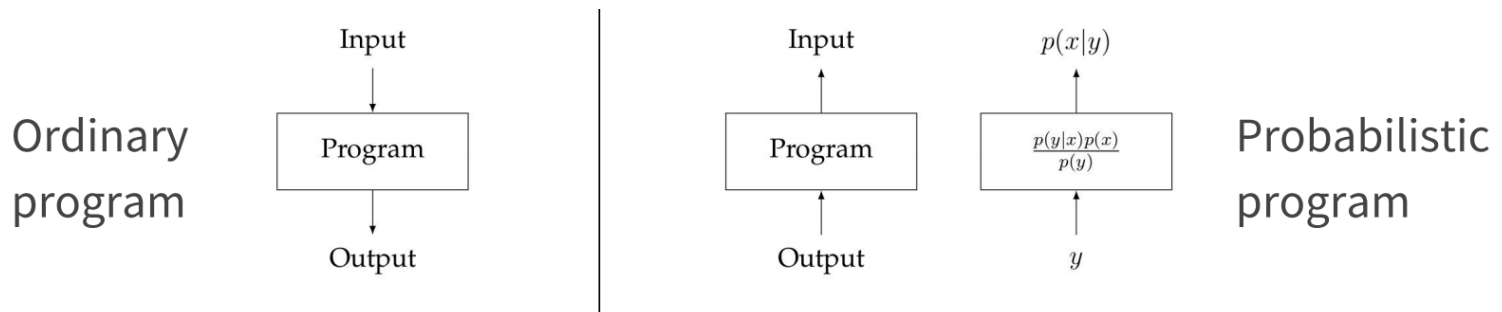
```
bool c1, c2;  
c1 = Bernoulli(0.5);  
c2 = Bernoulli(0.5);  
observe(c1 || c2);  
return(c1, c2);
```

# Inference


With a probabilistic program, we define a joint distribution of **unobserved** and **observed** variables  $p(x, y)$ .

