

Parameter inference for particle physics

using Magic!

Gert Kluge

PhD candidate at the University of Oslo

CERN School of Computing 2022 – Krakow
7th of September 2022

Particle physics in a nutshell: the search for $p(\boldsymbol{\vartheta}|\boldsymbol{x})$


Bayes theorem:

$$p(\boldsymbol{\vartheta}|\boldsymbol{x}) = \frac{p(\boldsymbol{x}|\boldsymbol{\vartheta})}{p(\boldsymbol{x})} p(\boldsymbol{\vartheta})$$

Particle physics in a nutshell: the search for $p(\boldsymbol{\vartheta}|\boldsymbol{x})$

Bayes theorem:

$$p(\boldsymbol{\vartheta}|\boldsymbol{x}) = \frac{p(\boldsymbol{x}|\boldsymbol{\vartheta})}{p(\boldsymbol{x})} p(\boldsymbol{\vartheta})$$




$$\boldsymbol{\vartheta} = \begin{matrix} m \\ \sigma \end{matrix}$$

Particle physics in a nutshell: the search for $p(\boldsymbol{\vartheta}|\boldsymbol{x})$

Bayes theorem:

$$p(\boldsymbol{\vartheta}|\boldsymbol{x}) = \frac{p(\boldsymbol{x}|\boldsymbol{\vartheta})}{p(\boldsymbol{x})} p(\boldsymbol{\vartheta})$$

$\boldsymbol{x} =$



$\boldsymbol{\vartheta} =$


φ_0
 γ
 E_{cut}

Particle physics in a nutshell: the search for $p(\boldsymbol{\vartheta}|\boldsymbol{x})$

Bayes theorem:



$$p(\boldsymbol{\vartheta}|\boldsymbol{x}) = \frac{p(\boldsymbol{x}|\boldsymbol{\vartheta})}{p(\boldsymbol{x})} p(\boldsymbol{\vartheta})$$

$\boldsymbol{x} =$



$\boldsymbol{\vartheta} =$

φ_0
 γ
 E_{cut}



Particle physics in a nutshell: the search for $p(\boldsymbol{\vartheta}|\boldsymbol{x})$

Bayes theorem:

$$p(\boldsymbol{\vartheta}|\boldsymbol{x}) = \frac{p(\boldsymbol{x}|\boldsymbol{\vartheta})}{p(\boldsymbol{x})} p(\boldsymbol{\vartheta})$$


$\boldsymbol{\vartheta}$ = parameters of interest (m and g)

Marginal likelihood
(likelihood integrated over nuisance parameters)

Prior ("a priori" assumption)



Evidence = $\int d\boldsymbol{\theta} p(\boldsymbol{x}|\boldsymbol{\theta})$

$\boldsymbol{x} =$

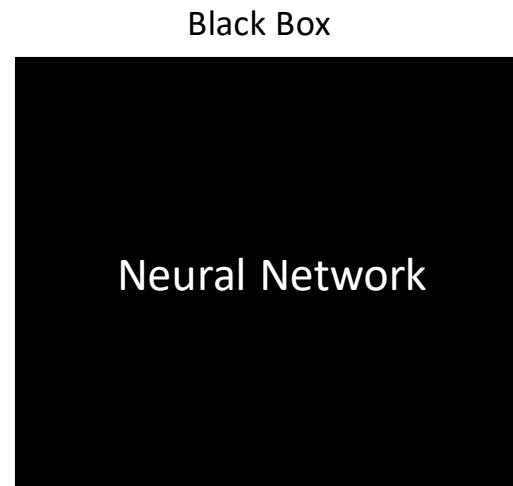


$\boldsymbol{\vartheta} =$

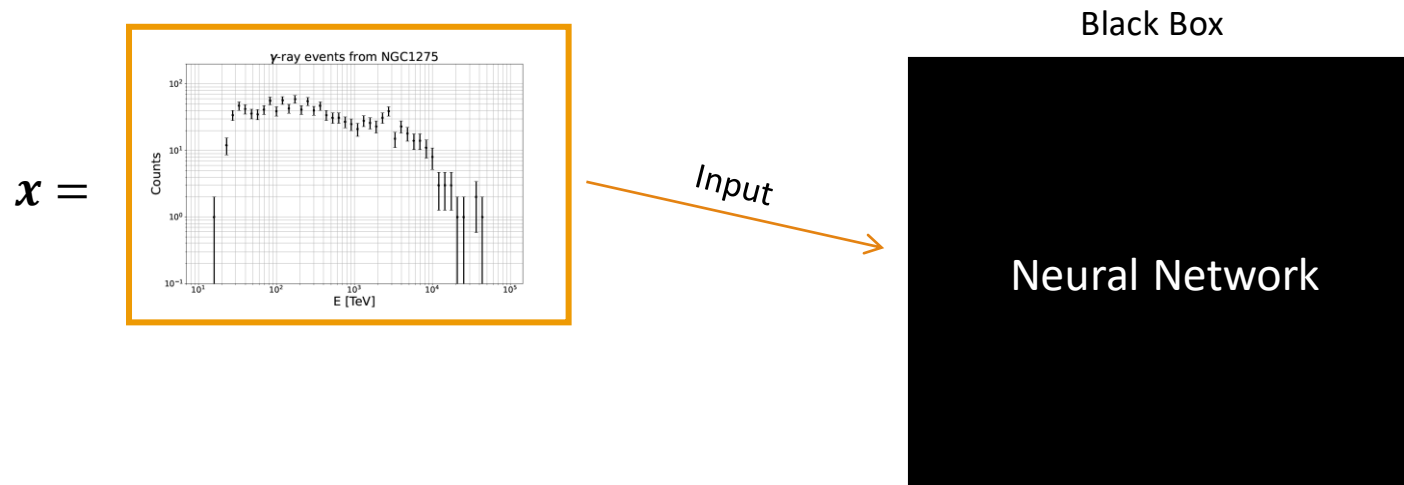
φ_0
 γ
 E_{cut}



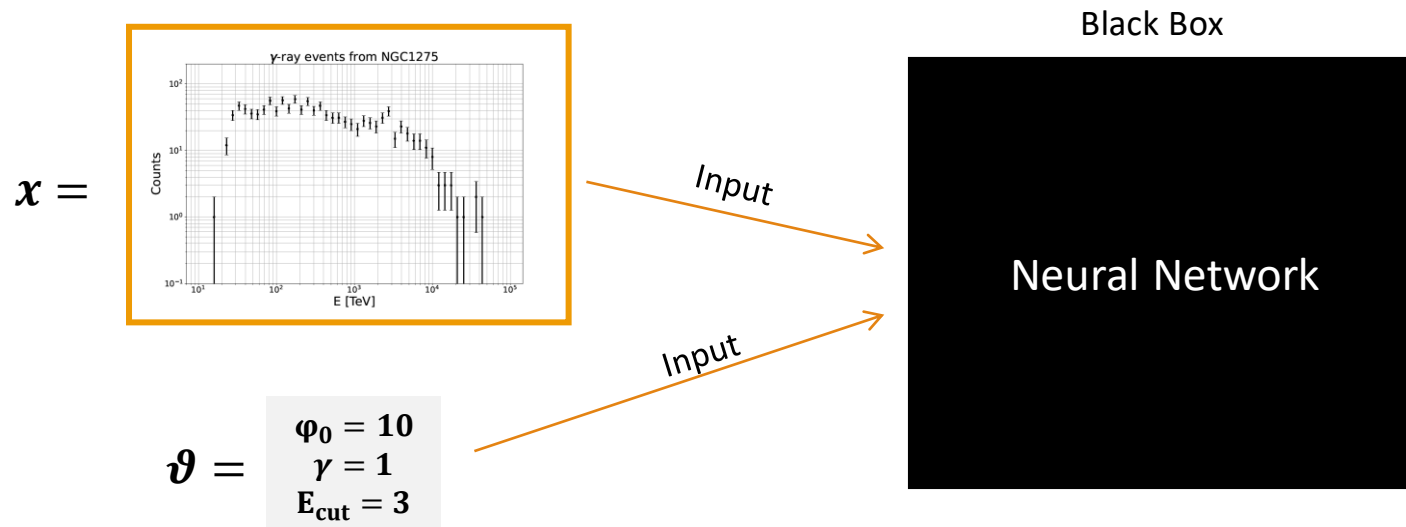
Neural networks can approximate posteriors!



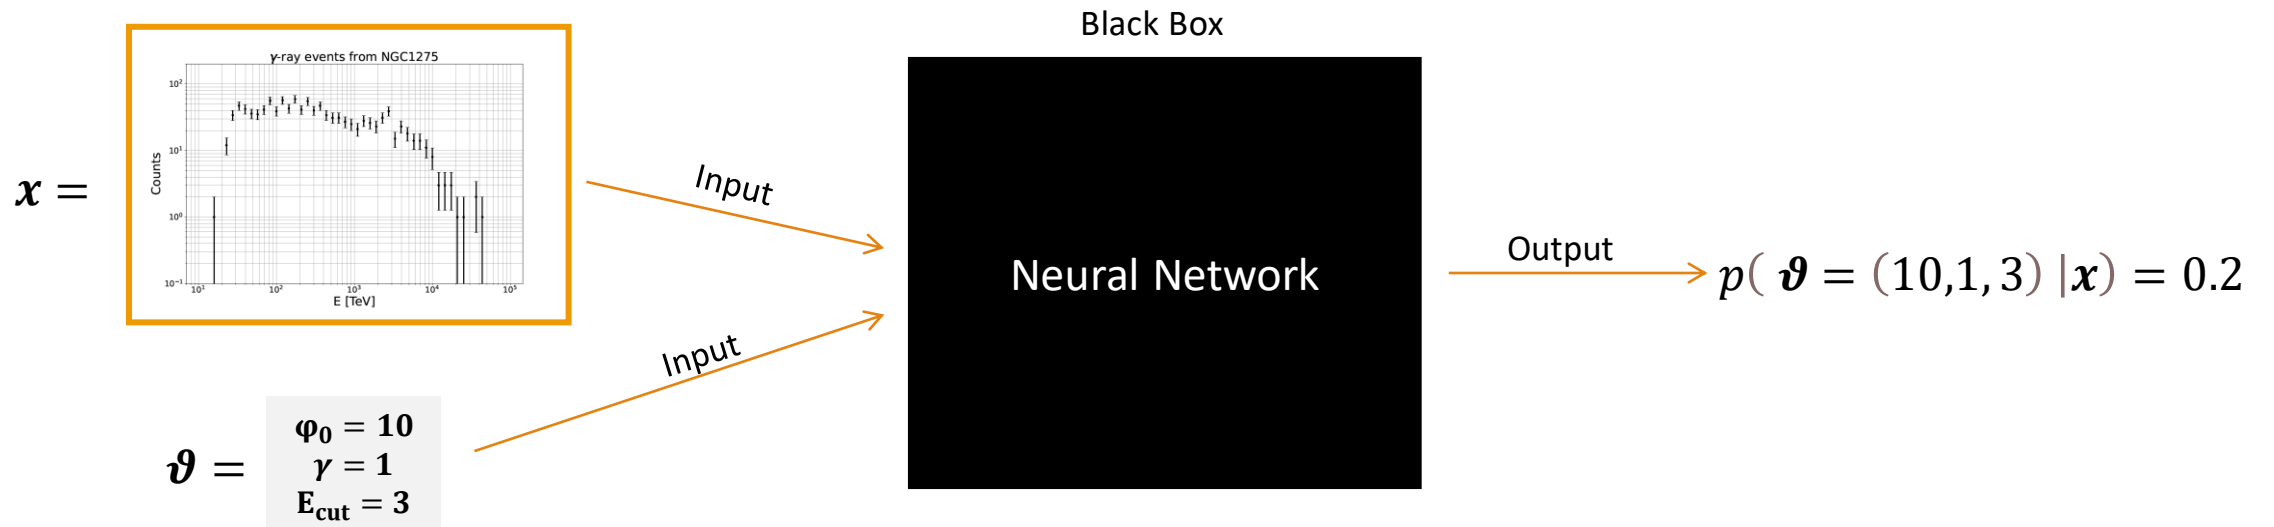
Neural networks can approximate posteriors!




