DIRECT LEARNING OF SYSTEMATICS-AWARE SUMMARY STATISTICS

Pablo de Castro (@pablodecm) and Tommaso Dorigo (@dorigo)

3rd August 2018 @ XIIIth QCHS Conference (Maynooth University - Ireland)

Statistical Methods for Physics Analysis in the XXI Century

AMVA4NewPhysics has received funding from European Union's Horizon 2020 Programme under Grant Agreement number 675440

OUTLINE OF THE TALK

1. Limitations of current ML techniques in HEP analyses
2. Why do we do signal vs background classification?
3. **INFERNO: Inference-Aware Neural Optimisation**
4. Results of proposed technique on simple example
5. Conclusions and Prospects

THE HEPML' THREE-BODY PROBLEM

The three main analysis components only share processed data, each step is carried out independently, without taking into account details about the remaining other two

The ultimate aim of analyses is **powerful statistical inference** (i.e. interval estimation or hyphotesis testing) based on the observed data

IS THIS THE BEST WE CAN DO?

MACHINE LEARNING WITHIN LHC ANALYSES

Most data analysis problems tacked with machine learning are cast as one the two canonical supervised learning tasks: classification or regression

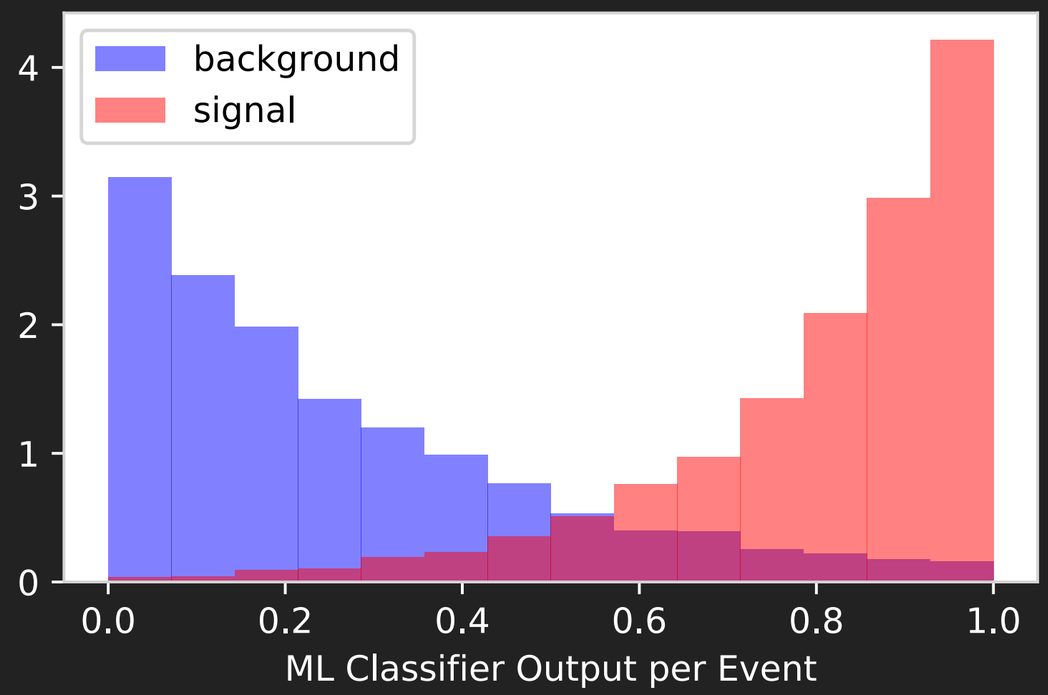

Event-by-Event Signal vs Background

- Higgs Kaggle challenge
- Every other LHC analysis

IS IT REALLY A CLASSIFICATION PROBLEM?

NO! IT IS A STATISTICAL INFERENCE PROBLEM

SIGNAL VS BACKGROUND? 🤔

Let us consider some i.i.d. observed data from a experiment $\mathcal{D} = \{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$

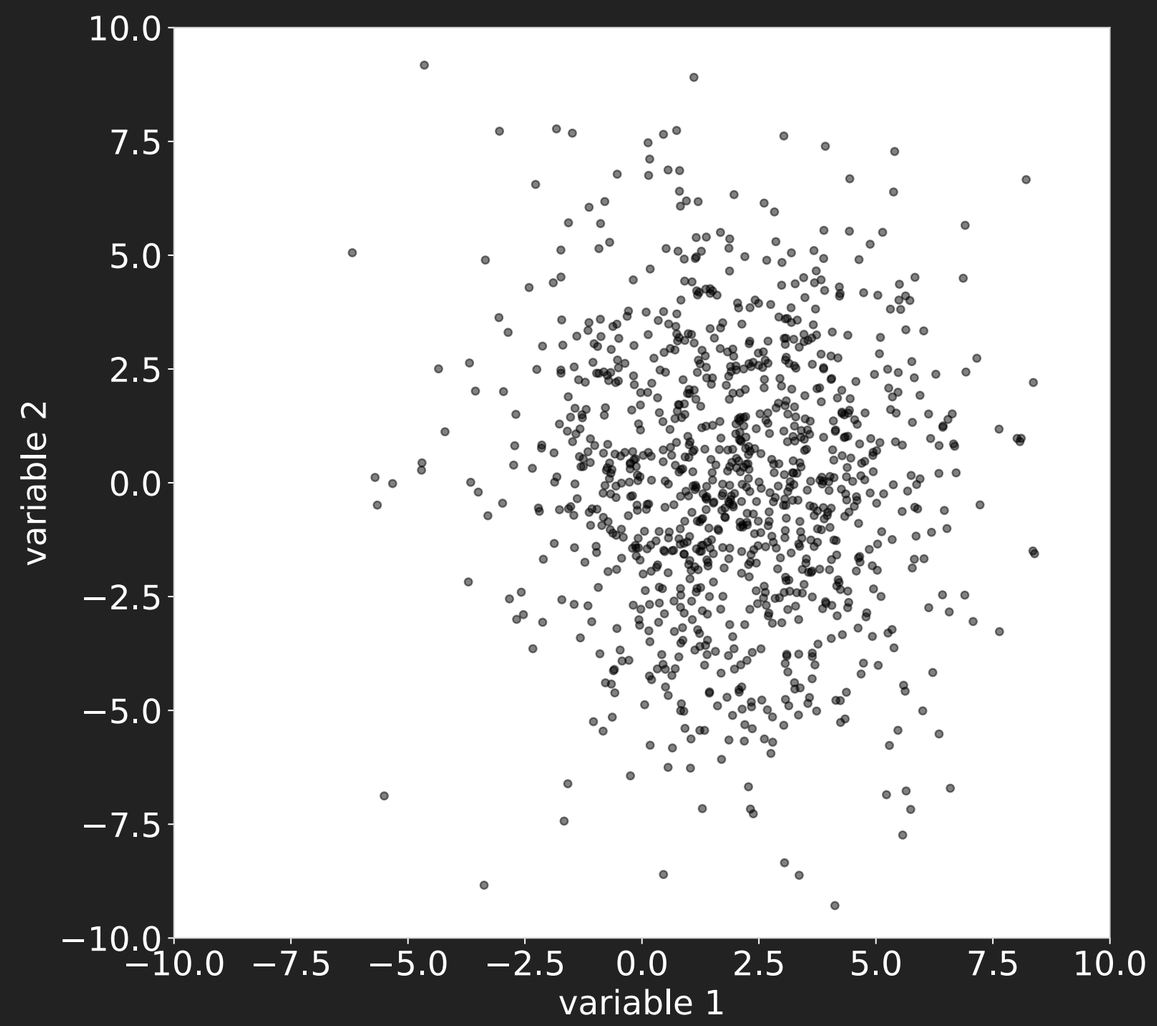

NO CLASS LABELS

Common problem → hyphotesis testing

Testing an alternate hyphotesis H_1 (e.g. theory that predicts new particle) against the null hyphotesis H_0 (e.g. Standard Model without that particle)

HOW TO APPROACH THIS PROBLEM?

SOLVED IN 1933: NEYMAN-PEARSON LEMMA

$$\Lambda(\mathcal{D}; H_0, H_1) = \prod_{\mathbf{x} \in \mathcal{D}} \frac{p(\mathbf{x}|H_0)}{p(\mathbf{x}|H_1)}$$

The **likelihood ratio** is the most powerful test (lowest Type II error) at a fixed significance level α between two simple hypotheses (i.e. that completely specify the distribution, **do not have additional parameters**)

MAIN COMPLICATION → $\mathbf{p(x|H_0)}$ ARE $\mathbf{p(x|H_1)}$ ARE INTRACTABLE

$\mathbf{p(x|model)}$ IS NOT KNOWN AT LHC EXPERIMENTS

Samples under different hypotheses can be simulated via complex physics-based MC programs but $p(\mathbf{x})$ cannot be directly evaluated → **LIKELIHOOD-FREE INFERENCE**