**Inference engines** give us distributions over unobserved variables, given observed variables (data)

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)}$$

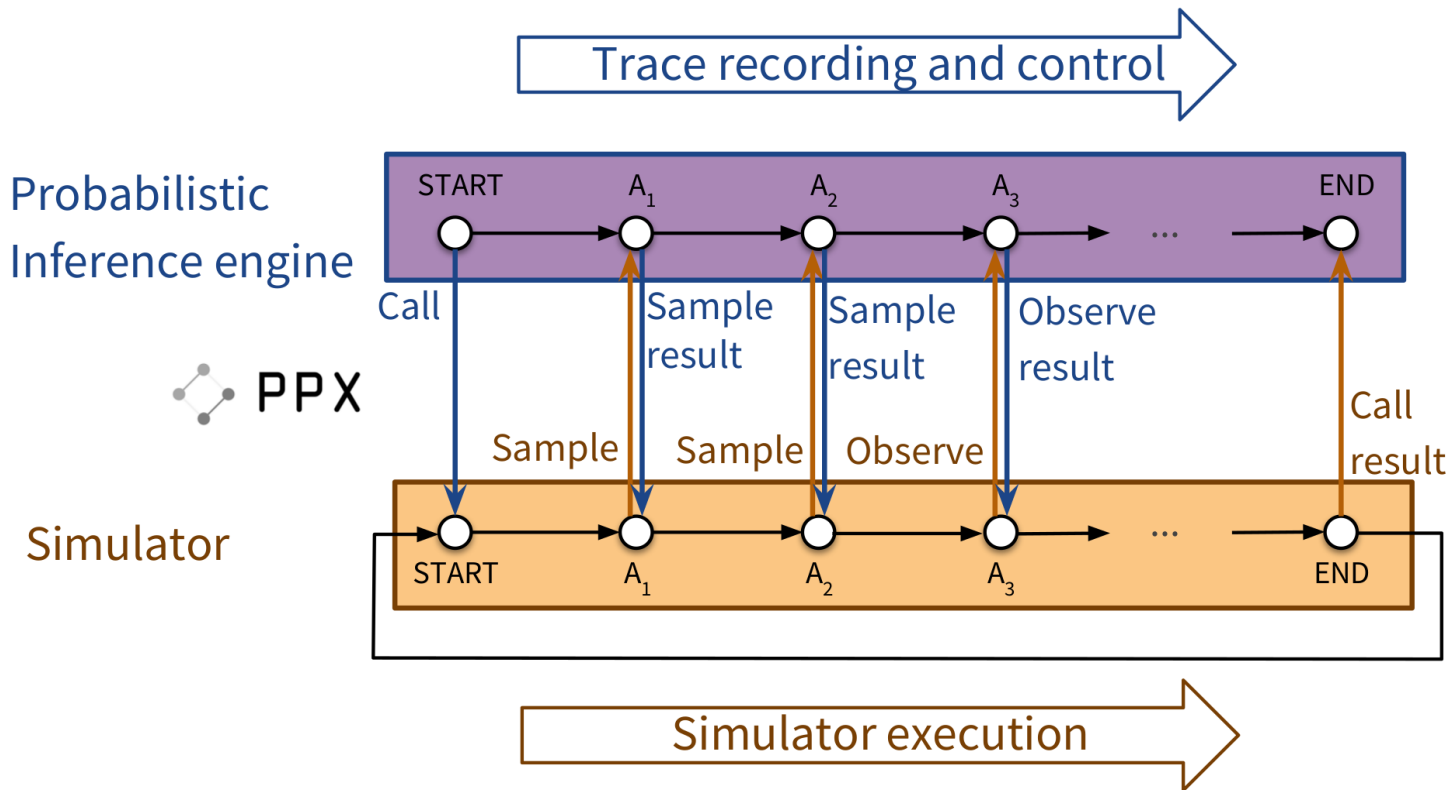


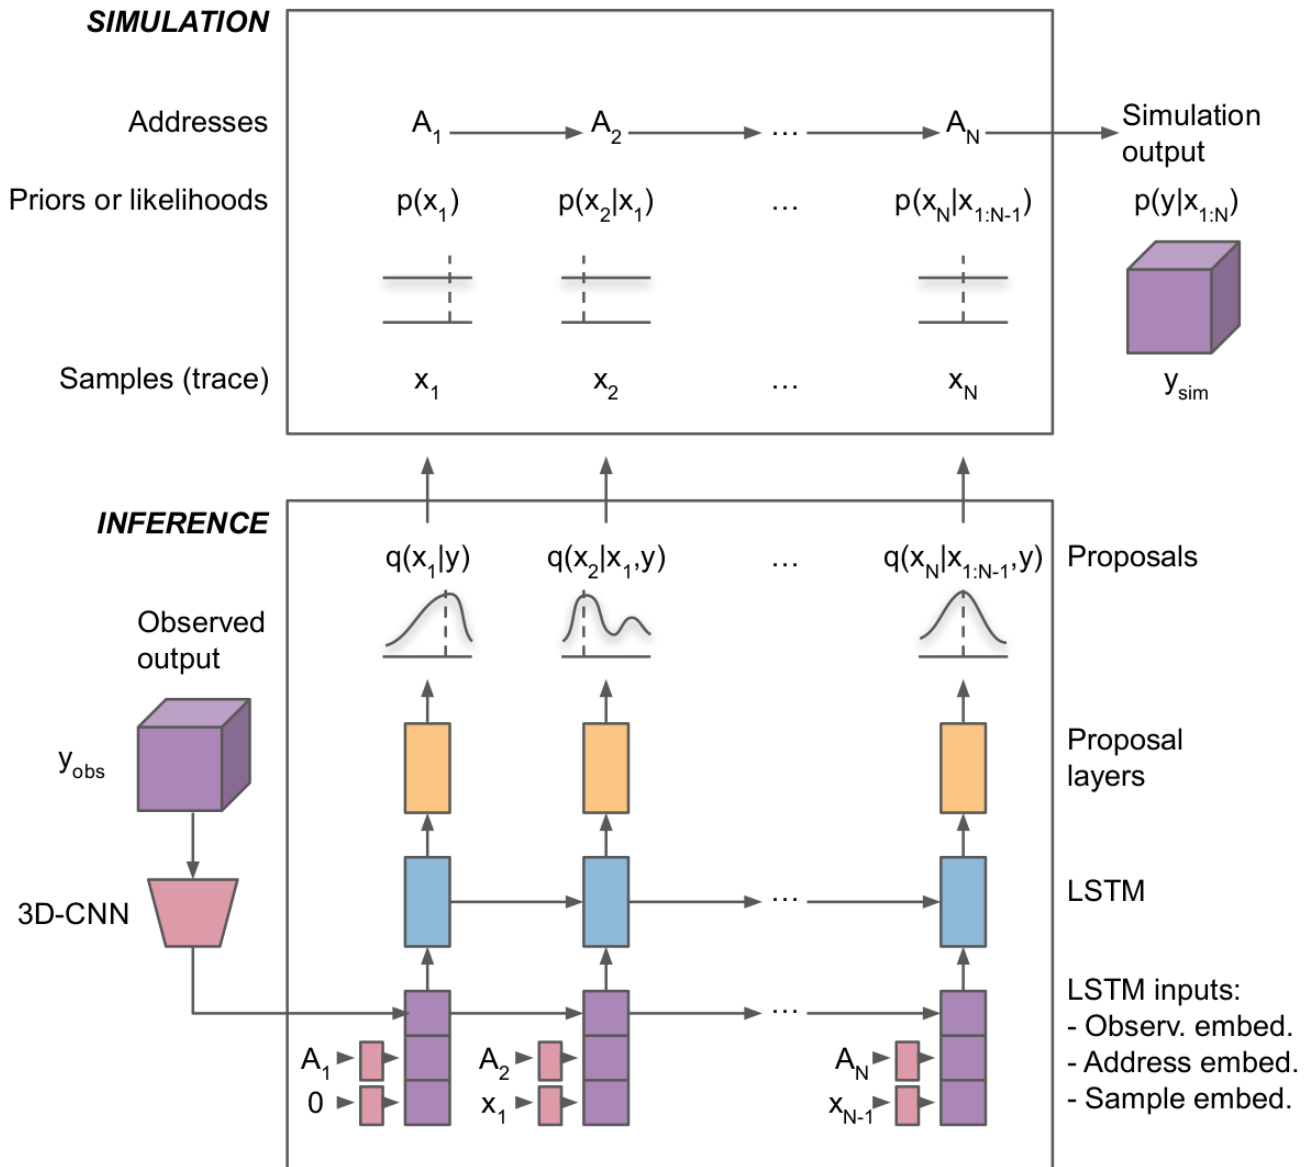
A stochastic simulator implicitly defines a probability distribution by sampling pseudo-random numbers. **Scientific simulators are probabilistic programs!**



