Generative Networks for Fast Simulation
Statistics for Generative Networks

Sascha Cassandra Diefenbacher, Heidelberg University

WMLQ 2026, Warsaw

September 7th, 2026

Baden-Württemberg Stiftung

UNIVERSITÄT HEIDELBERG
ZUKUNFT SEIT 1386

ML in Fundamental Physics

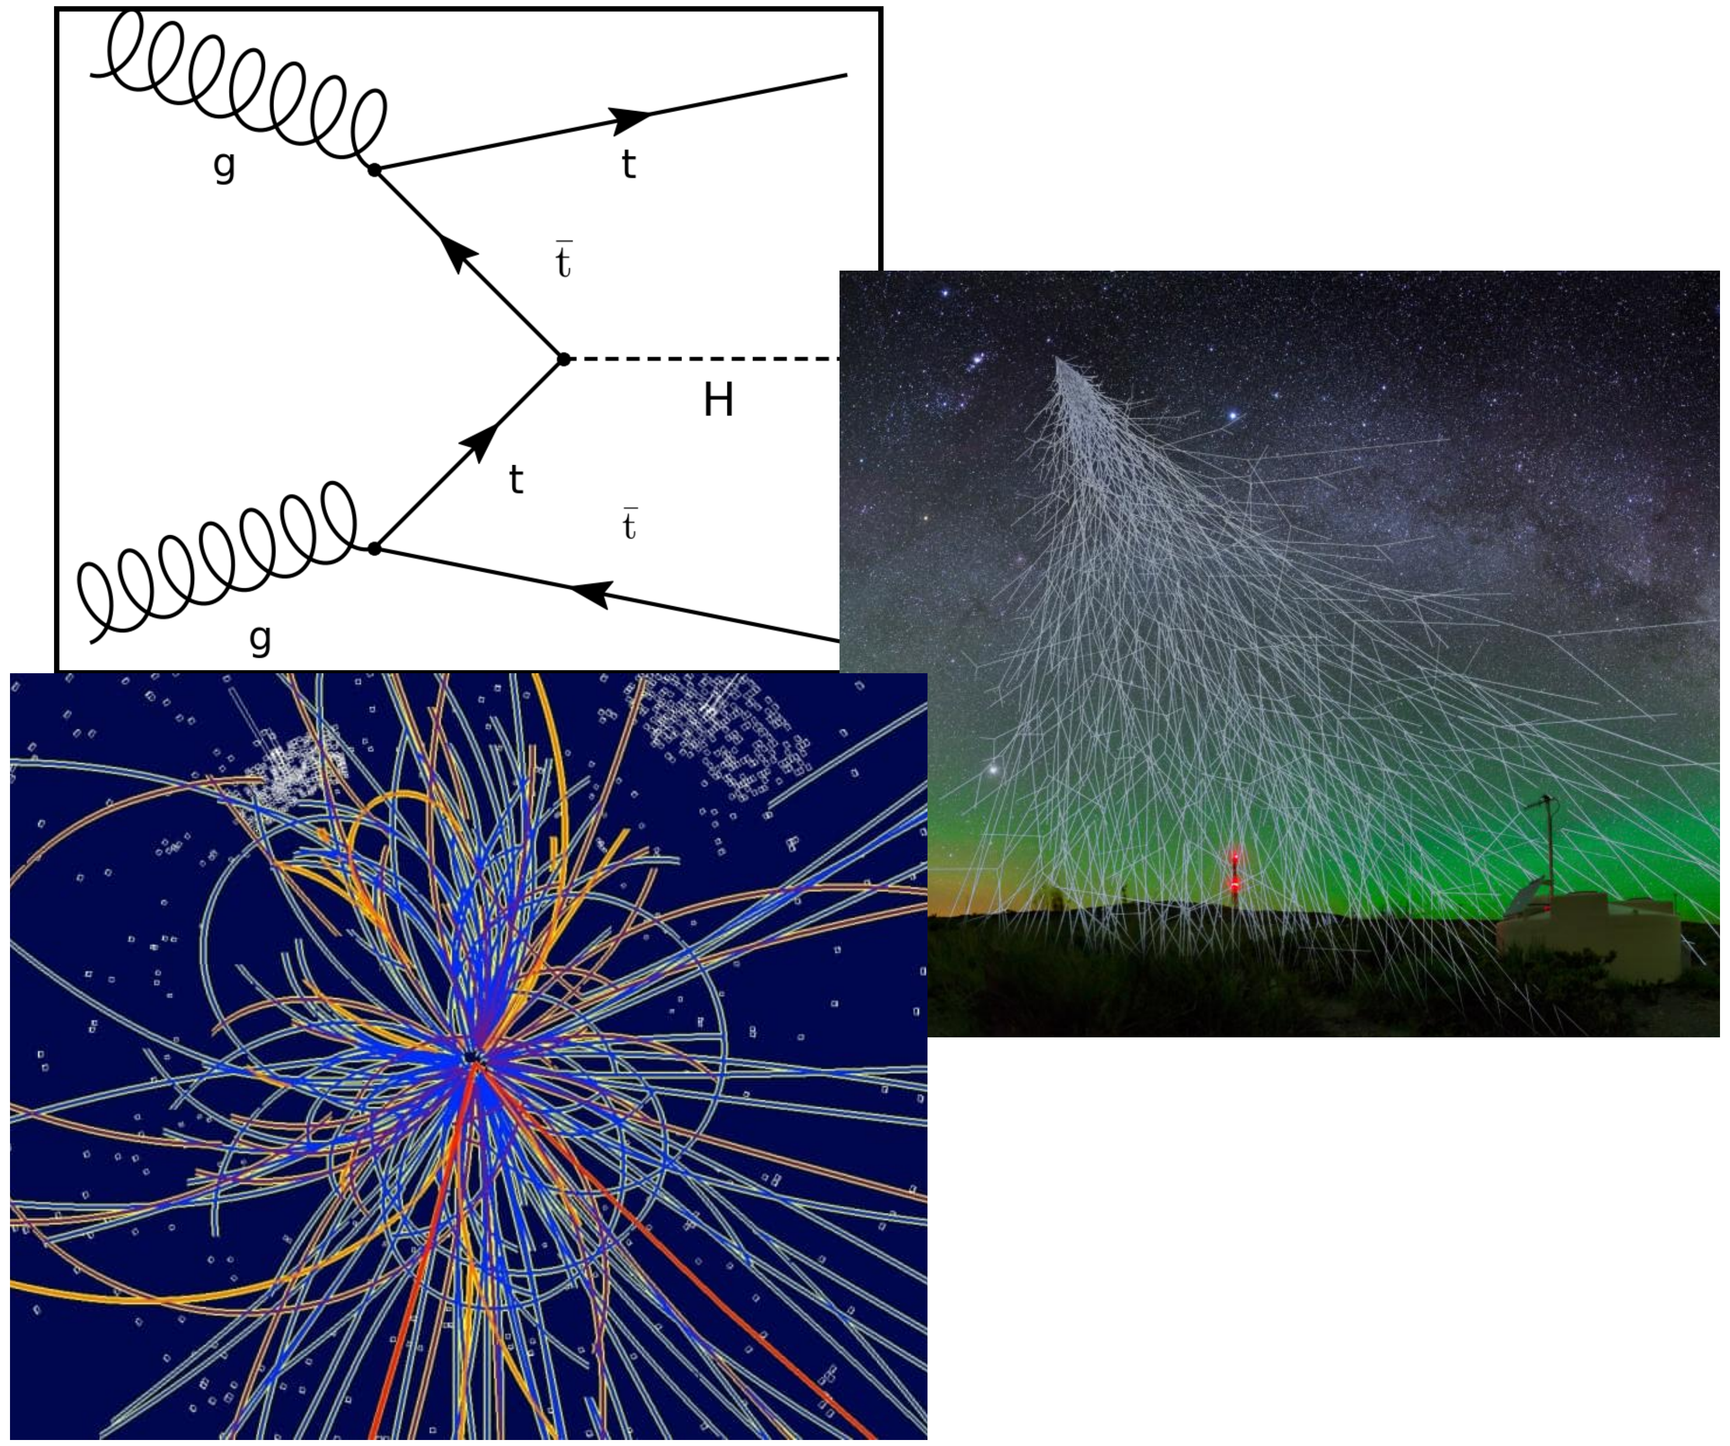

Challenge I: Complexity

- $t\bar{t}H$ events in terms of hadrons
- Air showers from ground observation
- Radio waves transformed to images

- Extremely high-dimensional
- Highly correlated
- Infer fundamental-theory parameters
  e.g. $t\bar{t}H$: masses, couplings, CP phase…