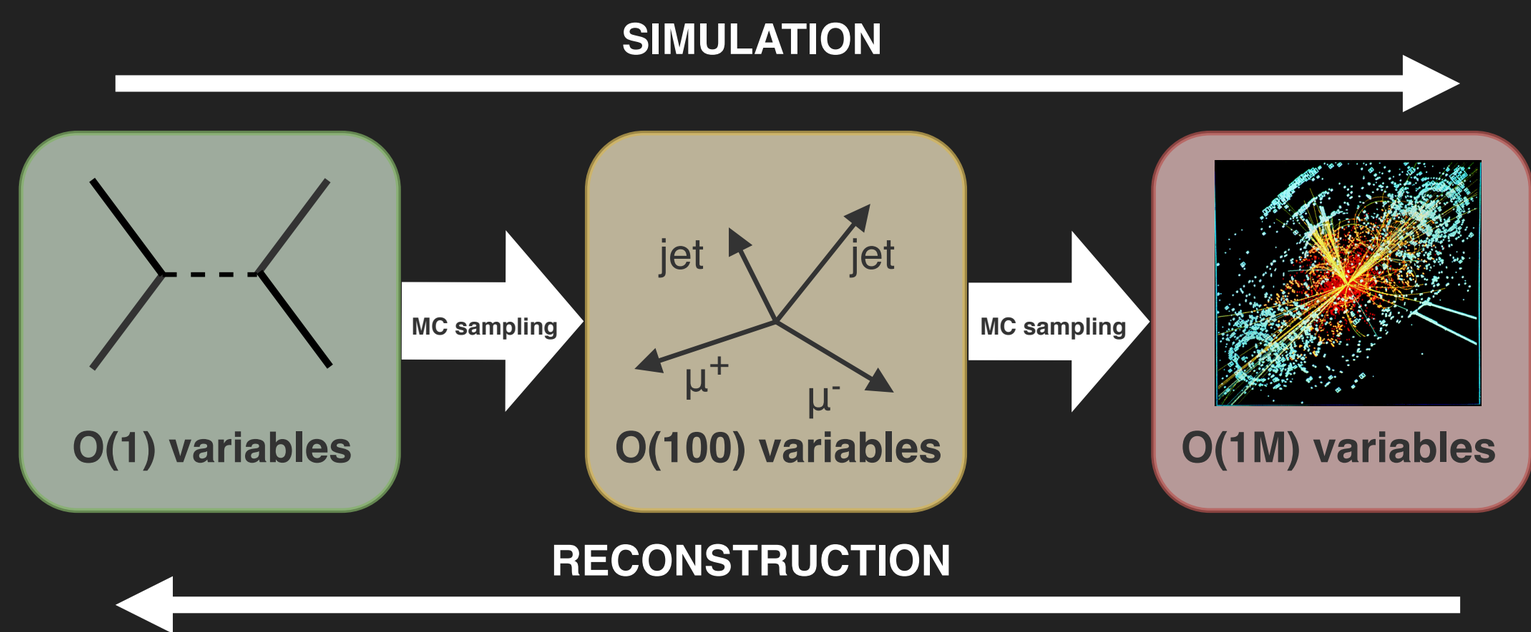

good approximations of $p(\mathbf{x})$ are unachievable due to curse of dimensionality

DIMENSIONALITY REDUCTION $\mathbf{R^n} \rightarrow \mathbf{R^{O(1)}}$

KEEPING AS MUCH USEFUL INFORMATION FOR INFERENCE AS POSSIBLE

THE MIXTURE RATIO TRICK

Alternate hyphotesis H_1 is commonly a mixture model of a background component $p_\mathrm{b}(\mathbf{x}) = p(\mathbf{x}|H_0)$ and a small fraction μ of signal component $p_\mathrm{s}(\mathbf{x})$

$$\Lambda^{-1} \sim \frac{p(\mathbf{x}|H_1)}{p(\mathbf{x}|H_0)} = \frac{(1-\mu)\cdot p_\mathrm{b}(\mathbf{x}) + \mu \cdot p_\mathrm{s}(\mathbf{x})}{p_\mathrm{b}(\mathbf{x})}$$

$$\Lambda^{-1} \sim 1 - \mu \cdot \left(\frac{p_\mathrm{s}(\mathbf{x})}{p_\mathrm{b}(\mathbf{x})} - 1 \right)$$

Therefore Λ is monotonically linearly decreasing with $p_\mathrm{s}(\mathbf{x})/p_\mathrm{b}(\mathbf{x})$

NEW TASK: APPROXIMATING $\mathbf{p_s(x)/p_b(x)}$

However, the generating distributions of background $p_{\mathrm{b}}(\mathbf{x})$ or signal $p_{\mathrm{s}}(\mathbf{x})$ are still not known

Only forward-simulated samples $\mathcal{S} = \{\mathbf{x}_0^{\mathrm{s}}, \ldots, \mathbf{x}_S^{\mathrm{s}}\}$ & $\mathcal{B} = \{\mathbf{x}_0^{\mathrm{b}}, \ldots, \mathbf{x}_B^{\mathrm{b}}\}$ are available

This synthetic dataset (2D Gaussian mixture) will be used extensively in this talk to exemplify different techniques

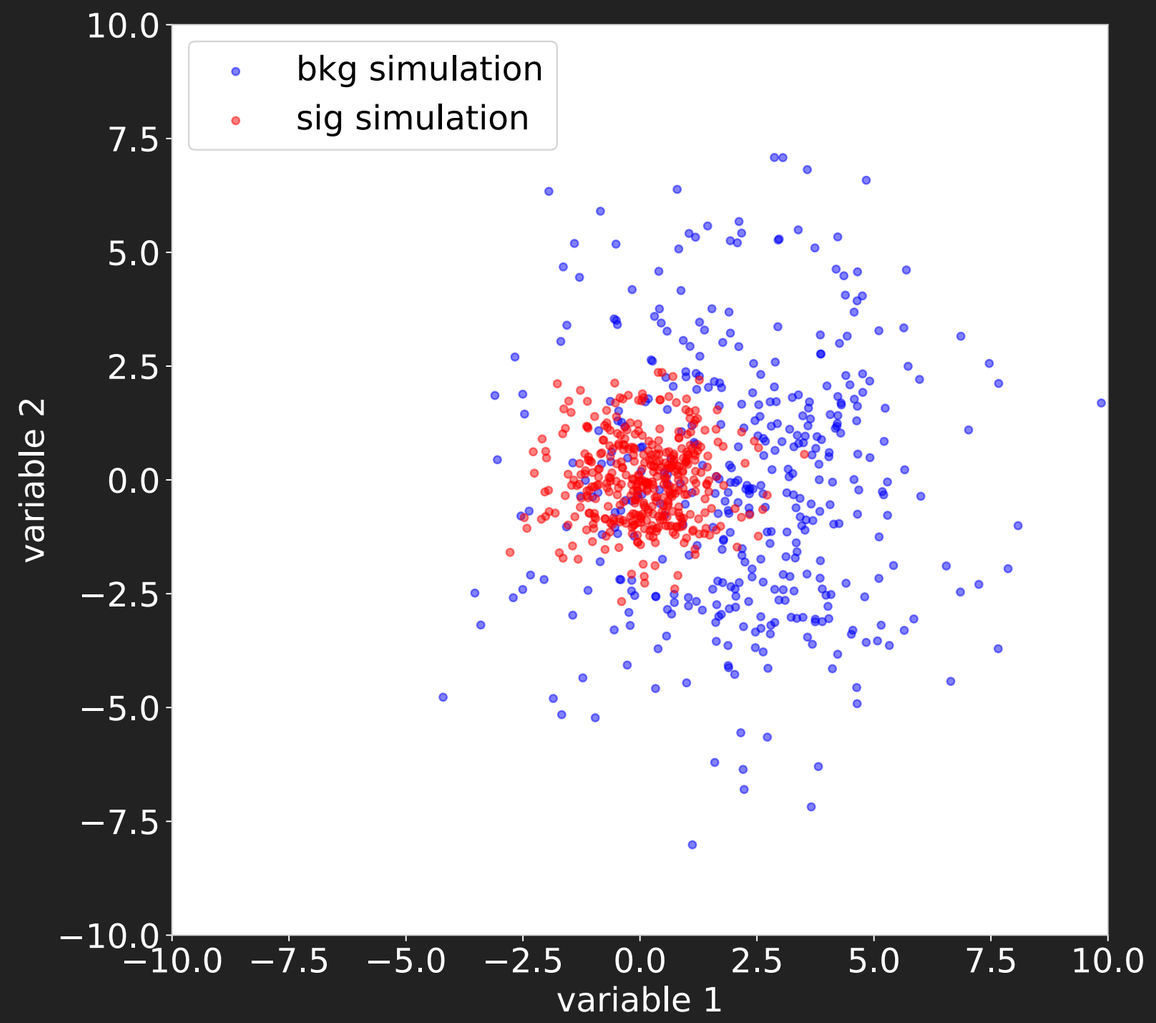

AMENABLE BY ML CLASSIFICATION → SIGNAL VS BACKGROUND

CLASSIFICATION AS SURROGATE TO OBTAIN SUMMARY STATISTICS

Most machine learning classification algorithms approximate the likelihood ratio $p_{\mathrm{s}}(\mathbf{x})/p_{\mathrm{b}}(\mathbf{x})$
(e.g. a deep neural network minimizing cross entropy loss $-\sum_i k_i \log y_i$)