Neural networks can approximate posteriors!



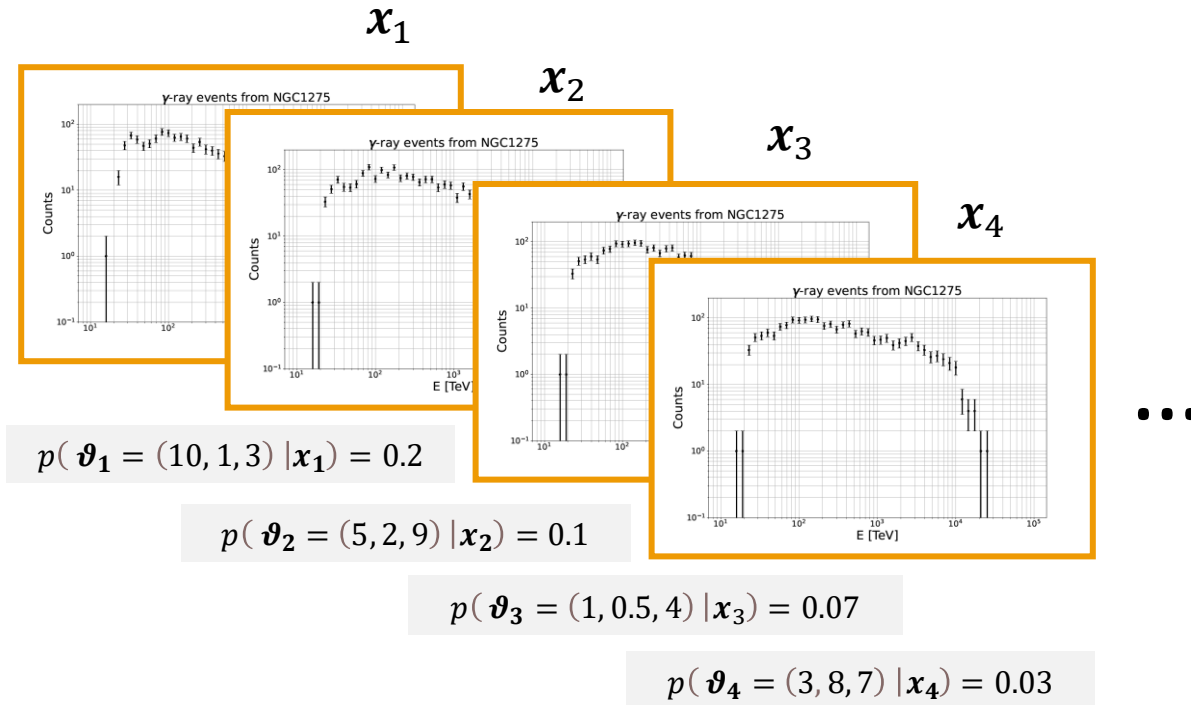
Neural networks can approximate posteriors!



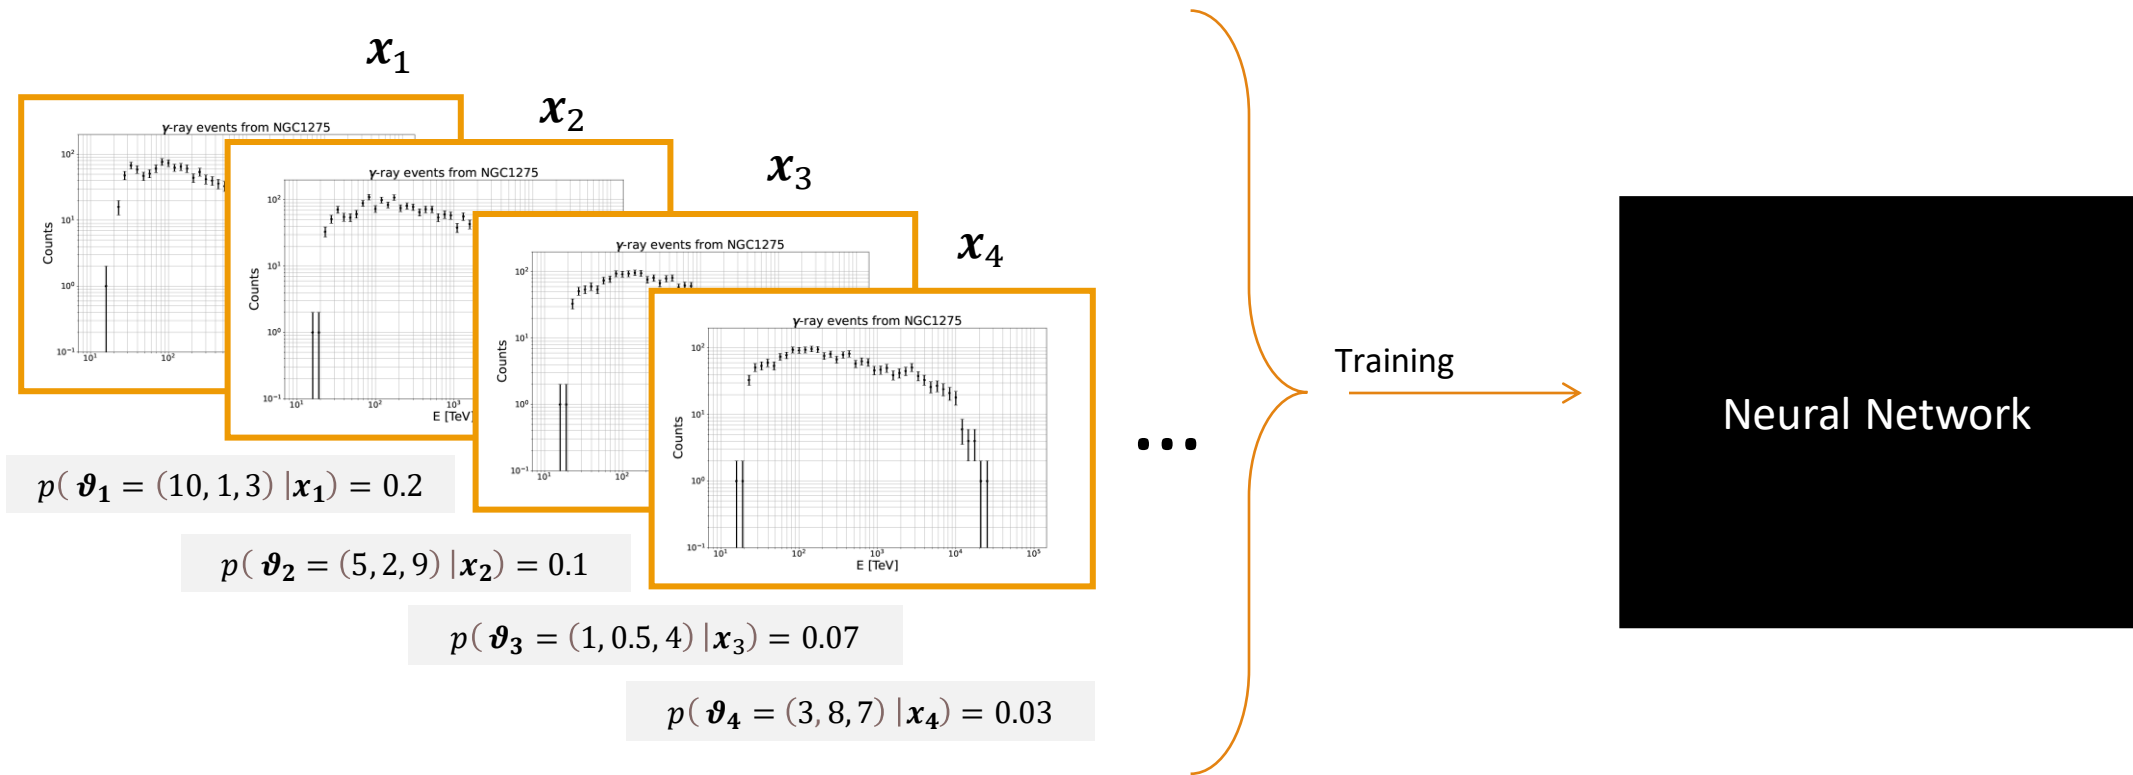
A network learns to do approximations by example:



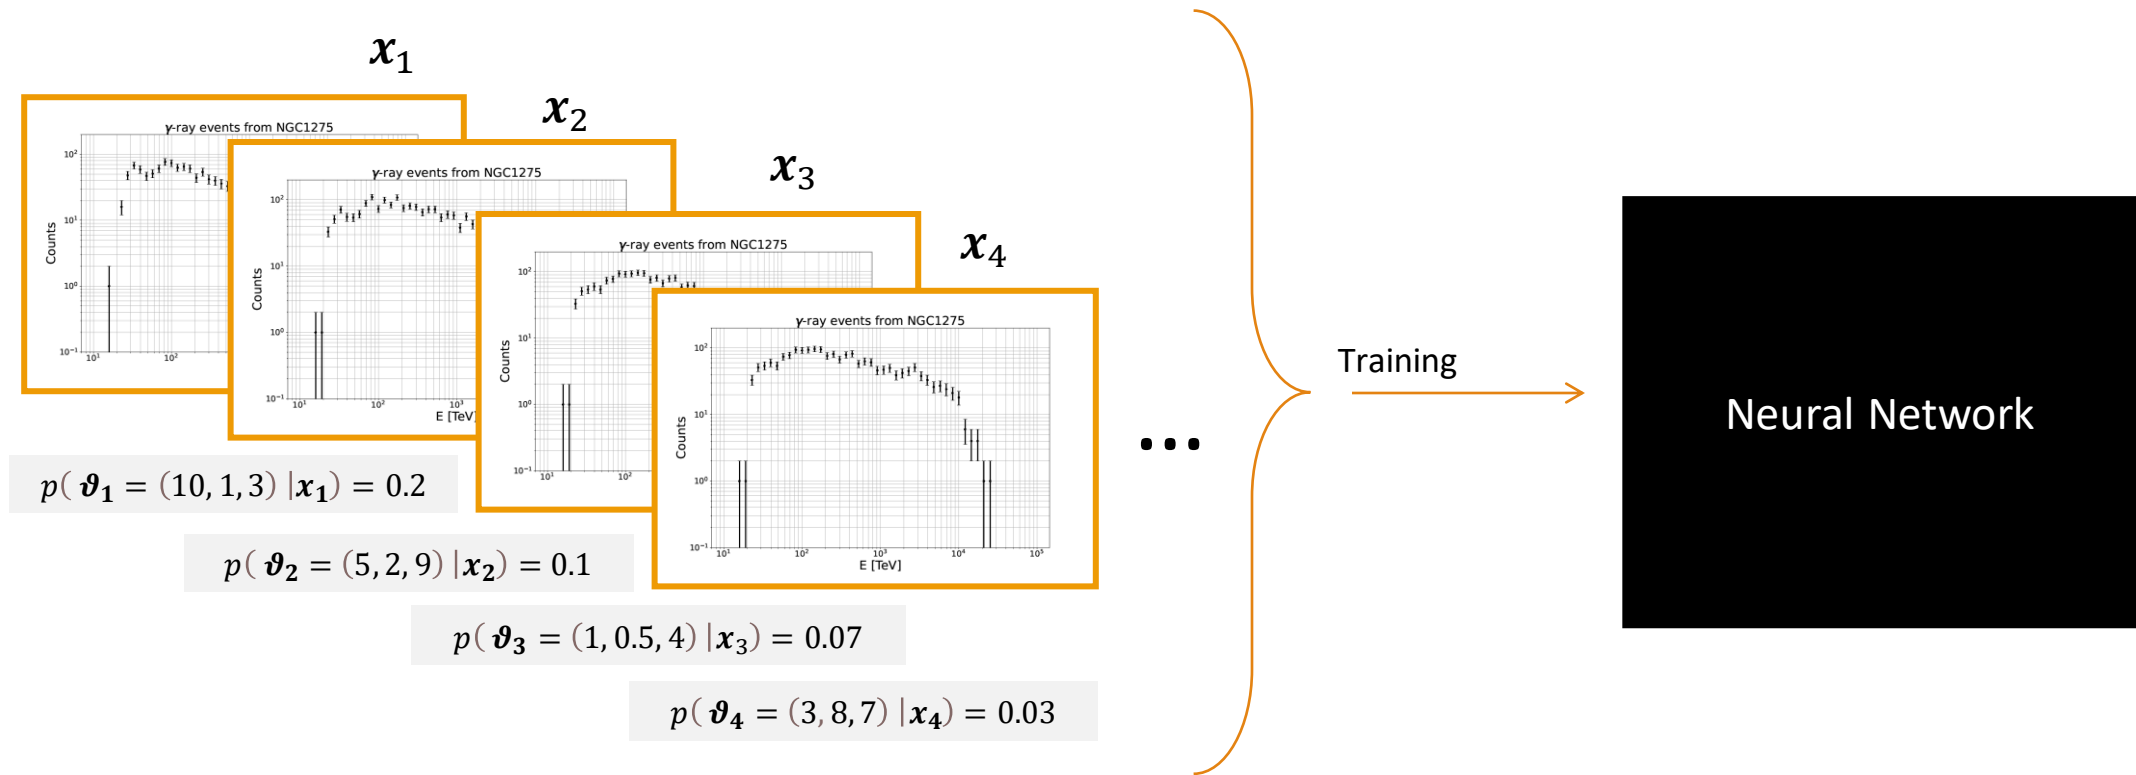
A network learns to do approximations by example:



A network learns to do approximations by example:




A network learns to do approximations by example:




... but this assumes that we can already calculate the posterior!


We can “trick” a network to learn the posterior implicitly



We can “trick” a network to learn the posterior implicitly



We can “trick” a network to learn the posterior implicitly




We can “trick” a network to learn the posterior implicitly

True

The observations are simulated according to the parameter values


x_1



$\varphi_0 = 10$
 $\gamma = 1$
 $E_{\text{cut}} = 3$

True


x_2



$\varphi_0 = 7$
 $\gamma = 8$
 $E_{\text{cut}} = 2$

True


x_3



$\varphi_0 = 2$
 $\gamma = 4$
 $E_{\text{cut}} = 3$

True

x_4



$\varphi_0 = 1.6$
 $\gamma = 20$
 $E_{\text{cut}} = 9$


True

...

False

The parameter values are chosen independently from the simulations


x_5



$\varphi_0 = 16$
 $\gamma = 1$
 $E_{\text{cut}} = 5$

False


x_6



$\varphi_0 = 2$
 $\gamma = 1$
 $E_{\text{cut}} = 9$

False


x_7



$\varphi_0 = 2.5$
 $\gamma = 18$
 $E_{\text{cut}} = 2$

False

x_8



$\varphi_0 = 5$
 $\gamma = 5.6$
 $E_{\text{cut}} = 2$


False

...

We can “trick” a network to learn the posterior implicitly


True

The observations are simulated according to the parameter values

 x_1 


$\varphi_0 = 10$
 $\gamma = 1$
 $E_{\text{cut}} = 3$

True

 x_2 


$\varphi_0 = 7$
 $\gamma = 8$
 $E_{\text{cut}} = 2$

True

 x_3 

$\varphi_0 = 2$
 $\gamma = 4$
 $E_{\text{cut}} = 3$


True

 x_4 

$\varphi_0 = 1.6$
 $\gamma = 20$
 $E_{\text{cut}} = 9$


True

...

 x_5 


$\varphi_0 = 16$
 $\gamma = 1$
 $E_{\text{cut}} = 5$

False

 x_6 


$\varphi_0 = 2$
 $\gamma = 1$
 $E_{\text{cut}} = 9$

False

 x_7 

$\varphi_0 = 2.5$
 $\gamma = 18$
 $E_{\text{cut}} = 2$

False

 x_8 

$\varphi_0 = 5$
 $\gamma = 5.6$
 $E_{\text{cut}} = 2$

False

...

Training

Neural Network

False


The parameter values are chosen independently from the simulations

We can “trick” a network to learn the posterior implicitly

True

The observations are simulated according to the parameter values


x_1



$\varphi_0 = 10$
 $\gamma = 1$
 $E_{\text{cut}} = 3$

True


x_2



$\varphi_0 = 7$
 $\gamma = 8$
 $E_{\text{cut}} = 2$

True


x_3



$\varphi_0 = 2$
 $\gamma = 4$
 $E_{\text{cut}} = 3$

True

x_4




$\varphi_0 = 1.6$
 $\gamma = 20$
 $E_{\text{cut}} = 9$

True

...


x_5



$\varphi_0 = 16$
 $\gamma = 1$
 $E_{\text{cut}} = 5$

False


x_6



$\varphi_0 = 2$
 $\gamma = 1$
 $E_{\text{cut}} = 9$

False


x_7



$\varphi_0 = 2.5$
 $\gamma = 18$
 $E_{\text{cut}} = 2$

False

x_8



$\varphi_0 = 5$
 $\gamma = 5.6$
 $E_{\text{cut}} = 2$

False

...