ML in Fundamental Physics

- LHC & SKA: O(PB) per second

- Challenging to select
- Expensive to analyze
- Needs even more simulated data

Challenge II: Large Datasets

ML in Fundamental Physics

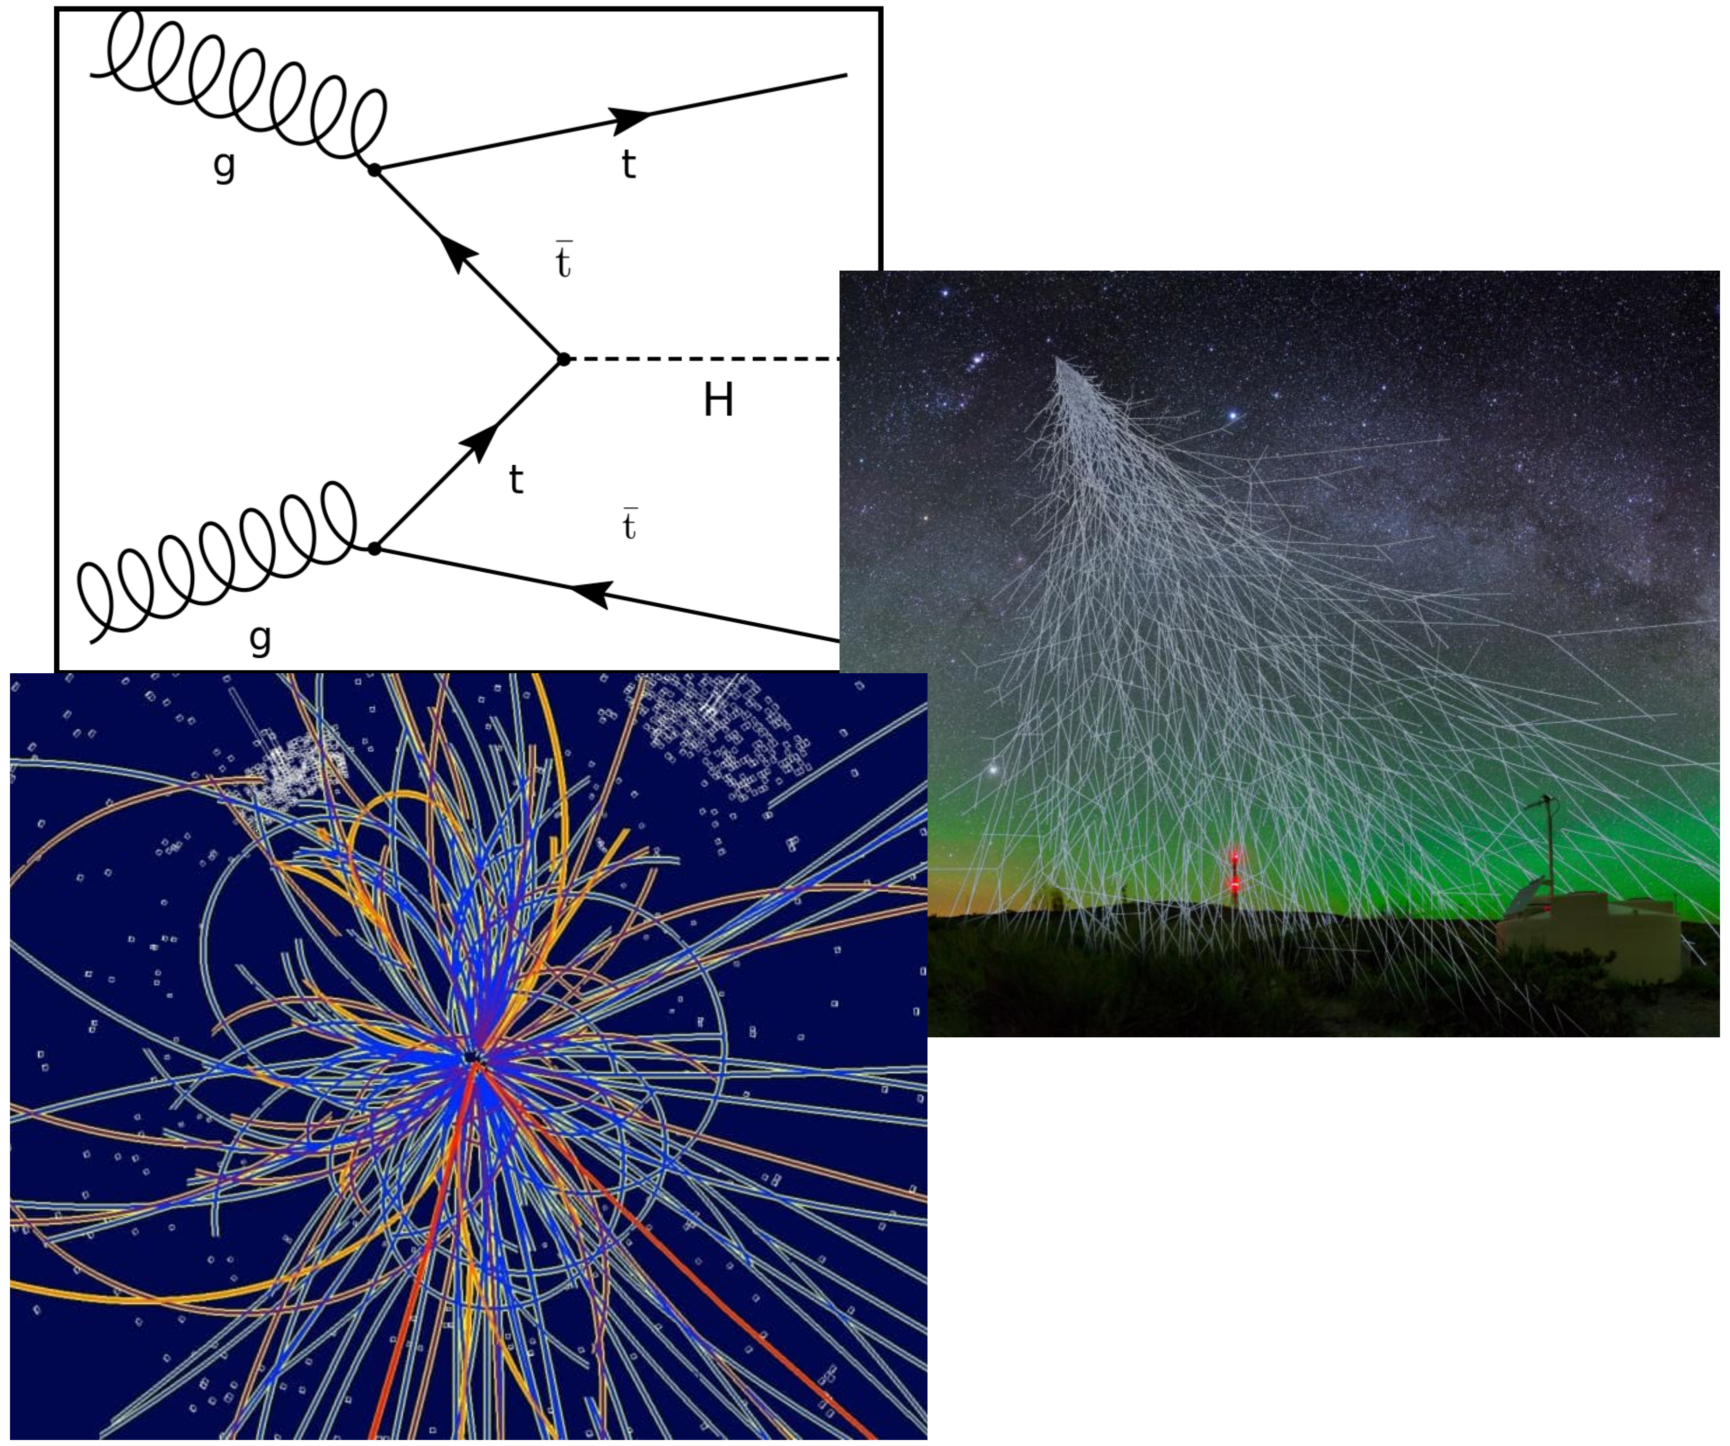

Challenge I: Complexity

Challenge II: Large Datasets

ML in Fundamental Physics

Solution:
Machine Learning