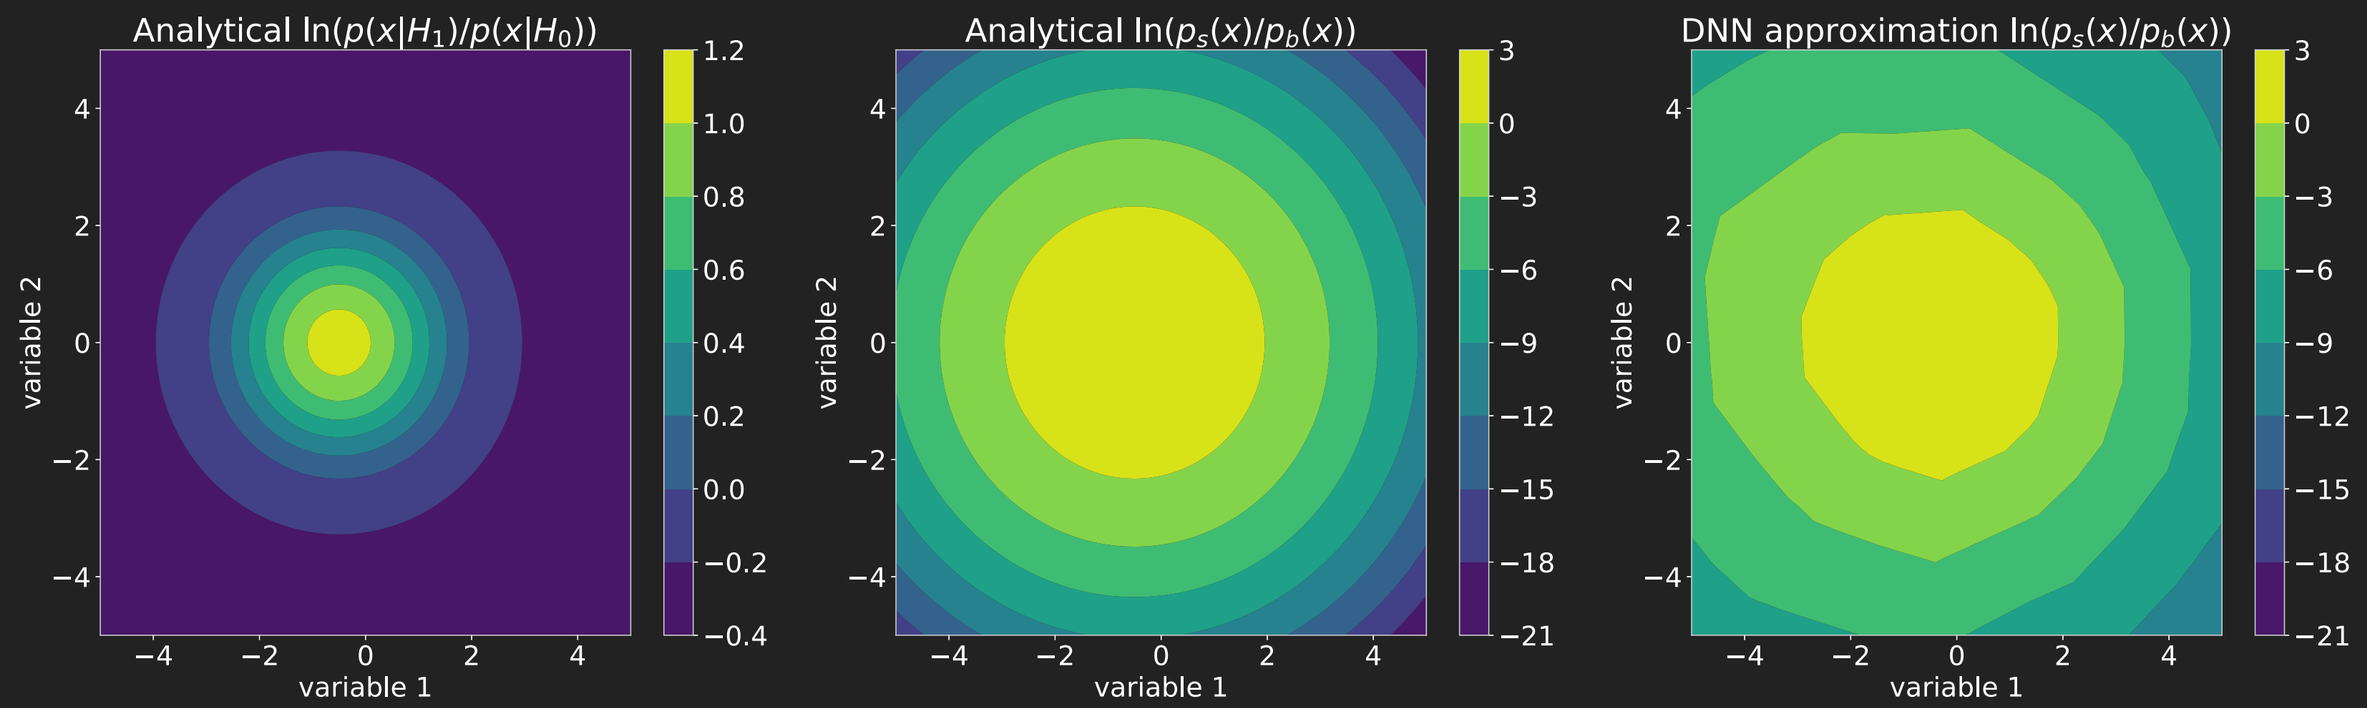

A direct exploitation of this for inference is carried out in "Approximating Likelihood Ratios with Calibrated Discriminative Classifiers" by K. Cranmer et al.
Work further extendend in arXiv:1805.12244 (and cited articles therein) by J. Brehmer et al to use also the joint score.

CLASSIFIER-BASED INFERENCE

A trained ML classifer $d(\mathbf{x})$ is an uncalibrated approximation of $p_\mathrm{s}(\mathbf{x})/p_\mathrm{b}(\mathbf{x})$

How can it be used for statistical inference from observed data $\mathcal{D}$?

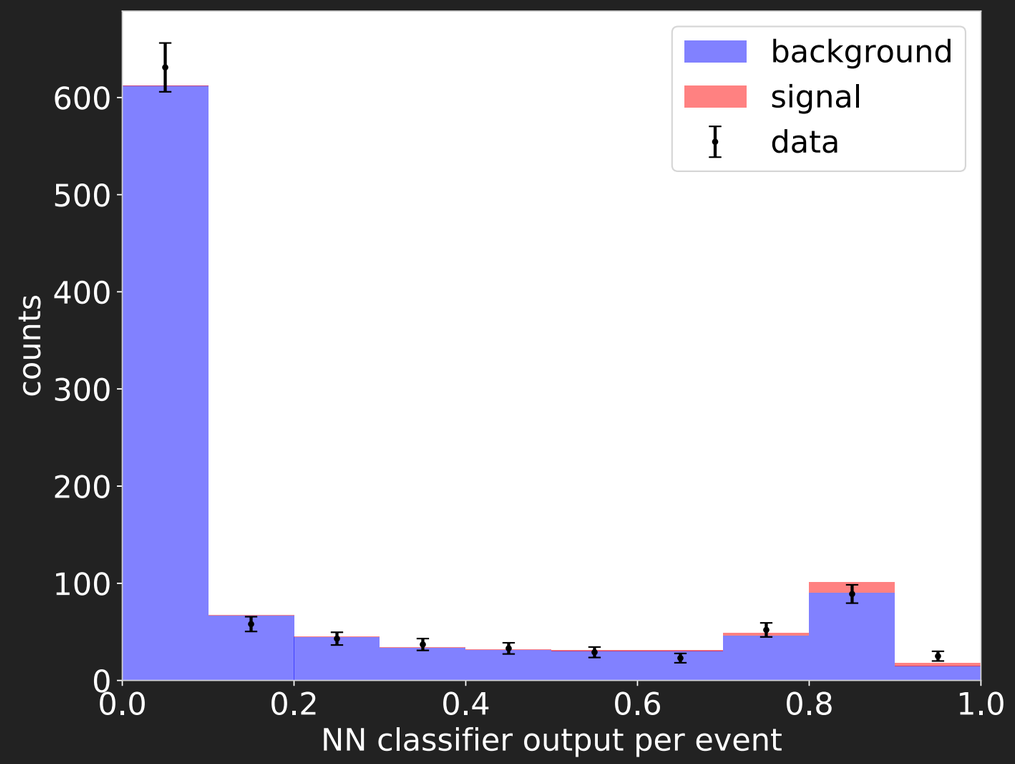

1-D → **cut or histogram to build a Poisson counts non-parametric likelihood**

$$\mathcal{L}(\mu) = \prod_{\mathrm{i}\in\mathrm{bins}} \mathrm{Pois}(n_\mathrm{i}|\mu \cdot s_\mathrm{i} + b_\mathrm{i})$$

which can be used for further inference, such as measuring μ given observed $\mathcal{D}$

Choice of binning can get quite tricky. For cut and count "Consistent optimization of AMS by logistic loss minimization" by W. Kotlowski is quite insteresting (theoretical proof of classification consistency as surrogate).

