


# Accurate Surrogate Amplitudes with Calibrated Uncertainties

Henning Bahl


based on  
2412.12069 in collaboration with Nina Elmer, Luigi Favaro,  
Manuel Haußmann, Tilman Plehn, and Ramon Winterhalder



Particle Theory Seminar, Würzburg, 24.7.2025

# The challenge ahead


- general trend: larger-and-larger experiments collecting more-and-more data
- e.g. LHC: already enormous dataset will be further enlarged by a factor  $\sim 10$
- costs for future experiments increasing




Fully exploit the available data!



- new analysis methods
- theory precision  $\simeq$  experimental precision
- in particular: high-precision MC simulation




# The particle physics workflow




ML can help with each of these steps by increasing

- accuracy/performance and/or
- increase speed

# ML for particle physics




precision



speed




control



physics

# Amplitude surrogates

$$|\mathcal{M}|^2 \approx$$




# Case study for amplitude surrogates

- evaluating analytic expressions for amplitudes  $|\mathcal{M}|^2$  can be very expensive due to
  - higher-order corrections
  - large final-state multiplicities
- possible solution:
  - generate small training sample using full analytic expression
  - train a NN to approximate  $|\mathcal{M}|^2$
  - generate events using NN surrogate, which is much faster to evaluate

→ Does this work?

# Comparison to classical interpolation


[Bresó et al., 2412.09534]



→ ML surrogates outperform classical interpolation techniques

# Speed comparison

[Janßen et al., 2301.13562]



$$f_{\text{eff}} = \frac{T_{\text{standard}}}{T_{\text{surrogate}}}$$

dipole vs naïve:  
encode singularity structure of amplitudes

Large speed-ups possible!

Process	SHERPA default			with dipole-model surrogate				
	$t_{\text{ME}}[\text{ms}]$	$t_{\text{PS}}[\text{ms}]$	$\epsilon_{\text{full}}$	$t_{\text{surr}}[\text{ms}]$	$x_{\text{max}}$	$\epsilon_{1\text{st,surr}}$	$\epsilon_{2\text{nd,surr}}$	$f_{\text{eff}}$
$gg \rightarrow e^- e^+ gg d\bar{d}$	54	0.40	1.411 %	0.14	2.6	1.418 %	39 %	16
$gg \rightarrow e^- e^+ ggg d\bar{d}$	16 216	5.70	0.076 %	0.20	3.6	0.085 %	29 %	269

Can we also control the uncertainties?



# Learning uncertainties

*"All models are wrong, but some — those that know when they can be trusted — are useful!"*

— George Box (adapted)



# Regression with uncertainties



- statistical or epistemic uncertainty  $\hat{=}$  lack of training data
- systematic or aleatoric uncertainty  $\hat{=}$  noise in the data, lack in model expressivity

# Systematic uncertainty: heteroskedastic loss

- log-likelihood loss:

$$\mathcal{L} = - \sum_{x_i, A_i \in D_{\text{train}}} \log p(A_{\text{true}}(x_i) | x_i, \theta)$$

sum over training dataset

true amplitudes

NN parameters

phase-space point

- assume Gaussian likelihood:  $p(A|x, \theta) = \mathcal{N}(\bar{A}(x), \sigma_{\text{syst}}^2(x))$


- NN learns both:  $\bar{A}(x)$  and  $\sigma_{\text{syst}}(x)$

$$\Rightarrow \text{heteroskedastic loss: } \mathcal{L} = \sum_i \left[ \frac{(\bar{A}(x_i) - A_{\text{true}}(x_i))^2}{2\sigma_{\text{syst}}^2(x_i)} + \log(\sigma_{\text{syst}}(x_i)) \right]$$



- constant  $\sigma_{\text{syst}}$   $\rightarrow$  recovers MSE loss

# Statistical uncertainty: repulsive ensemble

- train ensemble of networks
- ensure convergence to correct posterior via repulsive interaction between ensemble members
- each network leads to slightly different result
- spread of network predictions  $\sim$  statistical uncertainty
- less data  $\rightarrow$  higher spread




# Repulsive ensemble + heteroskedastic loss



Combined learnable modelling of systematic and statistical uncertainties!