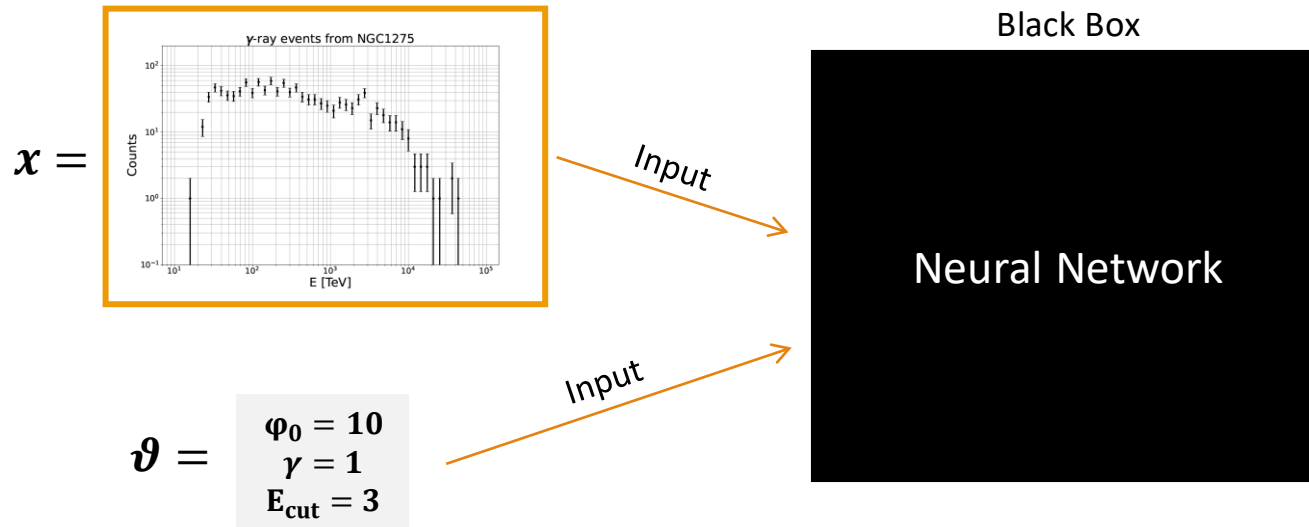
Training

Neural Network

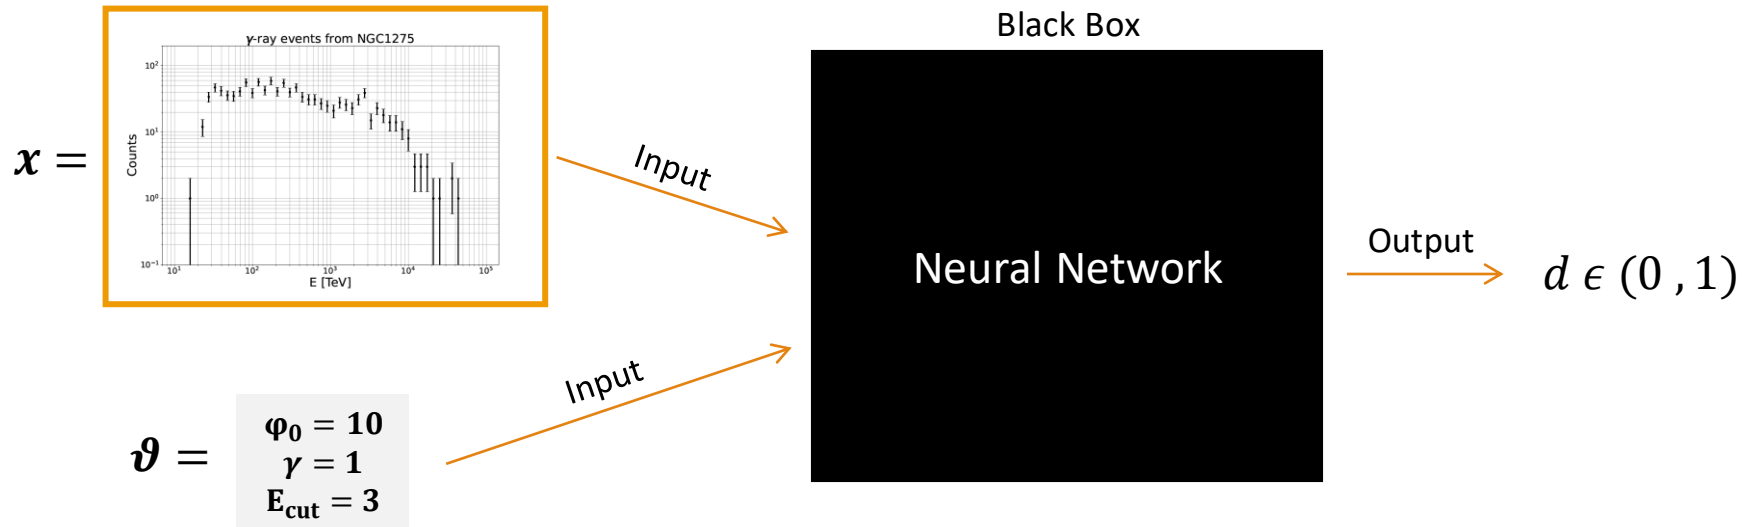
Note: We have to draw all of the parameters of interest from the prior distribution!

Also, the network must use the binary cross-entropy as cost function.

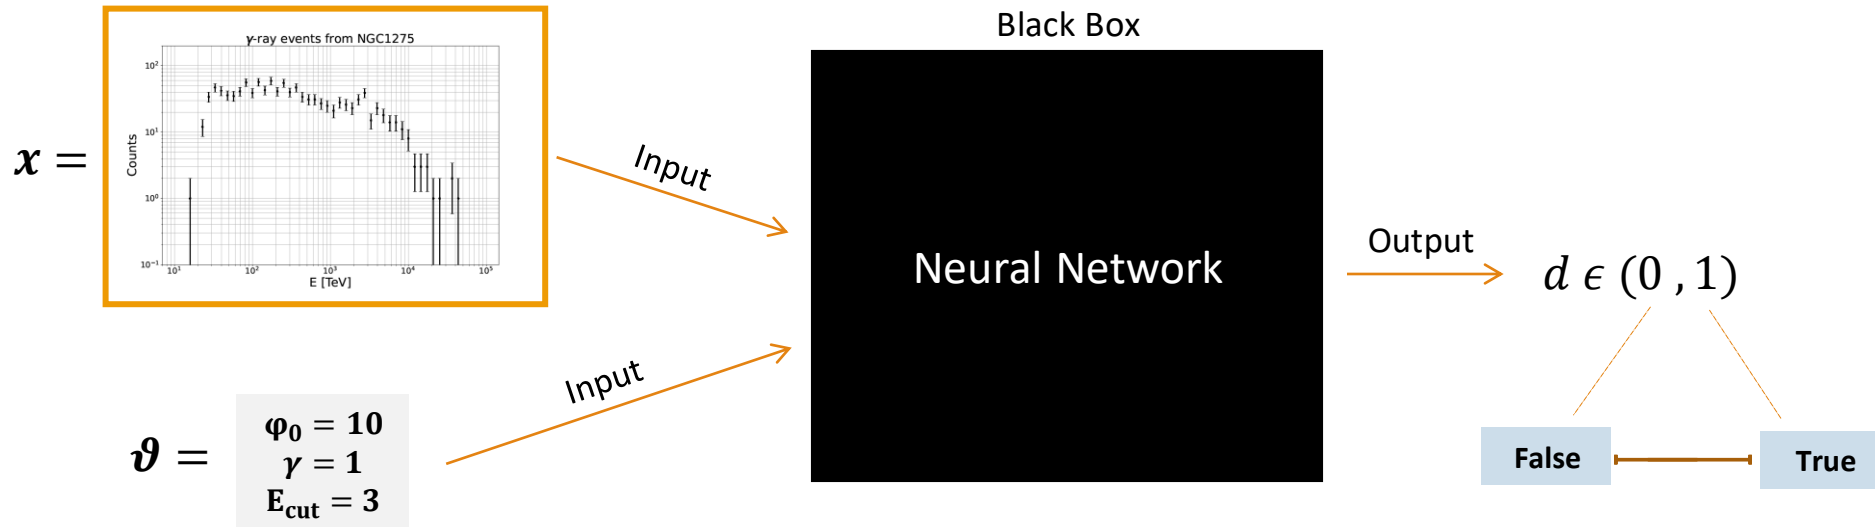
The network now has an implicit understanding of the posterior



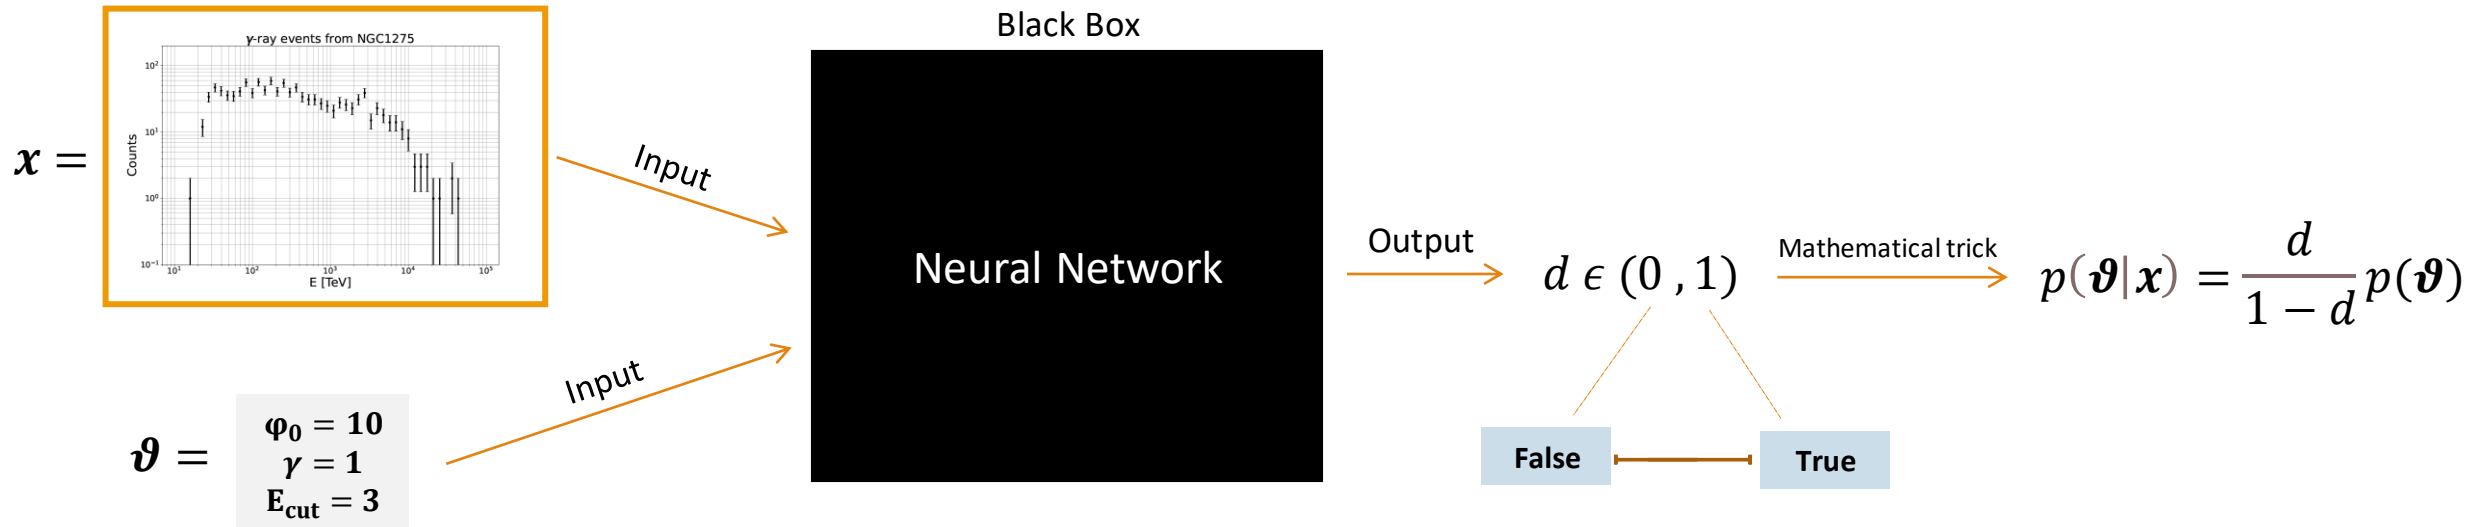
The network now has an implicit understanding of the posterior