MODELLING UNCERTAINTIES DEGRADE INFERENCE

Simulations are imperfect, mainly due to the limited information of the system being modelled

Lack of knowledge for inference accounted by additional unknown parameters (nuisance parameters ν)

Causes a degradation of classifier-based inference, leading to larger measurement uncertainties

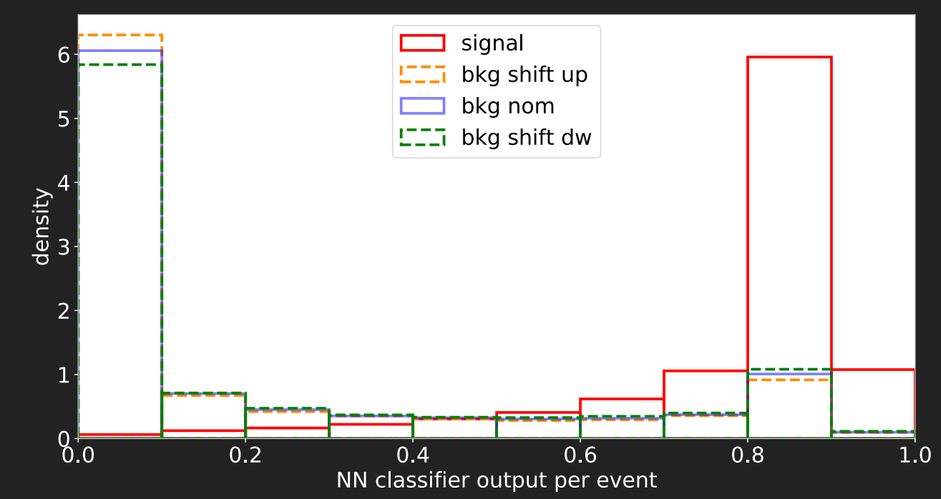

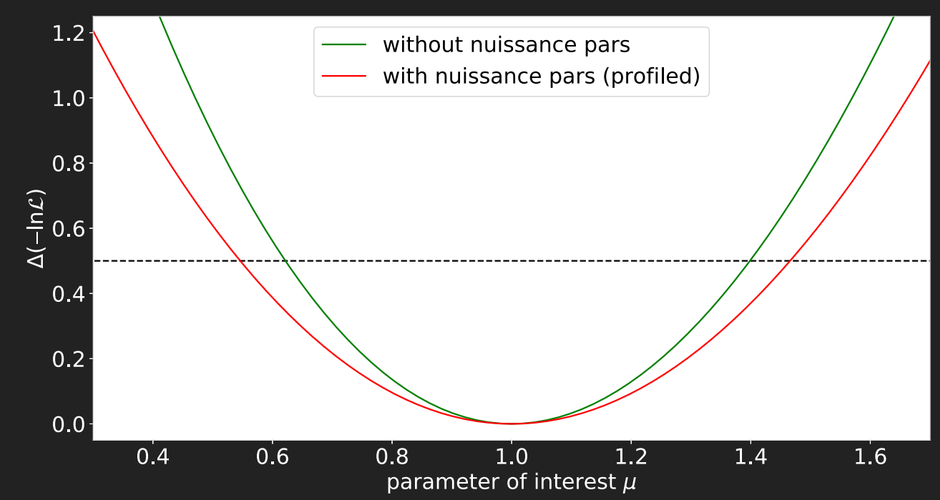

UPPER LIMIT OF ML USEFULNESS IN LHC ANALYSES

Classifiers can be made pivotal as described in "Learning to Pivot" by G. Louppe et al. A review/benchmarks on how to deal with systematics when using machine learning can be found in Adversarial learning to eliminate systematic errors: a case study in High Energy Physics by Victor Estrade et al NIPS2017.

CAN WE PUT IT ALL TOGETHER?

Embed some of the knowledge about modelling and statistical inference such as systematic uncertainties in the dimensionality-reduction step

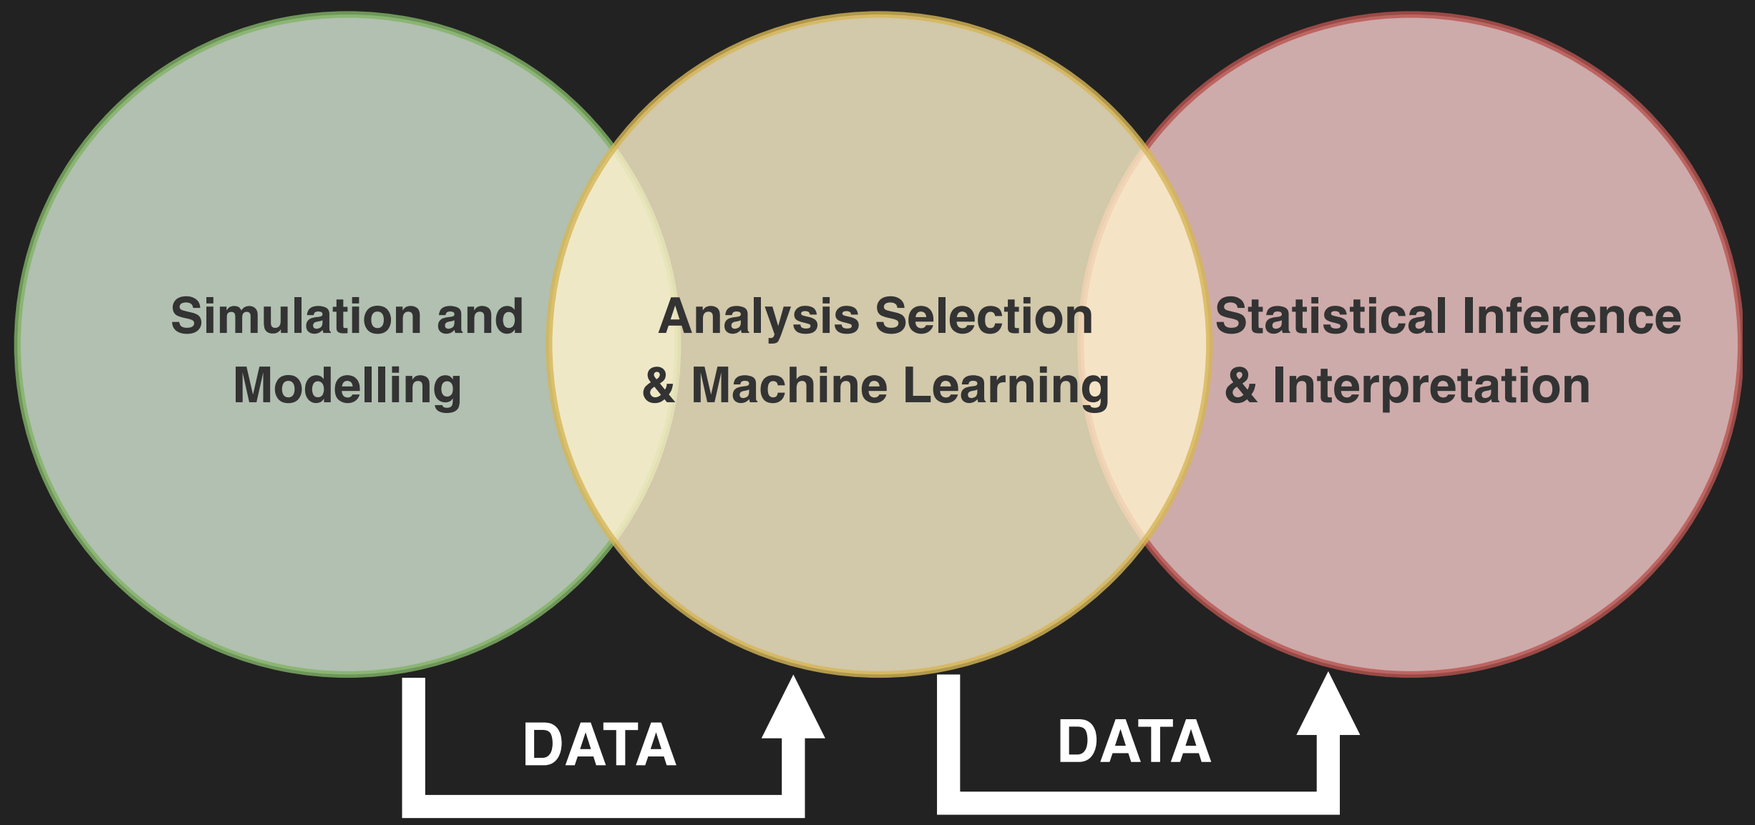

GLUE → AUTODIFF GRAPH FRAMEWORKS

CHEAP (BUT EXACT) DERIVATIVES ARE A BIG DEAL!

Autodiff Frameworks (e.g. TensorFlow)

- Highly parallel (CPU/GPU) or distributed
- Support higher order-gradients
  - Hessian very useful for inference
- **Statistical libraries available**
  - TensorFlow Distributions
  - TensorFlow Probability
  - Edward → probabilistic modelling
- **DNNs and Stochastic optimization**

```
import tensorflow as tf
ds = tf.contrib.distributions

bkg = ds.MultivariateNormalFullCovariance(loc=[2.,
        covariance_matrix=[[5.,0.],[0.,9.]],
        name="bkg")
sig = ds.MultivariateNormalDiag(loc=[0., 0.],
        scale_diag=[1., 1.],
        name="sig")

mu = tf.placeholder(shape=(),
        dtype=tf.float32, name="mu")
mix = ds.Mixture(cat=ds.Categorical(probs=[1.-mu, m
        components = [bkg, sig], name="mix")
```

Code specifying 2D Gaussian mixture synthetic model used all along the talk

Autodiff graph frameworks are also useful to speed up common inference tasks such as likelihood fits, toy generation or limit setting. Non machine learning uses within HEP include TensorFlowAnalysis or pyhf.