# Alternative: Bayesian NNs




- promote NN parameters to Gaussians  $q(\theta)$
- for each evaluation, sample from Gaussians
- learn means and widths

# $gg \rightarrow \gamma\gamma g$ dataset

- generated using Sherpa+Njet
- 1.1 million events, 70% used for training
- preprocessing:
  - learn logarithm of amplitude
  - rescale inputs and log-amplitudes to have zero mean and unit standard deviation
- amplitude symmetries
  - Lorentz-invariant
  - permutation-invariance w.r.t. identical particles



# Behavior of learned uncertainties





Test: apply different levels of Gaussian noise to amplitudes

- statistical uncertainty decreases with more training data
- systematic uncertainty converges to level of applied noise

$$A_{\text{train}} \sim \mathcal{N}(A_{\text{true}}, \sigma_{\text{train}}^2)$$
$$\sigma_{\text{train}} = f_{\text{smear}} A_{\text{true}}$$

# Extracting the noise level



with  $\sigma_{\text{noise}}^2 = \sigma_{\text{syst}}^2 - \sigma_{\text{syst},0}^2$  and  $\sigma_{\text{syst},0}^2$  being the systematic unc. due to limited NN expressivity.



Able to reliably extract noise level!

# Are these uncertainties calibrated?


- statistical uncertainties play minor role for amplitude regression
- define systematic pull:

$$t_{\text{syst}} = \frac{\langle A \rangle(x) - A_{\text{train}}(x)}{\sigma_{\text{syst}}(x)}$$




- if calibrated,  $t_{\text{syst}}$  distribution should follow  $\mathcal{N}(0, 1)$



Almost perfectly calibration → reliable uncertainty estimate




# Dependence on smearing




- BNN and deterministic models  
→ well calibrated for different smearing levels
- repulsive ensemble overestimates uncertainty for low smearing  
→ see below

# Pushing for precision

activations, layers, and domain knowledge



# Enhancing NN expressivity




Can we improve accuracy for low smearing?




Improve NN expressivity:

- different activations, NN structures
- more layers,
- exploit symmetries

# Activation functions and NN structure



# Adding more layers




adding more layers




convergence towards noise level

# Encoding our physics knowledge



Large gain in NN accuracy! Also found uncertainties still to be well calibrated.





# Ensembling and biases

The wisdom (and biases) of the crowd.



[HB, Elmer, Plehn, Winterhalder, to appear]

# Ensembling and biases

- ensembling averages out noise → widely used
- but, there can be systematic biases




# Extracting the biases



ensembling increases precision, but gains level off for  $N_{\text{ens}} \gtrsim 100$

# Miscalibration of ensemble systematic



$$A_{\text{NN}}(x) = \frac{1}{N_{\text{ens}}} \sum_{i=1}^{N_{\text{ens}}} A_{\text{NN},i}(x)$$
$$\sigma^2(x) = \frac{1}{N_{\text{ens}}} \sum_{i=1}^{N_{\text{ens}}} \sigma_i^2(x).$$



- averaging systematic unc. of ensemble members  $\sim$  average uncertainty of each member
- but ensemble more precise

# Fixing ensemble systematic unc.

$$A_{\text{NN}}(x) = \frac{1}{N_{\text{ens}}} \sum_{i=1}^{N_{\text{ens}}} A_{\text{NN},i}(x)$$

$$\mathcal{L}_{\sigma} = \frac{1}{B} \sum_{b=1}^B \left[ \frac{|A_{\text{train}}(x_b) - A_{\text{NN}}(x_b)|^2}{2\sigma^2(x_b)} + \log \sigma(x_b) \right]$$

solution: train separate network with loss  $\mathcal{L}_{\sigma}$  to predict systematic uncertainty of ensemble



lower noise levels  $\rightarrow$  non-Gaussian biases become visible

# Localized smearing

Is the systematic uncertainty able to pick-up local noise?