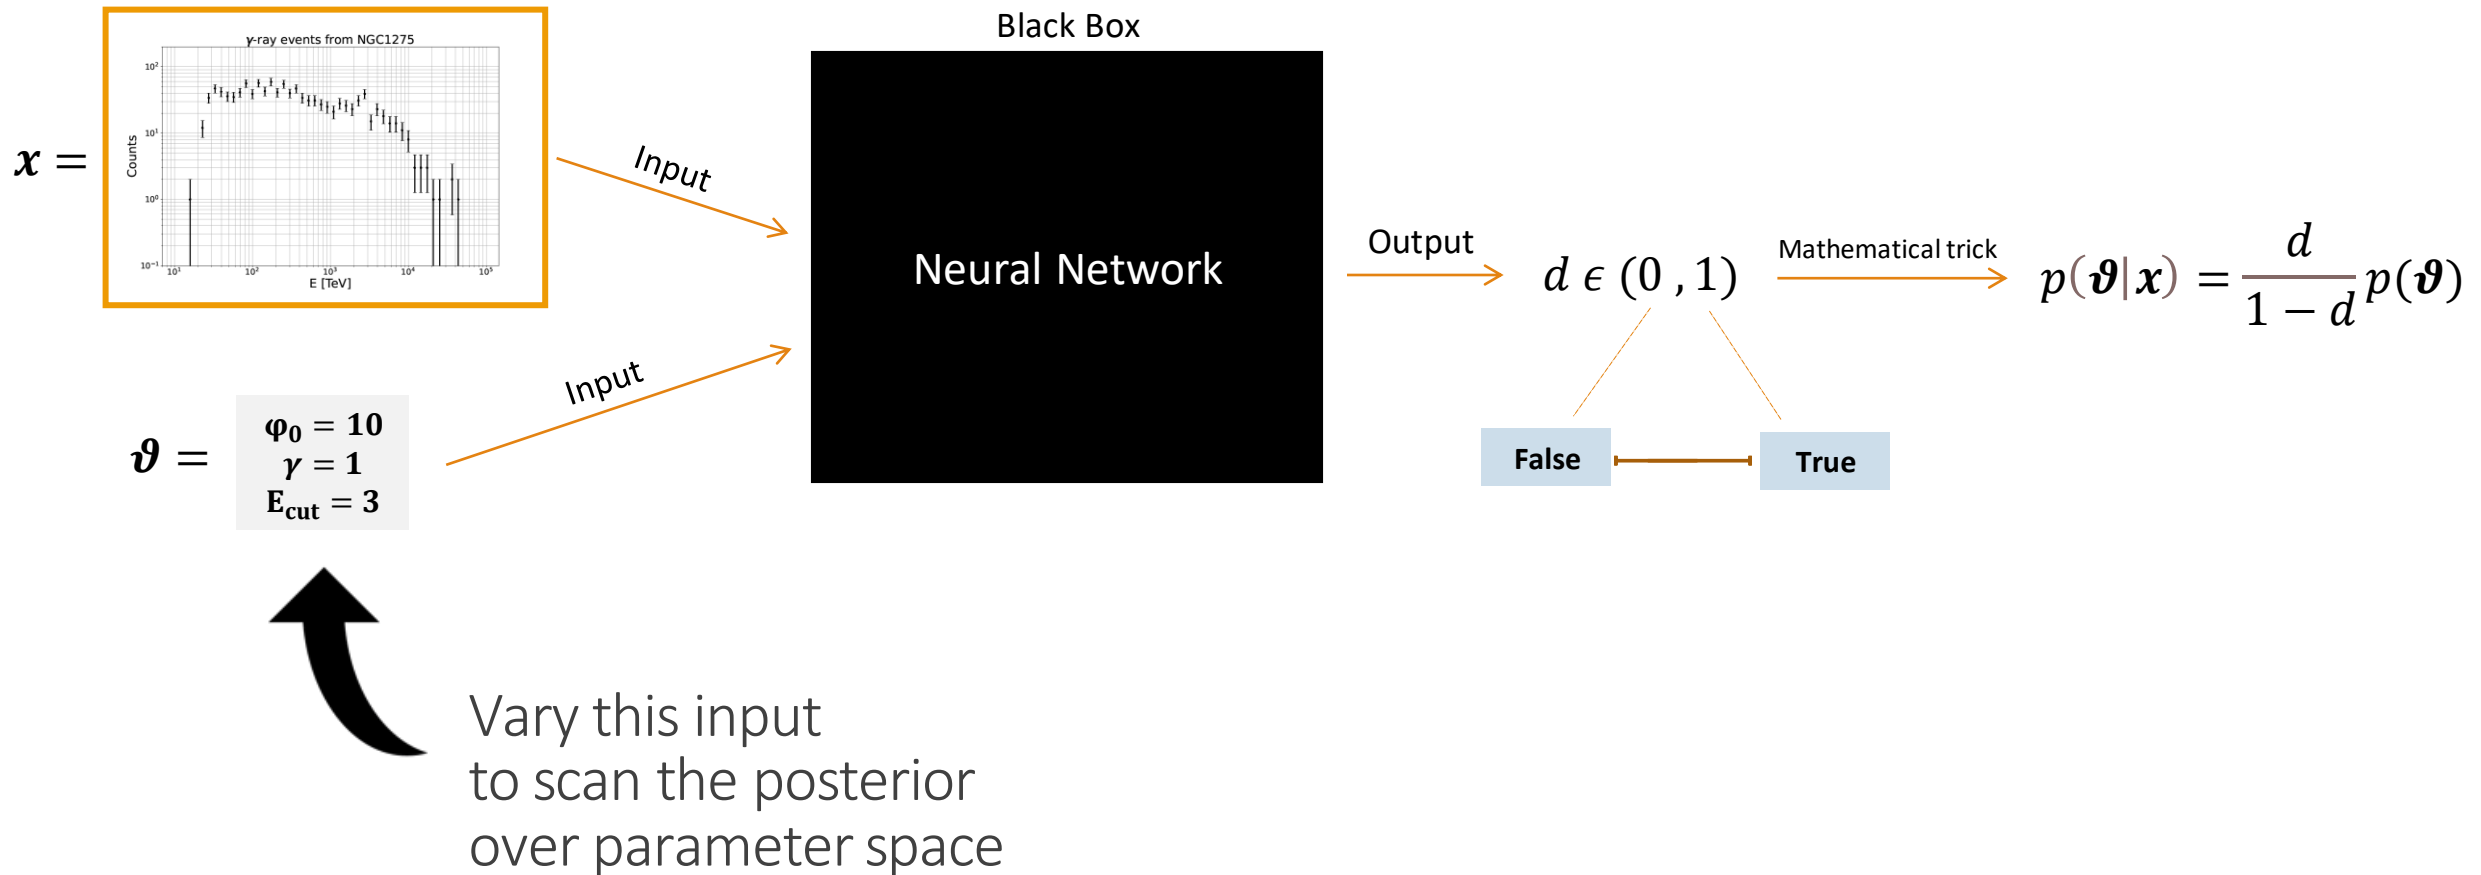
The network now has an implicit understanding of the posterior