END-TO-END DIFFERENTIABILITY FOR LHC ANALYSES

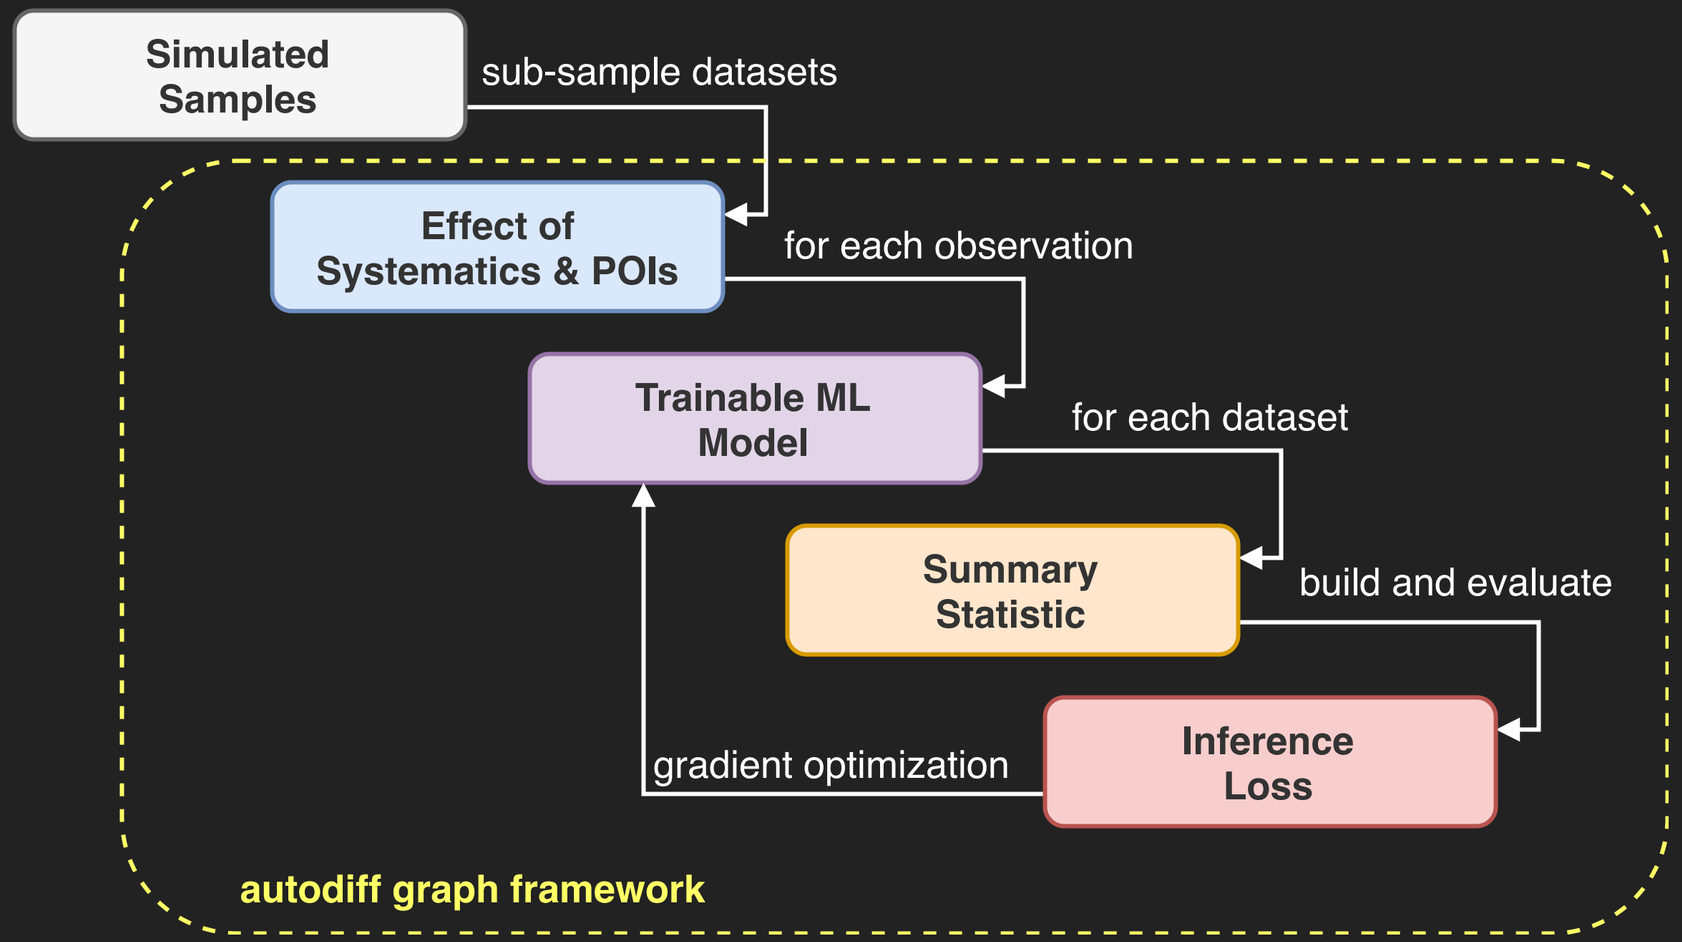

Within this general framework, several approaches are possible, focus here is

DIRECT LEARNING OF SYSTEMATICS-AWARE SUMMARY STATISTICS

MODELLING THE EFFECT OF SYSTEMATIC UNCERTAINTIES

Differentiable approximation of the effect of parameters of interest θ and nuisance parameters ν over the training datasets

A (potentially non-linear) function that depends on the details of problem and transforms event features and/or weights

- jet/tau/muon energy scale
- PDF/QCD uncertainties

Could also depend on simulation/latent variables, such the event category (S/B)

Simple example: shift for background observations (i.e. $\mathbf{x}' = \mathbf{x} + \nu \cdot \mathbf{s}$)

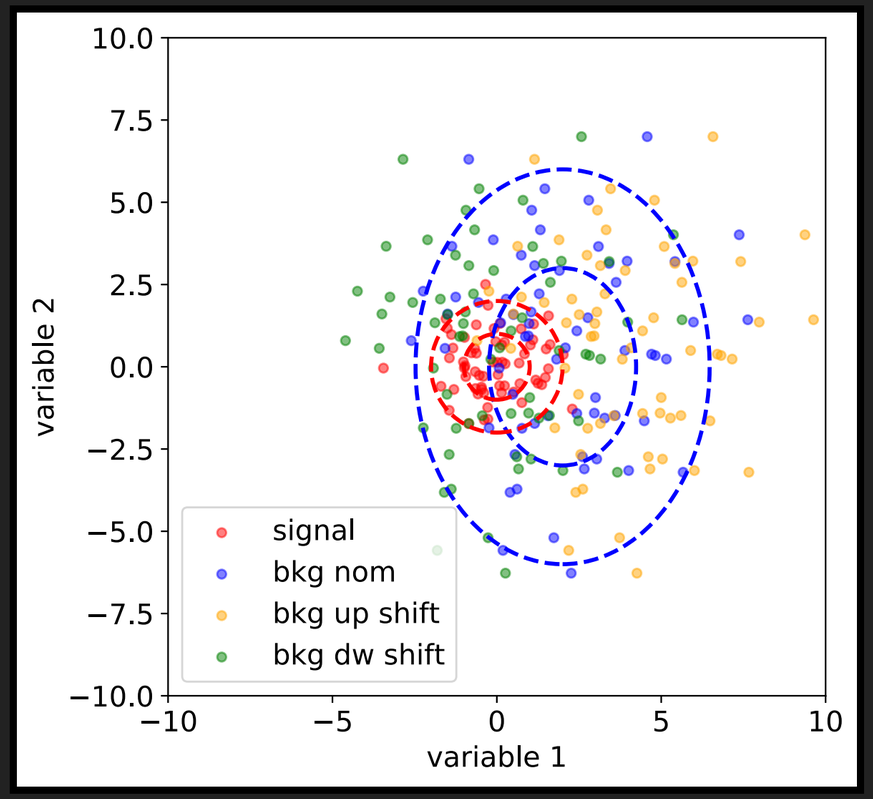

TRAINABLE PARAMETRIZED MODEL

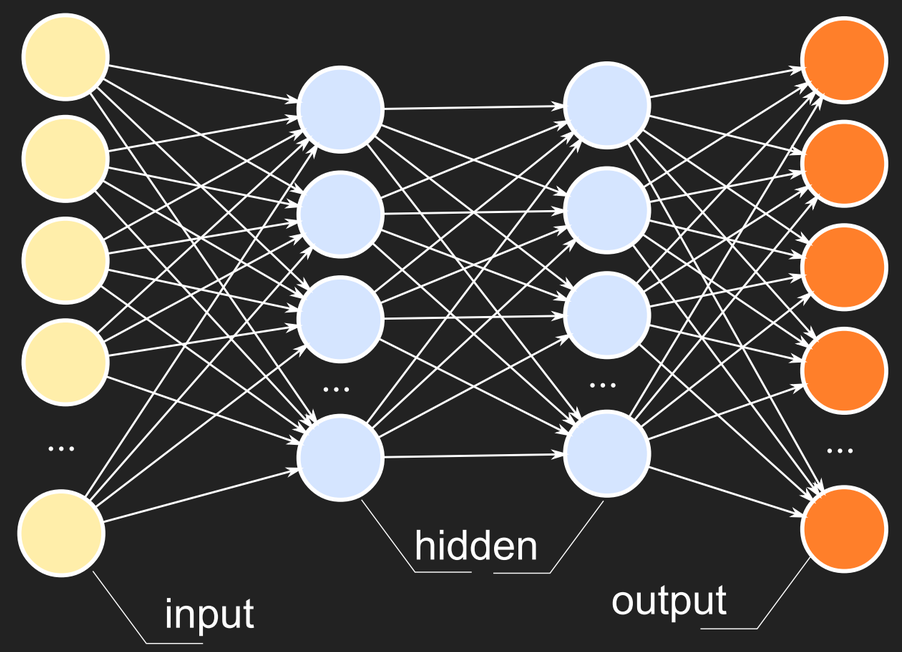

Any (deep) neural network will do

Could in principle re-use the same tweaks, techniques and architectures as for standard supervised deep learning

A small two-hidden layer MLP (10 units each, ReLU activation, glorot_normal initd) has been used for examples here.