[HB et al., to appear]

# Box smearing

- analytic amplitude evaluation might loss precision close to thresholds, kinematic edges, ...
- can the learned systematic uncertainties pick this up?
- test this by inducing local noise to our amplitudes
- simple start: box smearing with  $m_{\text{thresh}} = 200 \text{ GeV}$


$$A_{\text{train}}(x) = \begin{cases} \mathcal{N}(A_{\text{true}}(x), \epsilon A_{\text{true}}(x)) & \text{if } |m_{\gamma\gamma g}(x) - m_{\text{thresh}}| < w \\ A_{\text{true}} & \text{if } |m_{\gamma\gamma g}(x) - m_{\text{thresh}}| \geq w \end{cases}$$

# Box smearing — results




$\sigma_{\text{syst}}$  closely follows expected behavior



# Box smearing — calibration



also good calibration



# Peaked smearing



$$A_{\text{train}}(x) = \mathcal{N}\left(A_{\text{true}}(x), \frac{\epsilon m_{\text{thresh}}}{|m_{\gamma\gamma g}(x) - m_{\text{thresh}}|} A_{\text{true}}(x)\right) \rightarrow \text{well captured by systematic uncertainties}$$

# Conclusions



# Conclusions

- amplitude surrogates → speed up MC generation
- uncertainty-aware NNs allow for controlled modelling of
  - systematic uncertainties: data-inherent noise + model expressivity
  - statistical uncertainties: lack of data
- methods: heteroskedastic loss, repulsive ensemble, Bayesian NNs
- learned uncertainties reflect actual deviations from truth well
- encoding physics knowledge increase accuracy with still well-calibrated unc.



same techniques also applicable to all kind of other problems

# Appendix

# Kolmogorov-Arnold theorem

Any  $f: [0,1]^n \rightarrow \mathbb{R}$  can be written as a finite decomposition of univariate functions and addition.


$$f(x) = f(x_1, \dots, x_n) = \sum_{q=1}^{2n+1} \Phi_q \left( \sum_{p=1}^n \phi_{q,p}(x_p) \right)$$

where  $\phi_{q,p} : [0, 1] \rightarrow \mathbb{R}$  and  $\Phi_q : \mathbb{R} \rightarrow \mathbb{R}$ .


# Kolmogorov-Arnold Networks (KANs)

- KAN layer: 
$$x_{l+1} = \underbrace{\begin{pmatrix} \phi_{l,1,1}(\cdot) & \phi_{l,1,2}(\cdot) & \cdots & \phi_{l,1,n_l}(\cdot) \\ \phi_{l,2,1}(\cdot) & \phi_{l,2,2}(\cdot) & \cdots & \phi_{l,2,n_l}(\cdot) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{l,n_{l+1},1}(\cdot) & \phi_{l,n_{l+1},2}(\cdot) & \cdots & \phi_{l,n_{l+1},n_l}(\cdot) \end{pmatrix}}_{\equiv \Phi_l} x_l,$$
- KAN network:  $\text{KAN}(x) = (\Phi_{L-1} \circ \Phi_{L-2} \circ \dots \circ \Phi_1 \circ \Phi_0)x$
- GroupKAN layer:  $\text{activation}(x) \rightarrow \text{GroupKANlayer}_l(x) = \begin{pmatrix} \phi_{l,g_l(1)}(x_1) \\ \phi_{l,g_l(2)}(x_2) \\ \vdots \\ \phi_{l,g_l(n_l)}(x_{n_l}) \end{pmatrix}$
- GroupKAN network:  $\text{GroupKAN}(x) = (W_{L-1} \circ \text{GroupKANlayer}_{L-2} \circ W_{L-2} \circ \dots \circ \text{GroupKANlayer}_0 \circ W_0)x$
- Normal MLP network  $\text{MLP}(x) = (W_{L-1} \circ \text{activation} \circ W_{L-2} \circ \dots \circ \text{activation} \circ W_0)x$



# Kolmogorov-Arnold Networks (KANs)



# Learned activation functions



# BNN prior dependence



# Uncertainty overview