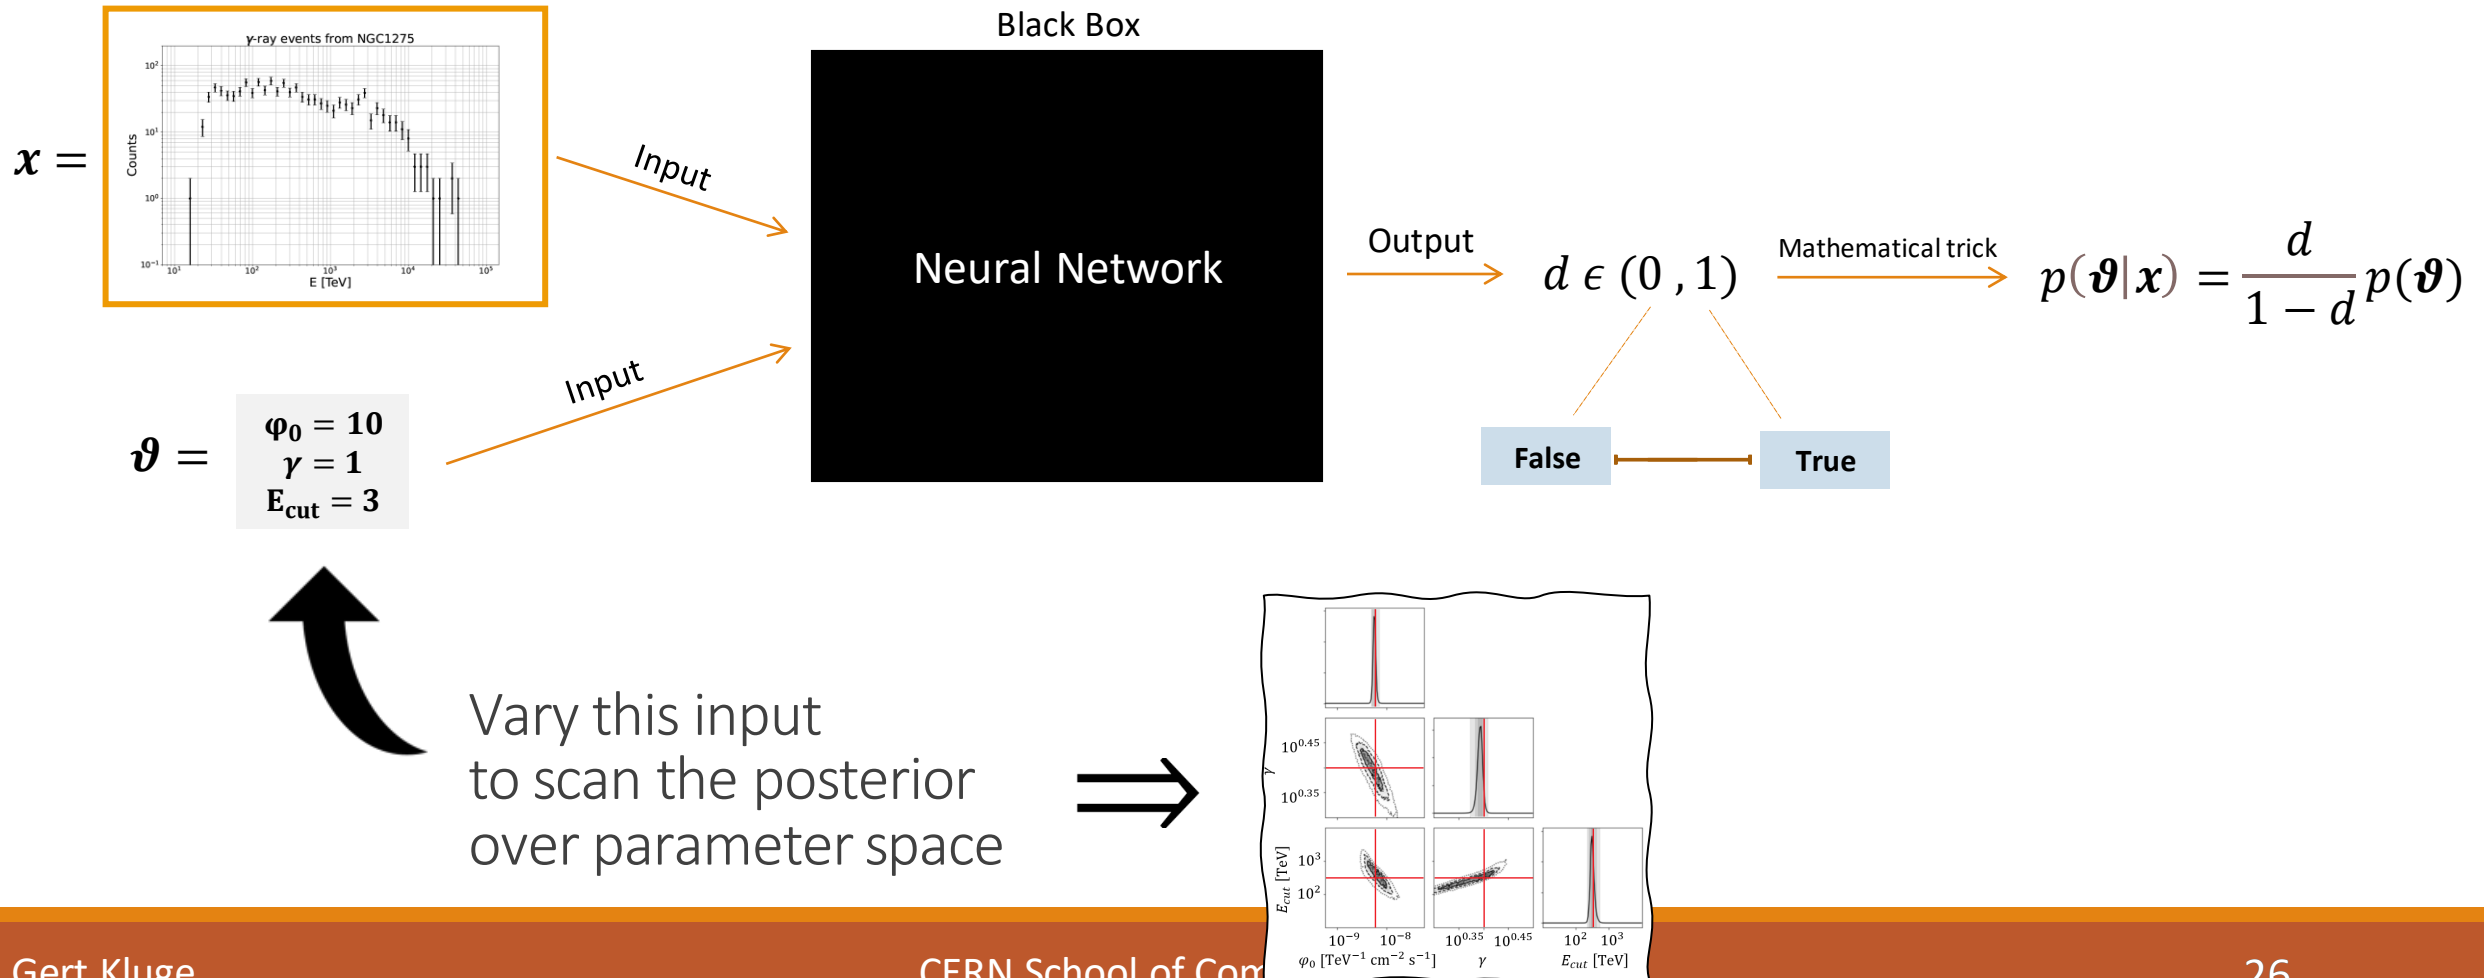
The network now has an implicit understanding of the posterior



The network now has an implicit understanding of the posterior



The network now has an implicit understanding of the posterior



Neural Ratio Estimation: some resources to get you started

- Original paper (to my knowledge) to introduce the concept:

<https://arxiv.org/abs/1903.04057>

- B. K. Miller, A. Cole, P. Forré, G. Louppe, and C. Weniger, “Truncated marginal neural ratio estimation”.

<https://arxiv.org/abs/2107.01214>

- B. K. Miller, A. Cole, G. Louppe, and C. Weniger, “Simulation-efficient marginal posterior estimation with swyft: stop wasting your precious time,”

<https://arxiv.org/abs/2011.13951>

Jun 2020

Likelihood-free MCMC with Amortized Approximate Ratio Estimators

Joeri Hermans¹ Volodimir Begy² Gilles Louppe¹

Abstract

Posterior inference with an intractable likelihood is becoming an increasingly common task in scientific domains which rely on sophisticated computer simulations. Typically, these forward models do not admit tractable densities forcing practitioners to make use of approximations. This work introduces a novel approach to address the intractability of the likelihood and the marginal

ratio of posterior densities between consecutive states in the Markov chain. This allows the posterior to be approximated numerically, provided that the likelihood $p(\mathbf{x}|\theta)$ and the prior $p(\theta)$ are tractable. We consider the equally common and more challenging setting, the so-called likelihood-free setup, in which the likelihood cannot be evaluated in a reasonable amount of time or has no tractable closed-form expression. However, drawing samples from the forward model is possible.

- SWYFT (a package under development that does NRE and more): <https://github.com/undark-lab/swyft>

One take-away:


- You may be able to do Bayesian parameter inference, without having to do impossible integrals, and without putting unrealistic constraints on your nuisance parameters.

One take-away:

- You may be able to do Bayesian parameter inference, without having to do impossible integrals, and without putting unrealistic constraints on your nuisance parameters.
- Basic requirements:
 1. It must be possible to efficiently simulate the experimental observations as a function of physical parameters (of interest and nuisance)
 2. It must be possible for a neural network to adequately distinguish between “matching” and “non-matching” pairs of observations and input parameters.

Backup

Particle physics in a nutshell: find the values of some free parameters

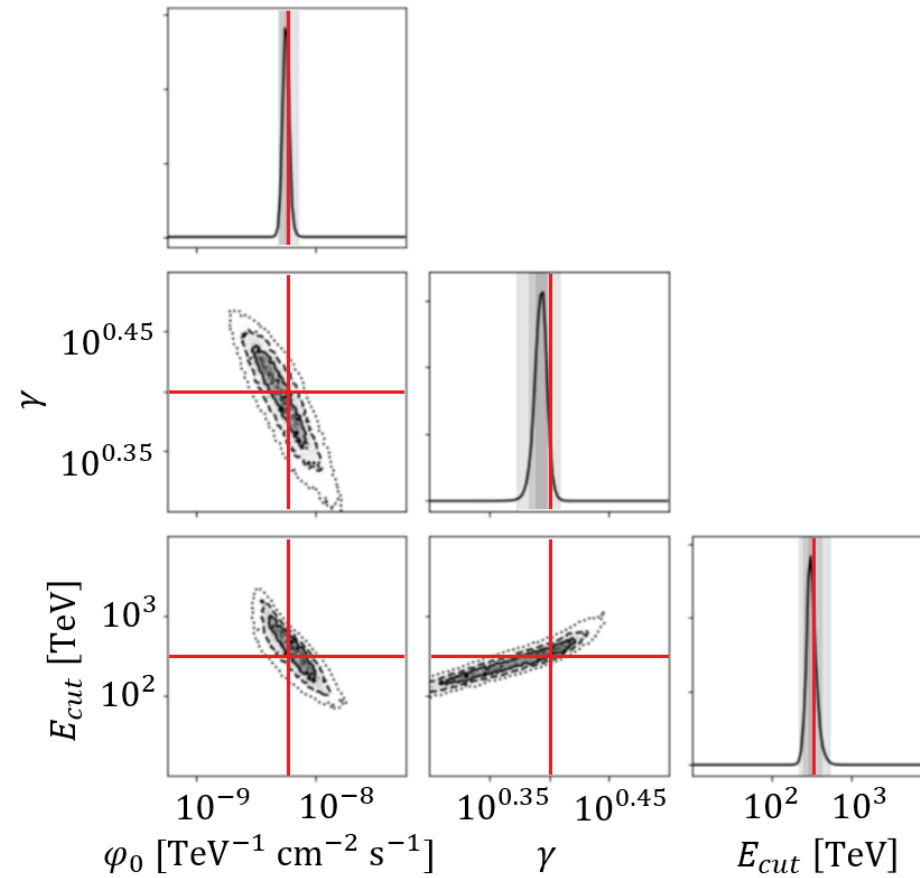


$$\varphi(E) = \varphi_0 \left(\frac{E}{E_0} \right)^\gamma e^{-E/E_{cut}} + \text{poisson noise}$$

- Parameters of interest:

- Amplitude φ_0
- Spectral index γ
- E_{cut}

Goal: Find the “probability-distributions” for the true parameter values



NRE is easy to implement without any knowledge of the underlying physics

NRE is easy to implement without any knowledge of the underlying physics

Workflow:

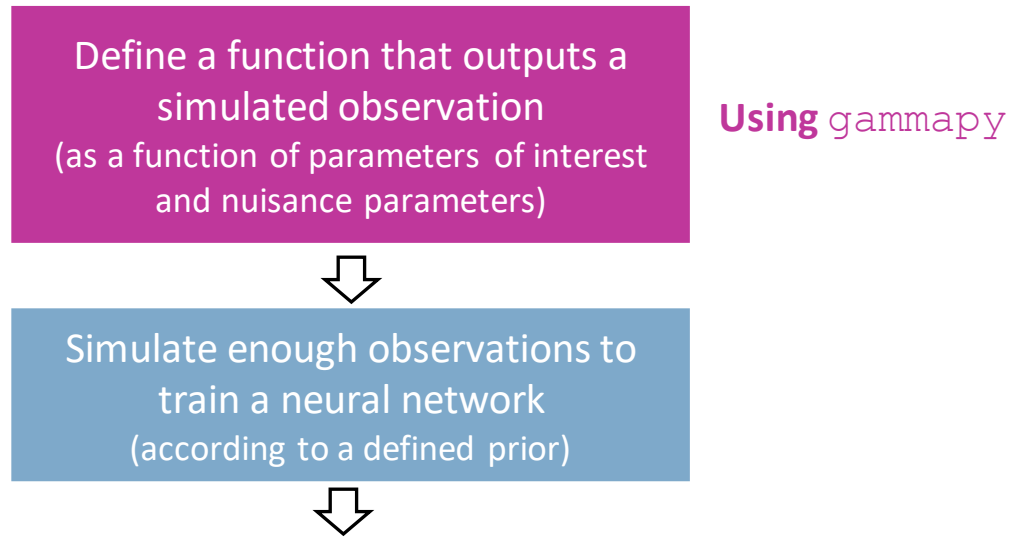
Define a function that outputs a
simulated observation
(as a function of parameters of interest
and nuisance parameters)

Using `gammapy`



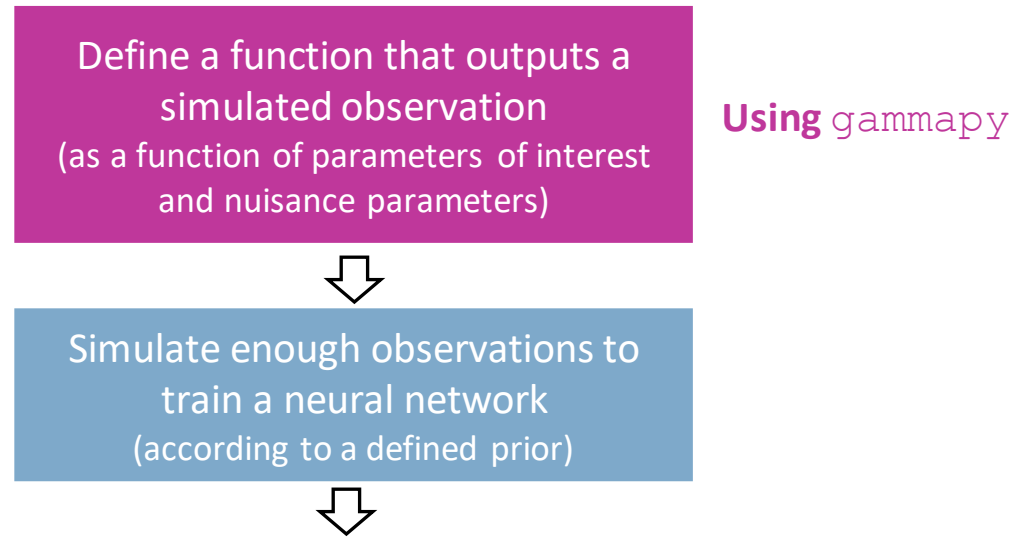
NRE is easy to implement without any knowledge of the underlying physics

Workflow:



NRE is easy to implement without any knowledge of the underlying physics

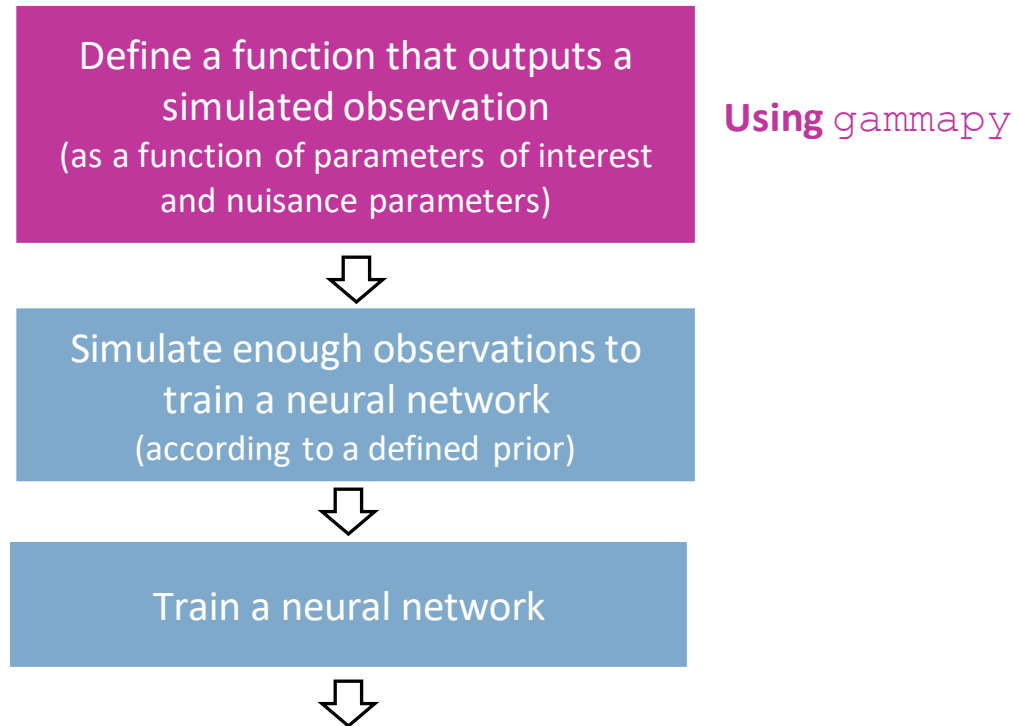
Workflow:



- The simulations implicitly contain the information on the relationship between parameters and observations

NRE is easy to implement without any knowledge of the underlying physics

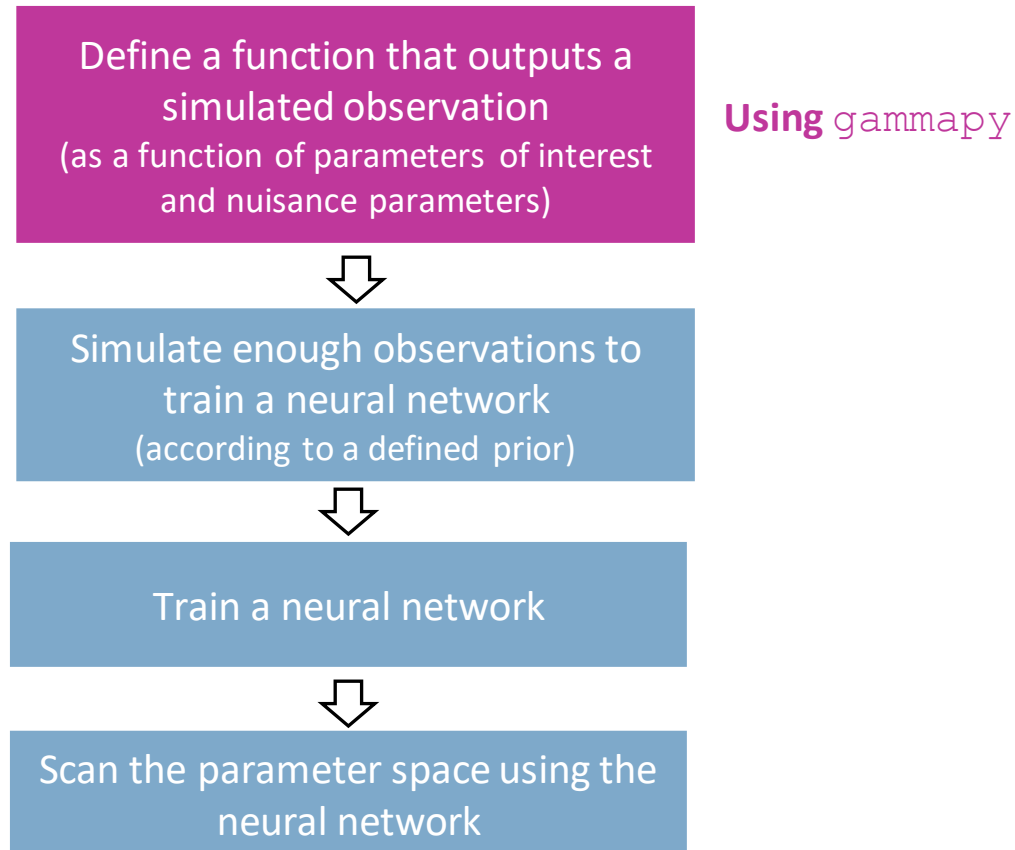
Workflow:



- The simulations implicitly contain the information on the relationship between parameters and observations