NEURAL NETWORK OUTPUT → SUMMARY STATISTIC

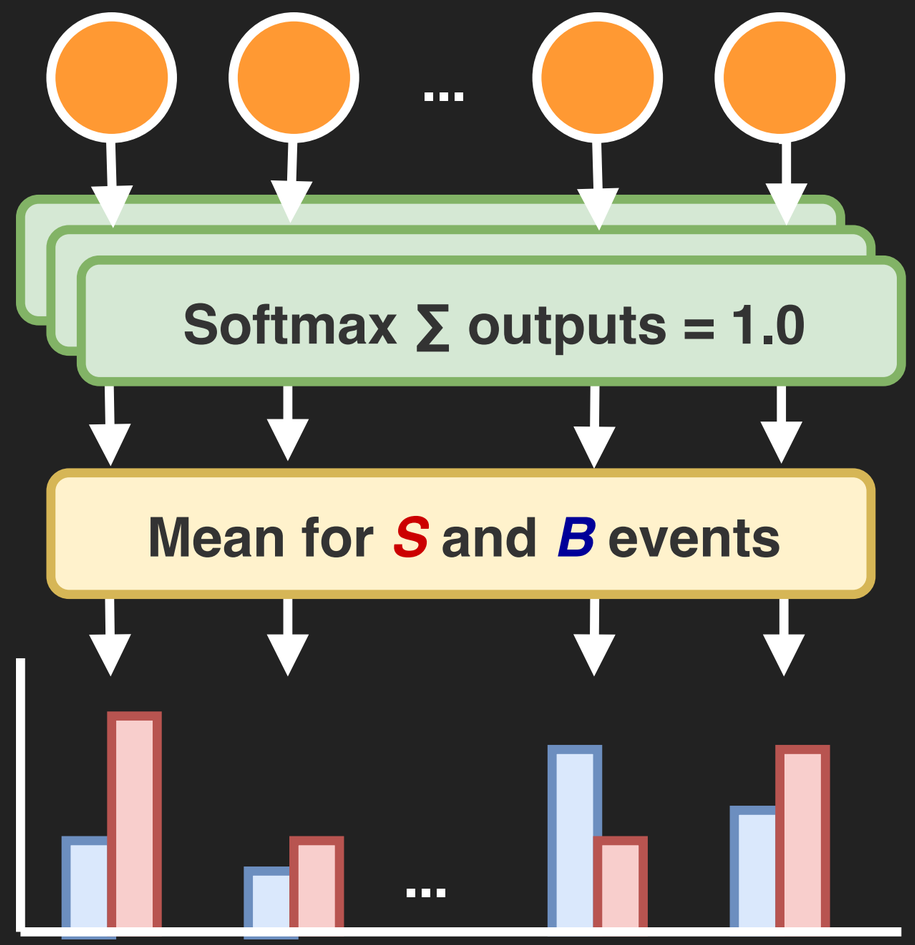

We can approximate a histogram-like summary statistic from the NN output applying softmax for each event and summing over each dataset

$$\mathcal{L}(\theta, \nu) = \prod_{\mathrm{i} \in \mathrm{bins}} \mathrm{Pois}(n_\mathrm{i} | \alpha_s \cdot s_\mathrm{i} + \alpha_b \cdot b_\mathrm{i})$$

The likelihood depends both on the neural network parameters and the statistical model parameters

INFERENCE-MOTIVATED LOSS FUNCTION

If we expand the negative log-likelihood around minimum (e.g. Asimov $n_\mathrm{i} = \alpha_s \cdot s_\mathrm{i} + \alpha_b \cdot b_\mathrm{i}$), due to Cramér-Rao bound:

$$\text{covariance} \geq \mathbf{H}^{-1}(-\ln \mathcal{L})$$

which can be computed via autodiff. Can use as loss function directly the variance bound on the parameters of interest