## Key insights

Let a neural network take full control of the internals of the simulation program by hijacking all calls to the random number generator.







# Taking control of Sherpa


## Experimental setup

- $\tau$  decay in Sherpa, 38 decay channels, coupled with an approximate calorimeter simulation in C++.
- Observations are 3D calorimeter depositions.
- Latent variables (Monte Carlo truth) of interest: decay channel,  $p_x$ ,  $p_y$ ,  $p_z$  momenta, final state momenta and IDs.

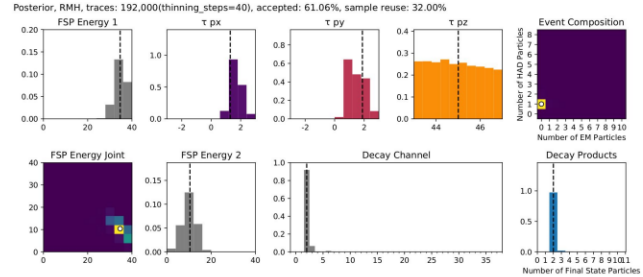




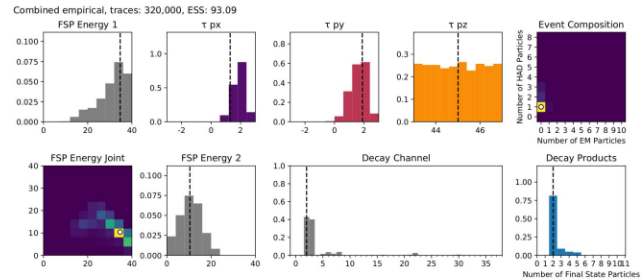
# Inference results



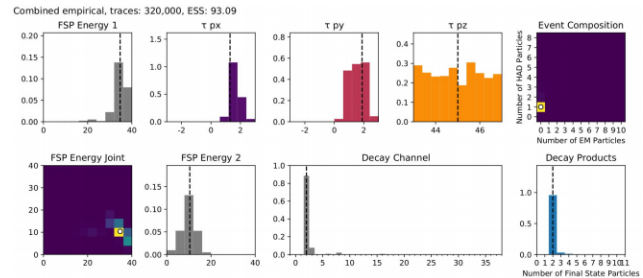
MCMC true posterior  
(7.7M single node)



IC proposal  
from trained NN




IC posterior  
after importance  
weighting





# Interpretability


Latent probabilistic structure of the 10 most frequent trace types:



# Interpretability

Latent probabilistic structure of the 10 most frequent trace types:

```
[forward(xt:: xarray_container<xt:: uvector<double, std:: allocator<double> >, (xt:: layout_type)1, xt:: svector<unsigned long, 4ul, std:: allocator<unsigned long>, true>, xt:: xtensor_expression_tag>)+0x5f; SherpaGenerator:: Generate()+0x36; SHERPA:: Sherpa:: GenerateOneEvent(bool)+0x2fa; SHERPA:: Event_Handler:: GenerateEvent(SHERPA:: eventtype:: code)+0x44d; SHERPA:: Event_Handler:: GenerateHadronDecayEvent(SHERPA:: eventtype:: code&)+0x45f; ATOOLS:: Random:: Get(bool, bool)+0x1d5; probprog_RNG:: Get(bool, bool)+0xf9]_Uniform_1
```

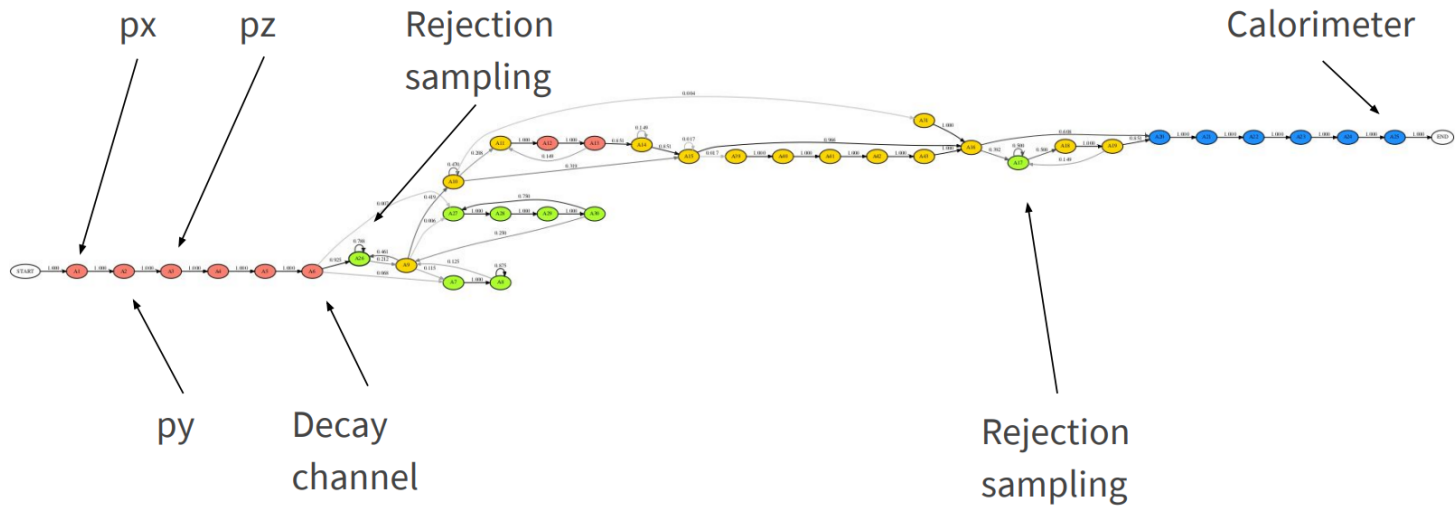


```
[forward(xt:: xarray_container<xt:: uvector<double, std:: allocator<double> >, (xt:: layout_type)1, xt:: svector<unsigned long, 4ul, std:: allocator<unsigned long>, true>, xt:: xtensor_expression_tag>)+0x5f; SherpaGenerator:: Generate()+0x36; SHERPA:: Sherpa:: GenerateOneEvent(bool)+0x2fa; SHERPA:: Event_Handler:: GenerateEvent(SHERPA:: eventtype:: code)+0x44d; SHERPA:: Event_Handler:: GenerateHadronDecayEvent(SHERPA:: eventtype:: code&)+0x982; SHERPA:: Event_Handler:: IterateEventPhases(SHERPA:: eventtype:: code&, double&)+0x1d2; SHERPA:: Hadron_Decays:: Treat(ATOOLS:: Blob_List*, double&)+0x975; SHERPA:: Decay_Handler_Base:: TreatInitialBlob(ATOOLS:: Blob*, METOOLS:: Amplitude2_Tensor*, std:: vector<ATOOLS:: Particle*, std:: allocator<ATOOLS:: Particle*> > const&)+0x1ab1; SHERPA:: Hadron_Decay_Handler:: CreateDecayBlob(ATOOLS:: Particle*)+0x4cd; PHASIC:: Decay_Table:: Select() const+0x9d7; ATOOLS:: Random:: GetCategorical(std:: vector<double, std:: allocator<double> > const&, bool, bool)+0x1a5; probprog_RNG:: GetCategorical(std:: vector<double, std:: allocator<double> > const&, bool, bool)+0x111]_Categorical(length_categories:38)_1
```