NRE is easy to implement without any knowledge of the underlying physics

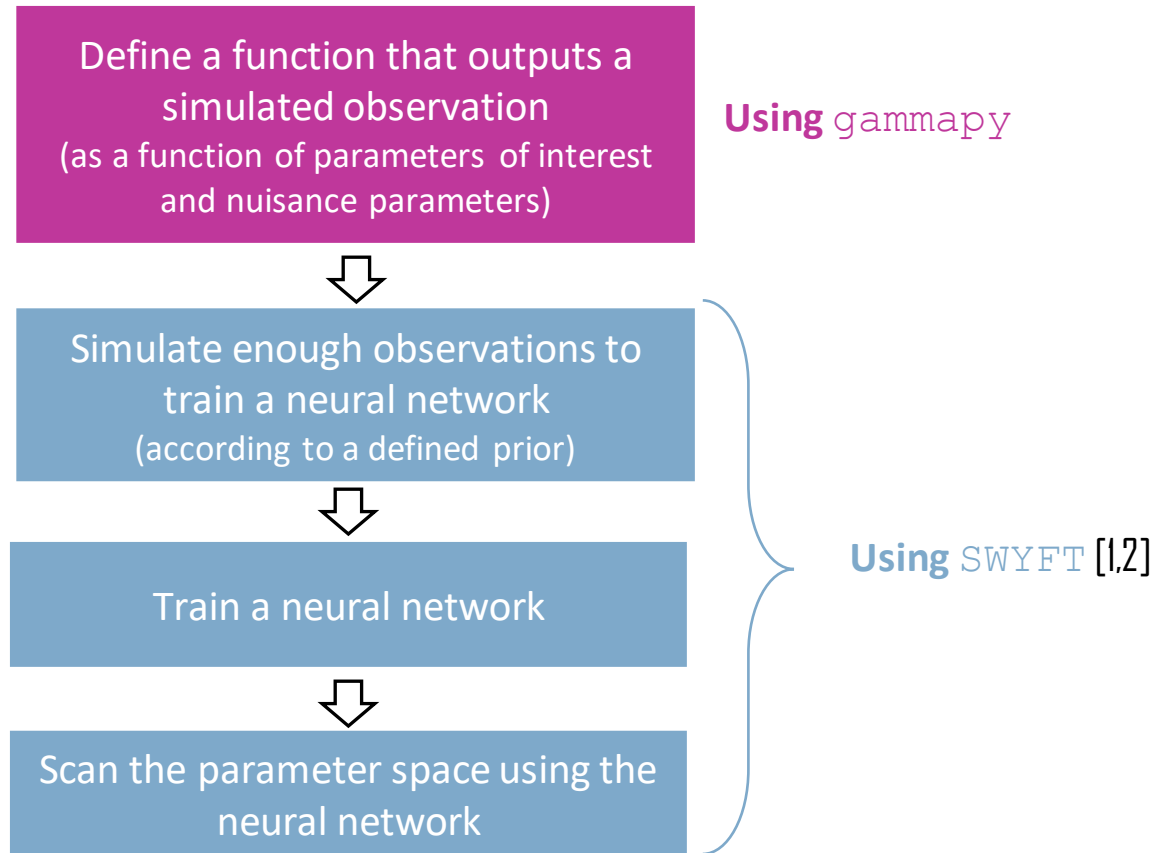
Workflow:



- The simulations implicitly contain the information on the relationship between parameters and observations

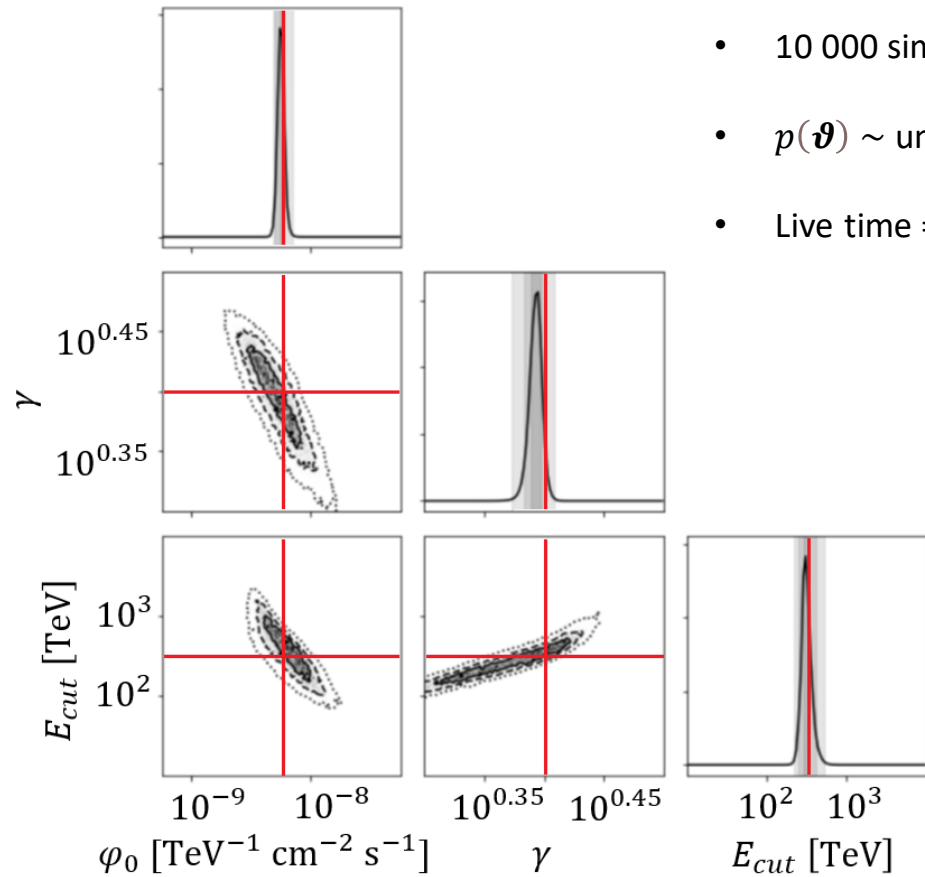
NRE is easy to implement without any knowledge of the underlying physics

Workflow:




- The simulations implicitly contain the information on the relationship between parameters and observations

Inference with NRE seems to be precise for the spectral fit




- 10 000 simulations used in training
- $p(\boldsymbol{\vartheta}) \sim \text{uniform on log scale}$
- Live time = 50 hr

Inference with NRE seems to be precise for the spectral fit




- 10 000 simulations used in training
- $p(\boldsymbol{\vartheta}) \sim \text{uniform on log scale}$
- Live time = 50 hr



The method seems particularly useful for ALP searches with gamma-telescopes


Expected result assuming ALPs with given mass m and coupling g (using `gammaALPs`):



The expected spectrum can be simulated using `gammapy` and `gammaALPs` (by M. Meyer)

The method seems particularly useful for ALP searches with gamma-telescopes

Expected result assuming ALPs with given mass m and coupling g (using `gammaALPs`):




- Parameters of interest:
 - Mass of ALPs, m
 - ALP-photon coupling, g

The expected spectrum can be simulated using `gammapy` and `gammaALPs` (by M. Meyer)

The method seems particularly useful for ALP searches with gamma-telescopes

Expected result assuming ALPs with given mass m and coupling g (using `gammaALPs`):




The expected spectrum can be simulated using `gammapy` and `gammaALPs` (by M. Meyer)

- Parameters of interest:
 - Mass of ALPs, m
 - ALP-photon coupling, g
- Nuisance parameters:
 - Amplitude
 - Spectral index
 - Cut-off energy
 - Magnetic field configuration
 - + 12 more related to configuration of NGC1275

The method seems particularly useful for ALP searches with gamma-telescopes


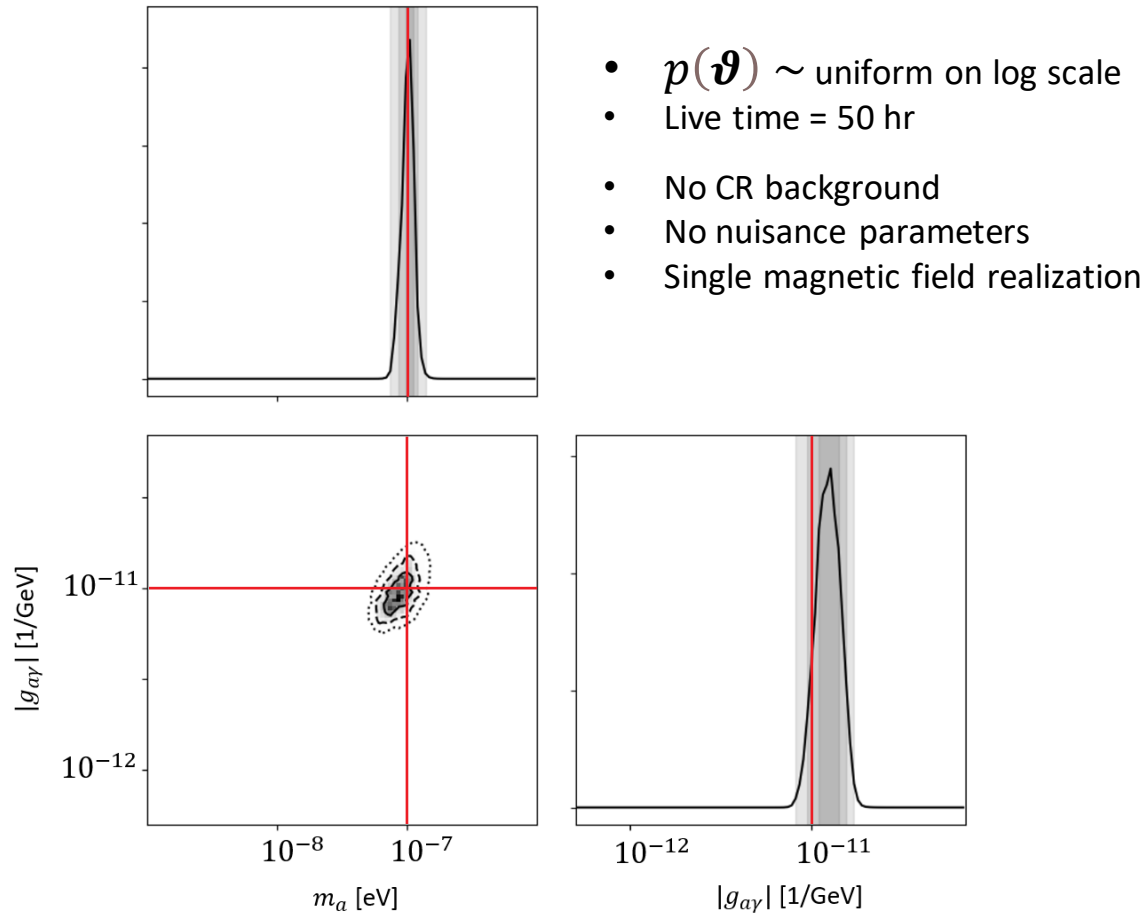
Expected result assuming ALPs with given mass m and coupling g (using `gammaALPs`):



The expected spectrum can be simulated using `gammapy` and `gammaALPs` (by M. Meyer)

- Parameters of interest:
 - Mass of ALPs, m
 - ALP-photon coupling, g
- Nuisance parameters:
 - Amplitude
 - Spectral index
 - Cut-off energy
 - Magnetic field configuration
 - + 12 more related to configuration of NGC1275
- For ALP searches, the frequentist test statistic does *not* obey Wilk's theorem!
 - Monte Carlo simulations are necessary to relate the TS to a significance of detection or exclusion **for each point in parameter space**.
 - Conventional computations are extremely expensive

Preliminary results indicate the method is suitable for ALP searches



Preliminary results indicate the method is suitable for ALP searches