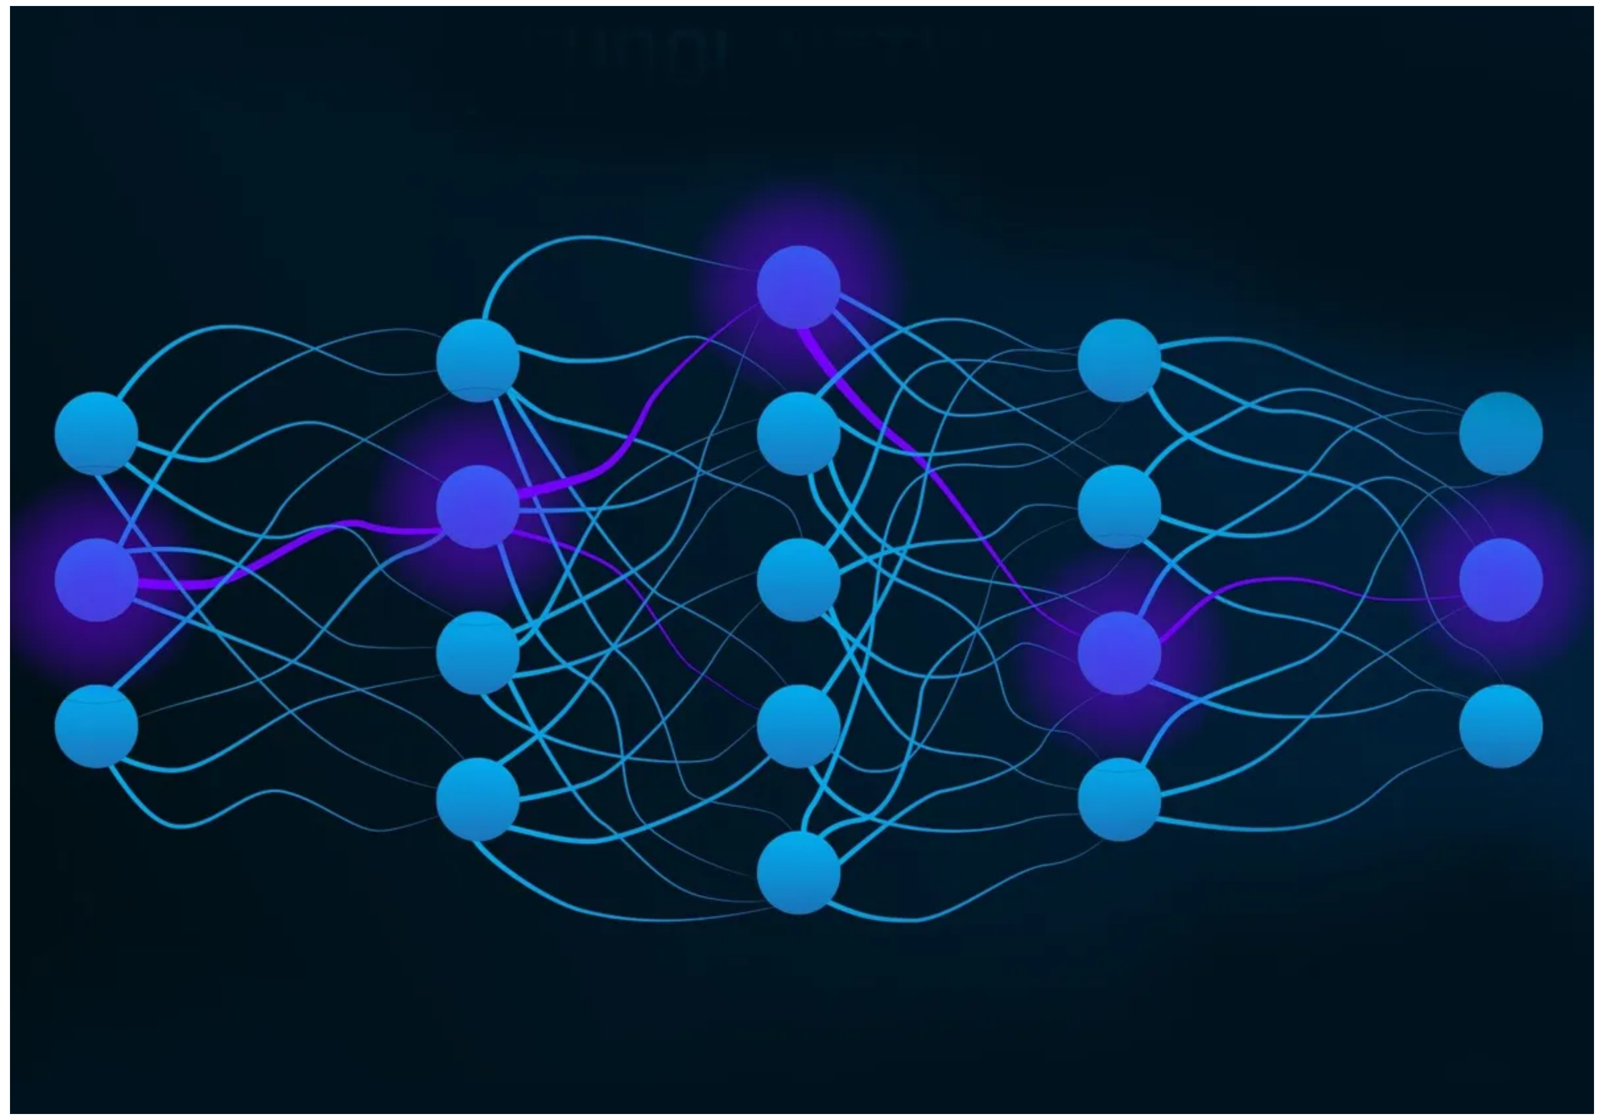

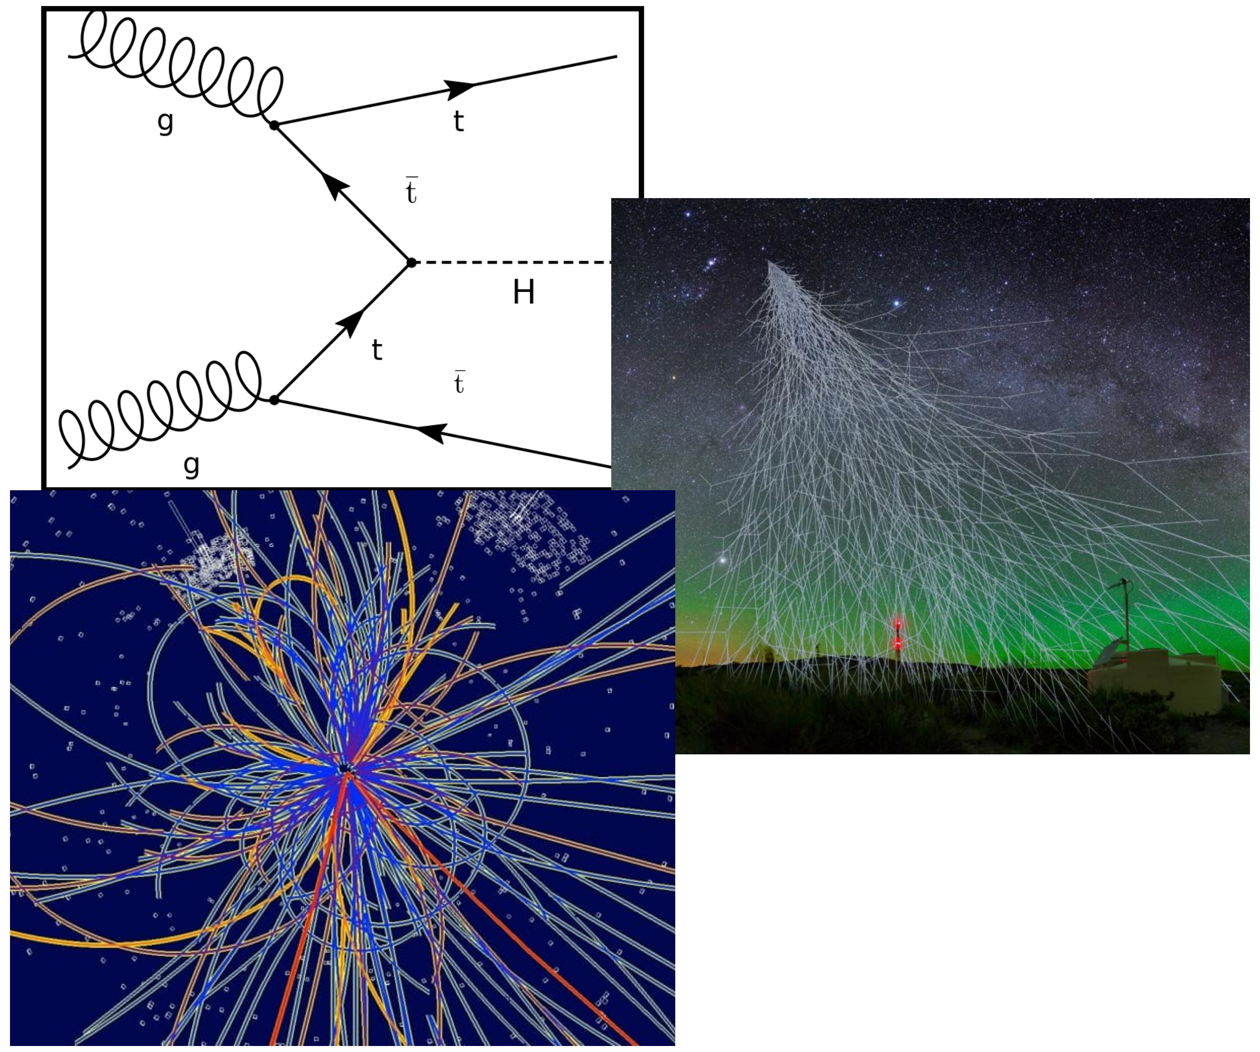

Challenge I: Complexity

Challenge II: Large Datasets

ML in Fundamental Physics

Solution:

Machine Learning

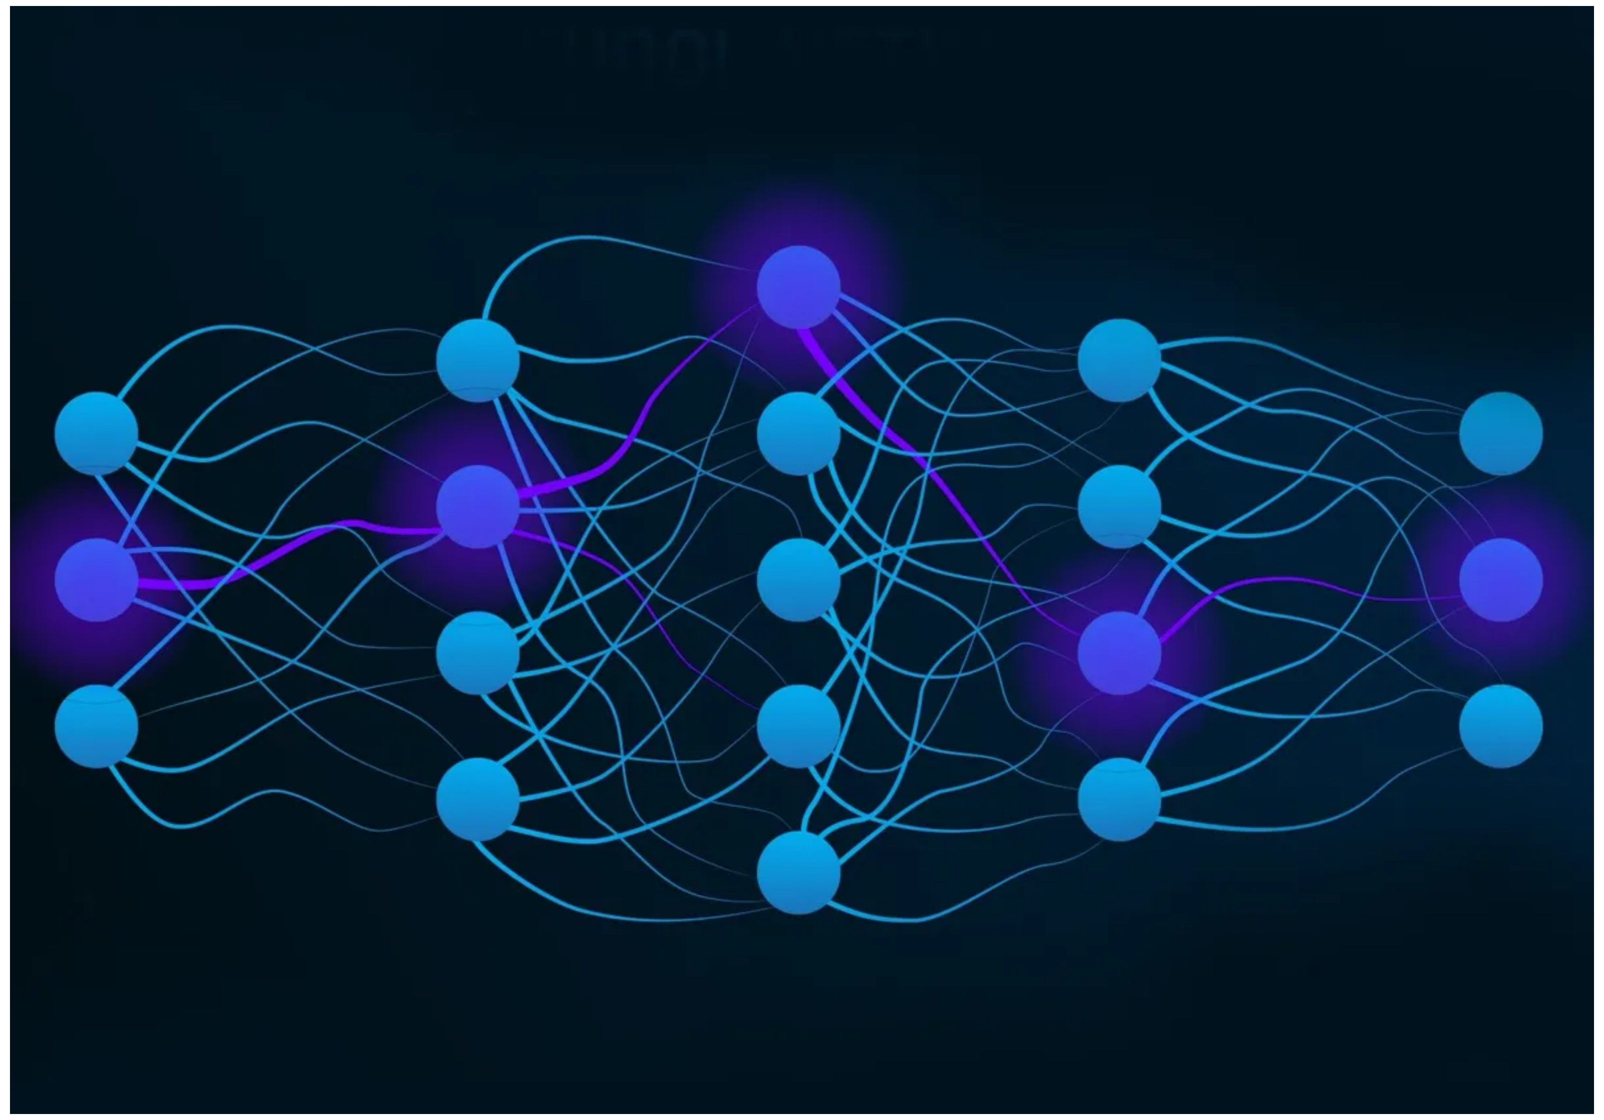

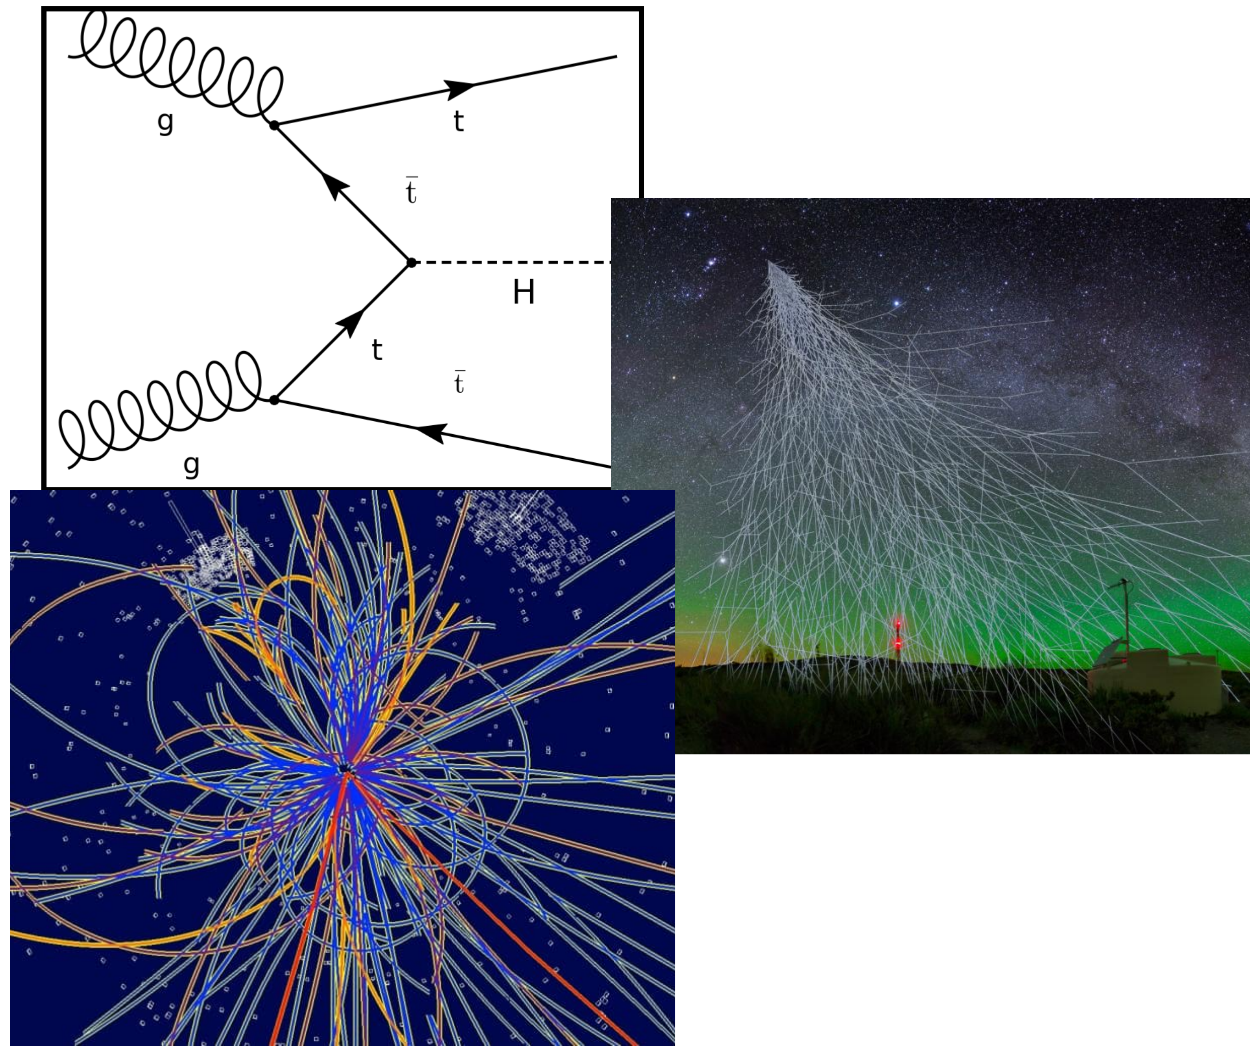

Understand Complexity

Using Large Datasets