# Interpretability

Latent probabilistic structure of the 25 most frequent trace types:




# Interpretability

Latent probabilistic structure of the 100 most frequent trace types:




# Interpretability

Latent probabilistic structure of the 250 most frequent trace types:

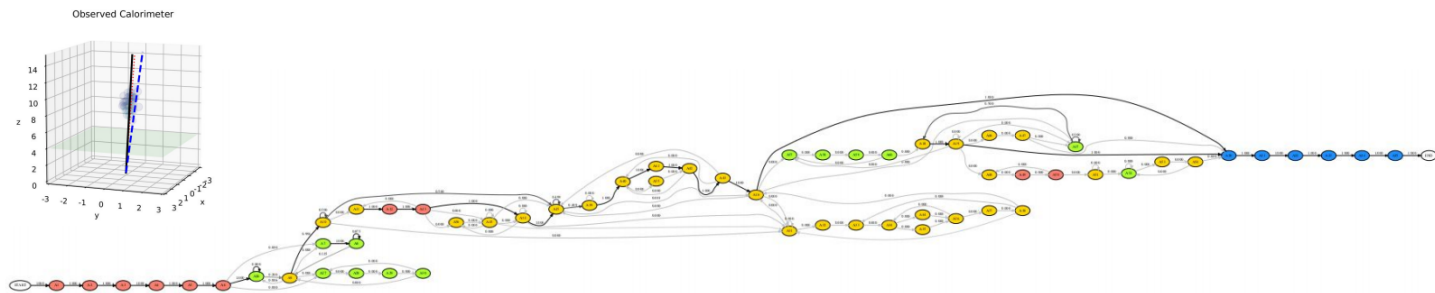




# Interpretability



(a) Prior execution  $p(\mathbf{x})$ .



(b) Posterior execution  $p(\mathbf{x}|\mathbf{y})$  conditioned on a given calorimeter observation  $\mathbf{y}$ .

# Summary

# Summary

- Much of modern science is based on "likelihood-free" simulations.
- The likelihood-ratio is central to many statistical inference procedures.
- Supervised learning enables likelihood-ratio estimation.
- Better likelihood-ratio estimates can be achieved by mining simulators.
- Probabilistic programming enables posterior inference in scientific simulators.

# Collaborators



Kyle Cranmer



Juan Pavez



Johann  
Brehmer



Joeri Hermans



Lukas  
Heinrich



Atılım Güneş  
Baydin



Wahid Bhimji



Frank Wood

# References

- Stoye, M., Brehmer, J., Louppe, G., Pavez, J., & Cranmer, K. (2018). Likelihood-free inference with an improved cross-entropy estimator. arXiv preprint arXiv:1808.00973.
- Baydin, A. G., Heinrich, L., Bhimji, W., Gram-Hansen, B., Louppe, G., Shao, L., ... & Wood, F. (2018). Efficient Probabilistic Inference in the Quest for Physics Beyond the Standard Model. arXiv preprint arXiv:1807.07706.
- Brehmer, J., Louppe, G., Pavez, J., & Cranmer, K. (2018). Mining gold from implicit models to improve likelihood-free inference. arXiv preprint arXiv:1805.12244.
- Brehmer, J., Cranmer, K., Louppe, G., & Pavez, J. (2018). Constraining Effective Field Theories with Machine Learning. arXiv preprint arXiv:1805.00013.
- Brehmer, J., Cranmer, K., Louppe, G., & Pavez, J. (2018). A Guide to Constraining Effective Field Theories with Machine Learning. arXiv preprint arXiv:1805.00020.
- Casado, M. L., Baydin, A. G., Rubio, D. M., Le, T. A., Wood, F., Heinrich, L., ... & Bhimji, W. (2017). Improvements to Inference Compilation for Probabilistic Programming in Large-Scale Scientific Simulators. arXiv preprint arXiv:1712.07901.
- Louppe, G., Hermans, J., & Cranmer, K. (2017). Adversarial Variational Optimization of Non-Differentiable Simulators. arXiv preprint arXiv:1707.07113.
- Cranmer, K., Pavez, J., & Louppe, G. (2015). Approximating likelihood ratios with calibrated discriminative classifiers. arXiv preprint arXiv:1506.02169.

The end.