$$\text{loss} \approx \mathrm{Var}(\mu) \quad \text{(expected)}$$

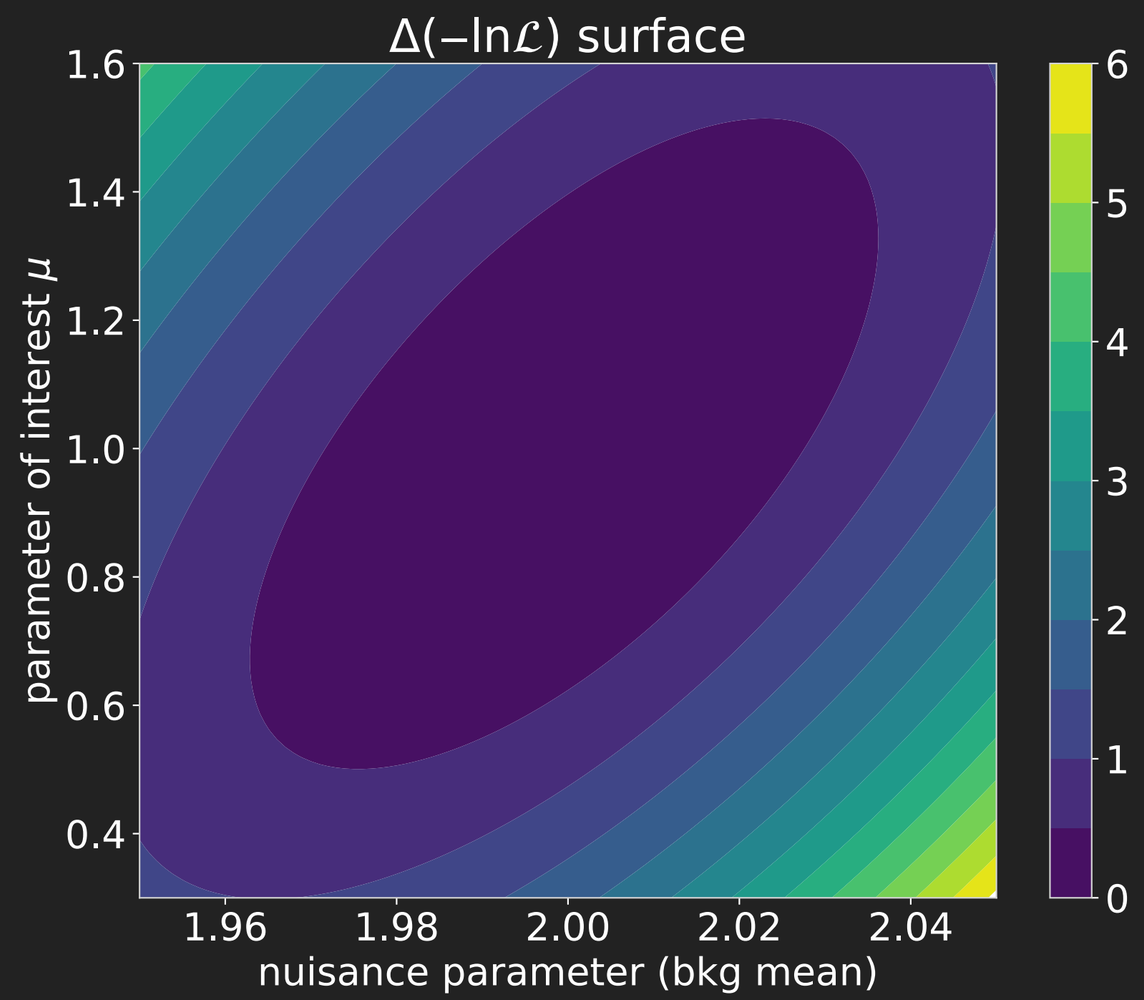

INFERNO: INFERENCE-AWARE NEURAL OPTIMISATION

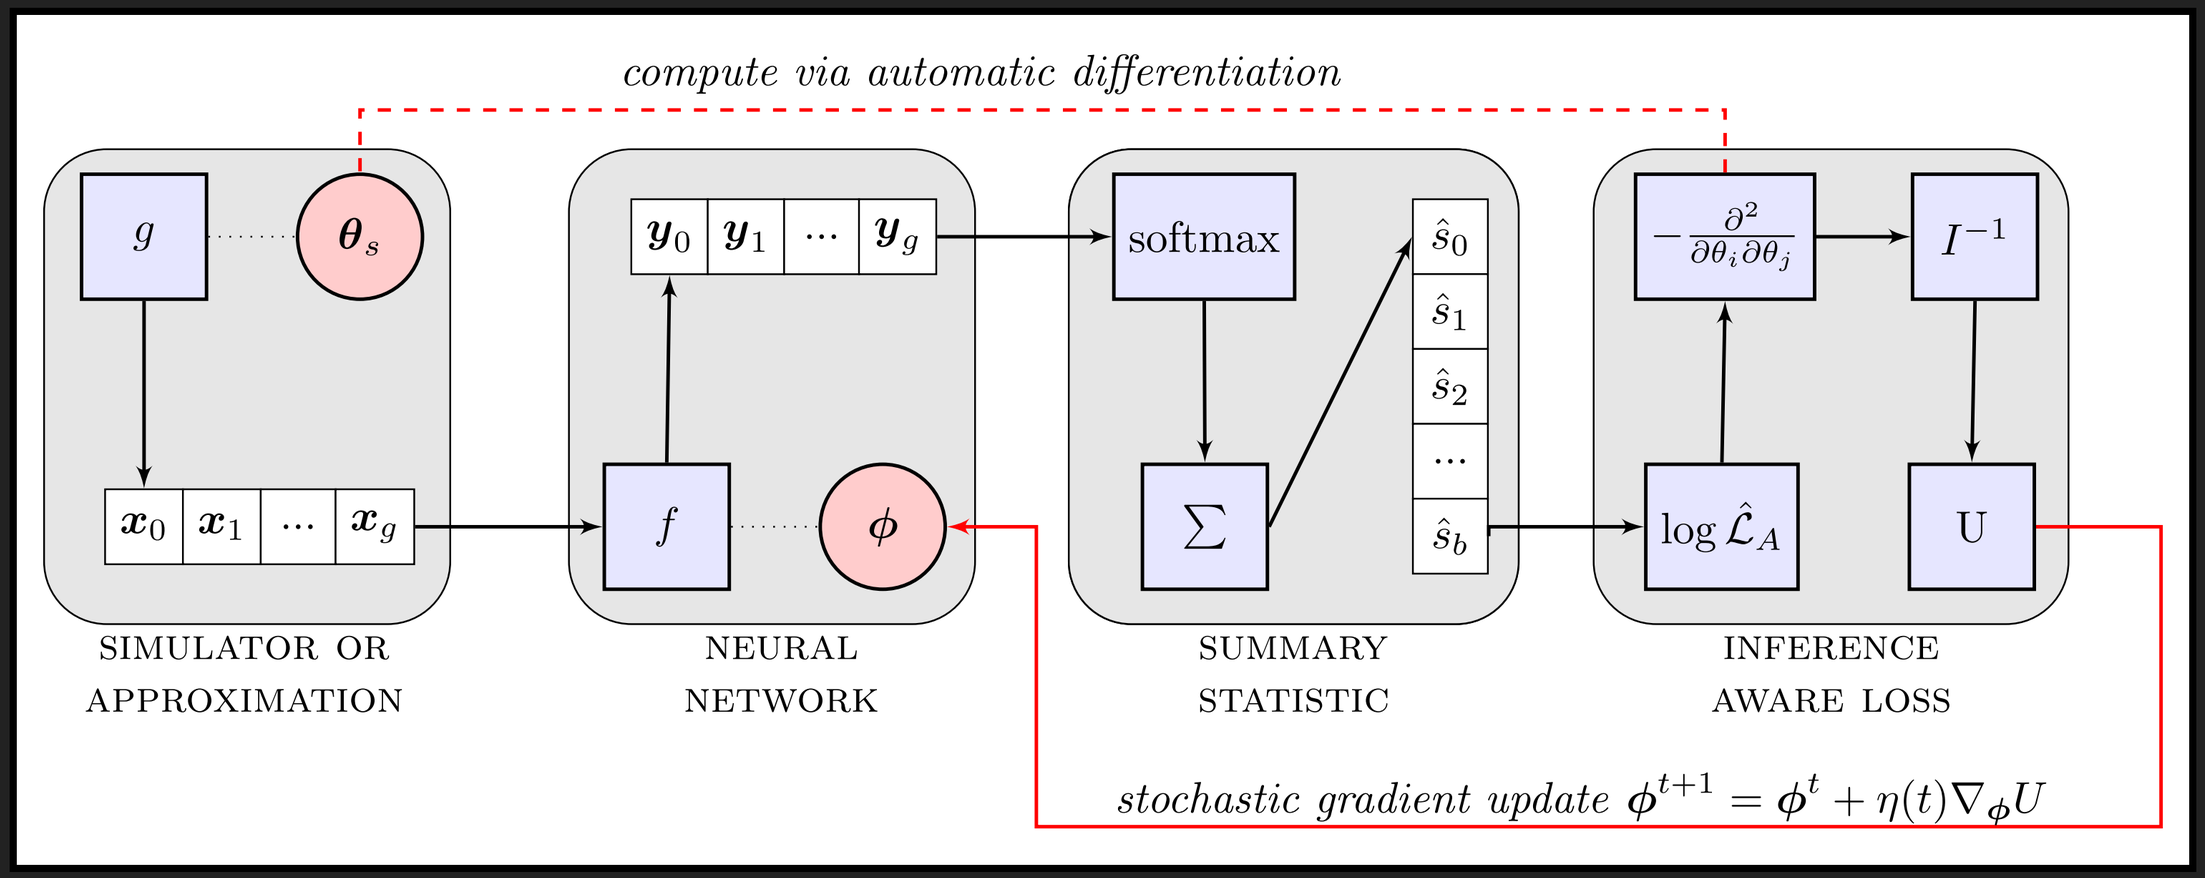

check arxiv.org/abs/1806.04743 for a detailed mathematical description

SYNTHETIC EXAMPLE IMPLEMENTATION

Applied on 2D Gaussian two-component mixture toy dataset, with unknown background mean in one of the coordinates → **one nuisance parameter**

Loss is non-decomposable, because it is dataset-based instead of event-based

Number of simulated events sub-sampled for each mini-batch is especially relevant to regulate gradient variance

Seems to converge independently on the initialization. Learning rate is a critical hyper-parameter though.

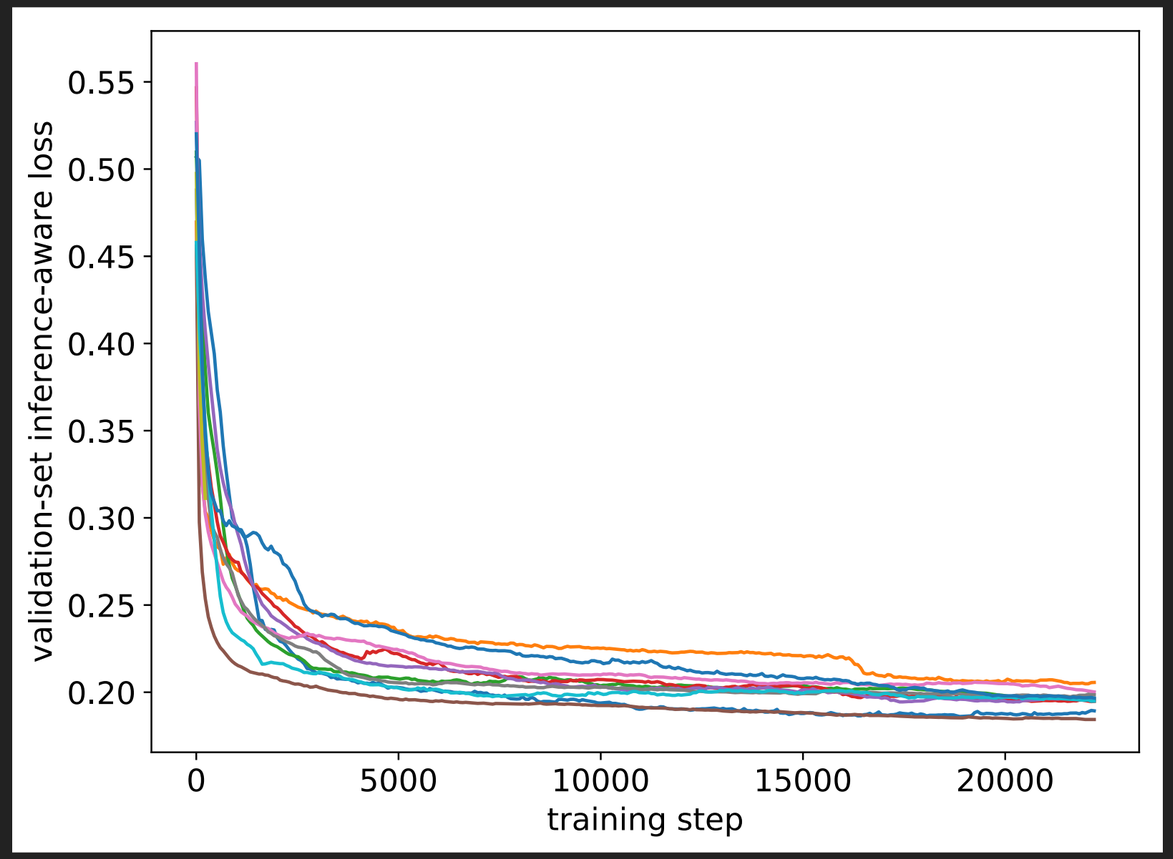

COMPARISON WITH CLASSIFICATION-BASED APPROACH

Expected signal strength uncertainties computed using the validation set (average of 10 random initialisations):

- cross-entropy: 0.444 ± 0.003
- **inference-aware: $\mathbf{0.437 \pm 0.008}$**

Early results are encouraging but additional study of the technique needed.

Now working on higher-dimensional and more nuisance parameters problems.

BETTER/EQUAL THAN CLASSIFICATION

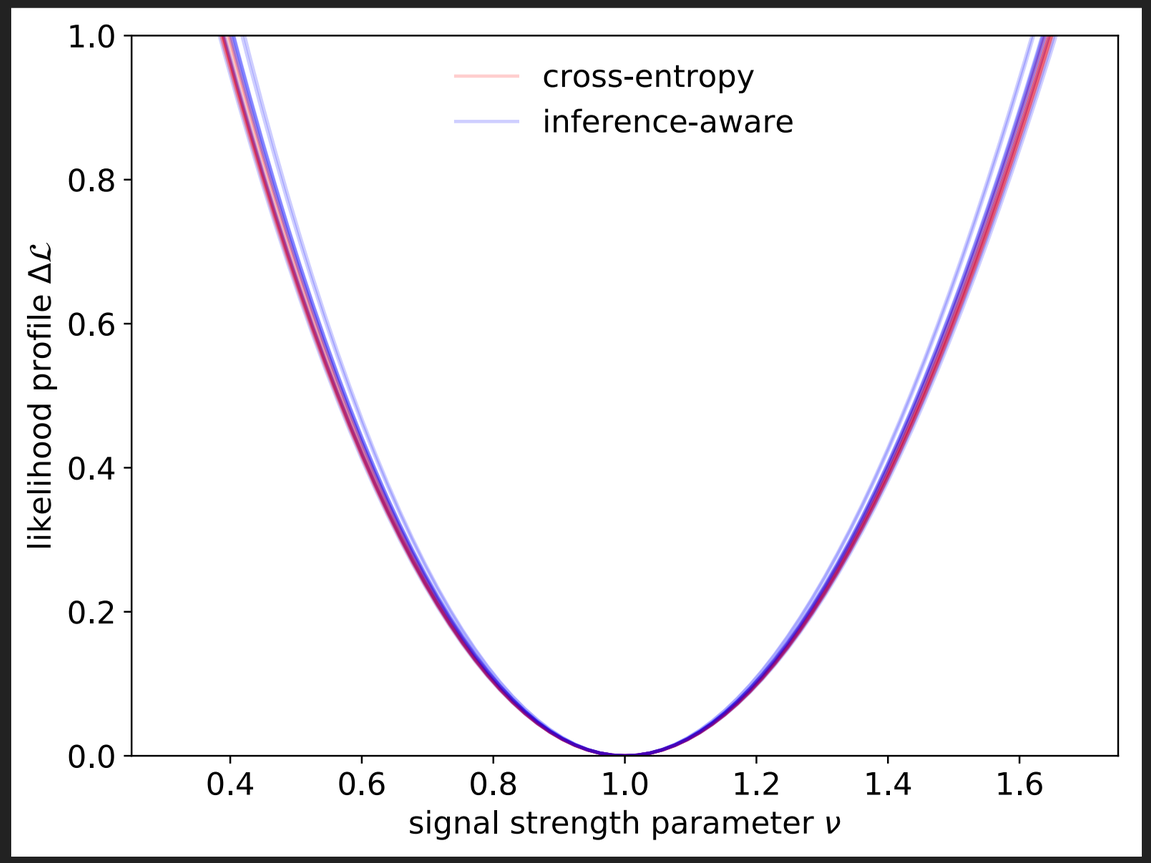

MORE DETAILS ON ARXIV PREPRINT

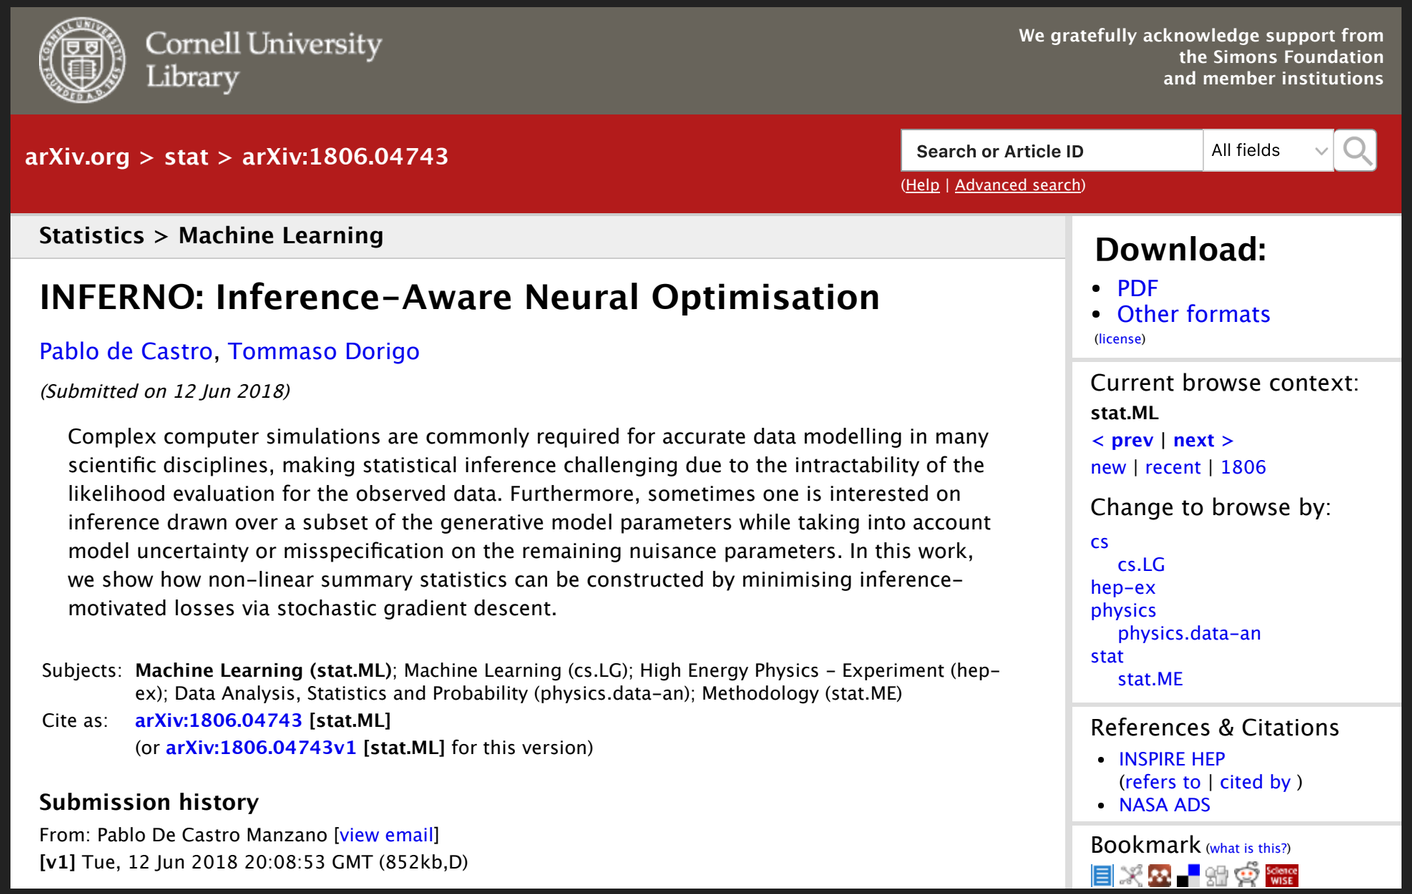

Cornell University Library

We gratefully acknowledge support from the Simons Foundation and member institutions

arXiv.org > stat > arXiv:1806.04743

Search or Article ID | All fields

(Help | Advanced search)

Statistics > Machine Learning

INFERNO: Inference-Aware Neural Optimisation

Pablo de Castro, Tommaso Dorigo

(Submitted on 12 Jun 2018)

Complex computer simulations are commonly required for accurate data modelling in many scientific disciplines, making statistical inference challenging due to the intractability of the likelihood evaluation for the observed data. Furthermore, sometimes one is interested on inference drawn over a subset of the generative model parameters while taking into account model uncertainty or misspecification on the remaining nuisance parameters. In this work, we show how non-linear summary statistics can be constructed by minimising inference-motivated losses via stochastic gradient descent.

Subjects: **Machine Learning (stat.ML)**; Machine Learning (cs.LG); High Energy Physics - Experiment (hep-ex); Data Analysis, Statistics and Probability (physics.data-an); Methodology (stat.ME)

Cite as: **arXiv:1806.04743 [stat.ML]**
(or **arXiv:1806.04743v1 [stat.ML]** for this version)

Submission history

From: Pablo De Castro Manzano [view email]
[v1] Tue, 12 Jun 2018 20:08:53 GMT (852kb,D)

Download:
- PDF
- Other formats

(license)

Current browse context:
stat.ML
< prev | next >
new | recent | 1806

Change to browse by:
cs
cs.LG
hep-ex
physics
physics.data-an
stat
stat.ME

References & Citations
- INSPIRE HEP (refers to | cited by)
- NASA ADS

Bookmark (what is this?)

recently published on a stat.ML preprint arxiv.org/abs/1806.04743
feedback is greatly welcomed (DM @pablodecm or pablo.decastro@cern.ch)

CONCLUSIONS AND PROSPECTS

Presented a machine learning approach that directly optimises an inference-guided non-decomposable loss accounting for the effect of model uncertainties

Flexibility of current autodiff frameworks allows the inclusion of nuisance parameters effect (via derivatives) over the training batches

The application of this approach to a realistic systematics-dominated benchmark (e.g. systematic-extended Higgs benchmark dataset) could shed some light on its real-world usefulness

Working on an update to the paper to be released together with TensorFlow implementation code with more involved examples.

THANK YOU
FOR YOUR ATTENTION

Will also be presenting this work at the poster session from 6pm,
you are all invited to pass by to discuss this and related work!