Particle Physics Simulation

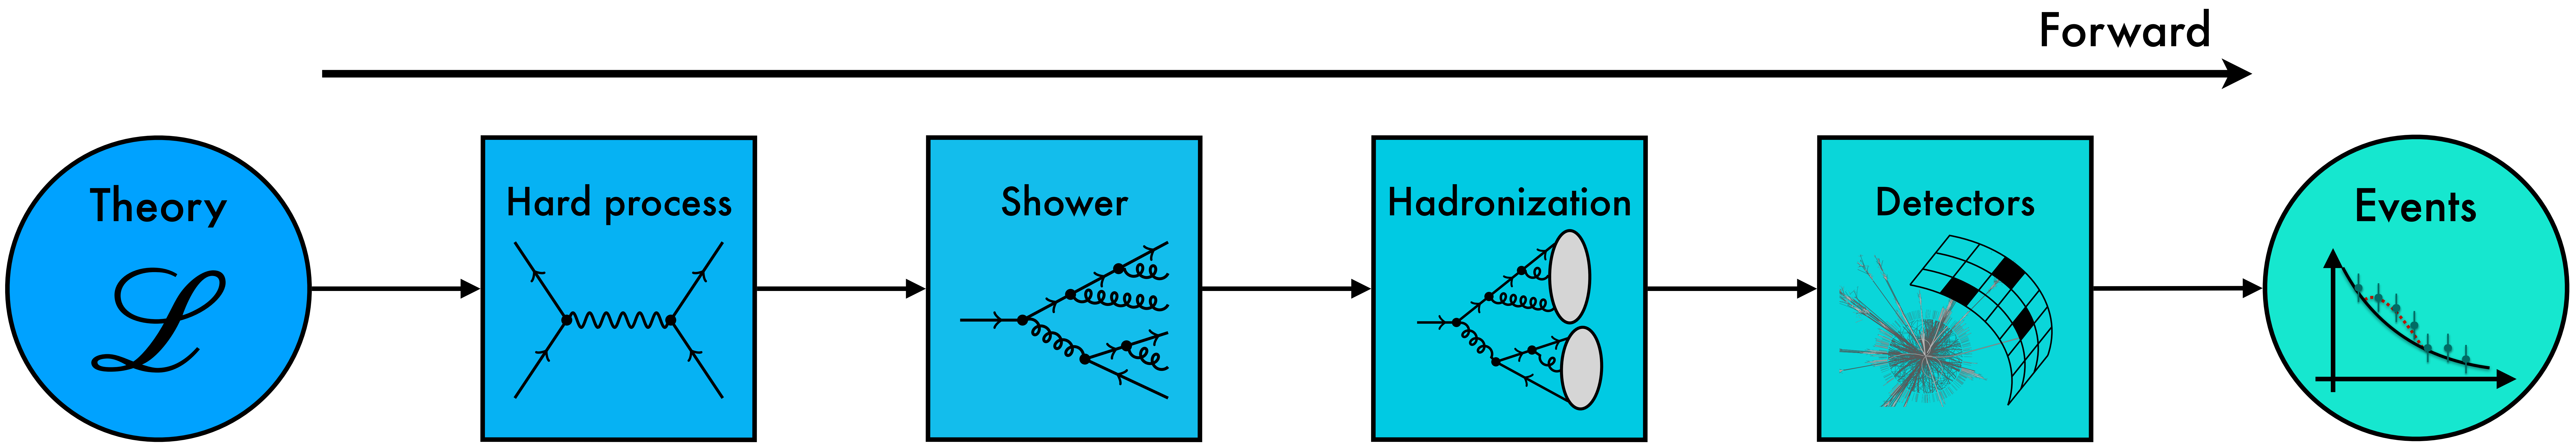

Hard Process

- Models initial particle collision
- Many processes with cross-sections
- Calculate cross-sections from theory model/Lagrangian
- Use Monte Carlo sampling
- Complex processes → expensive

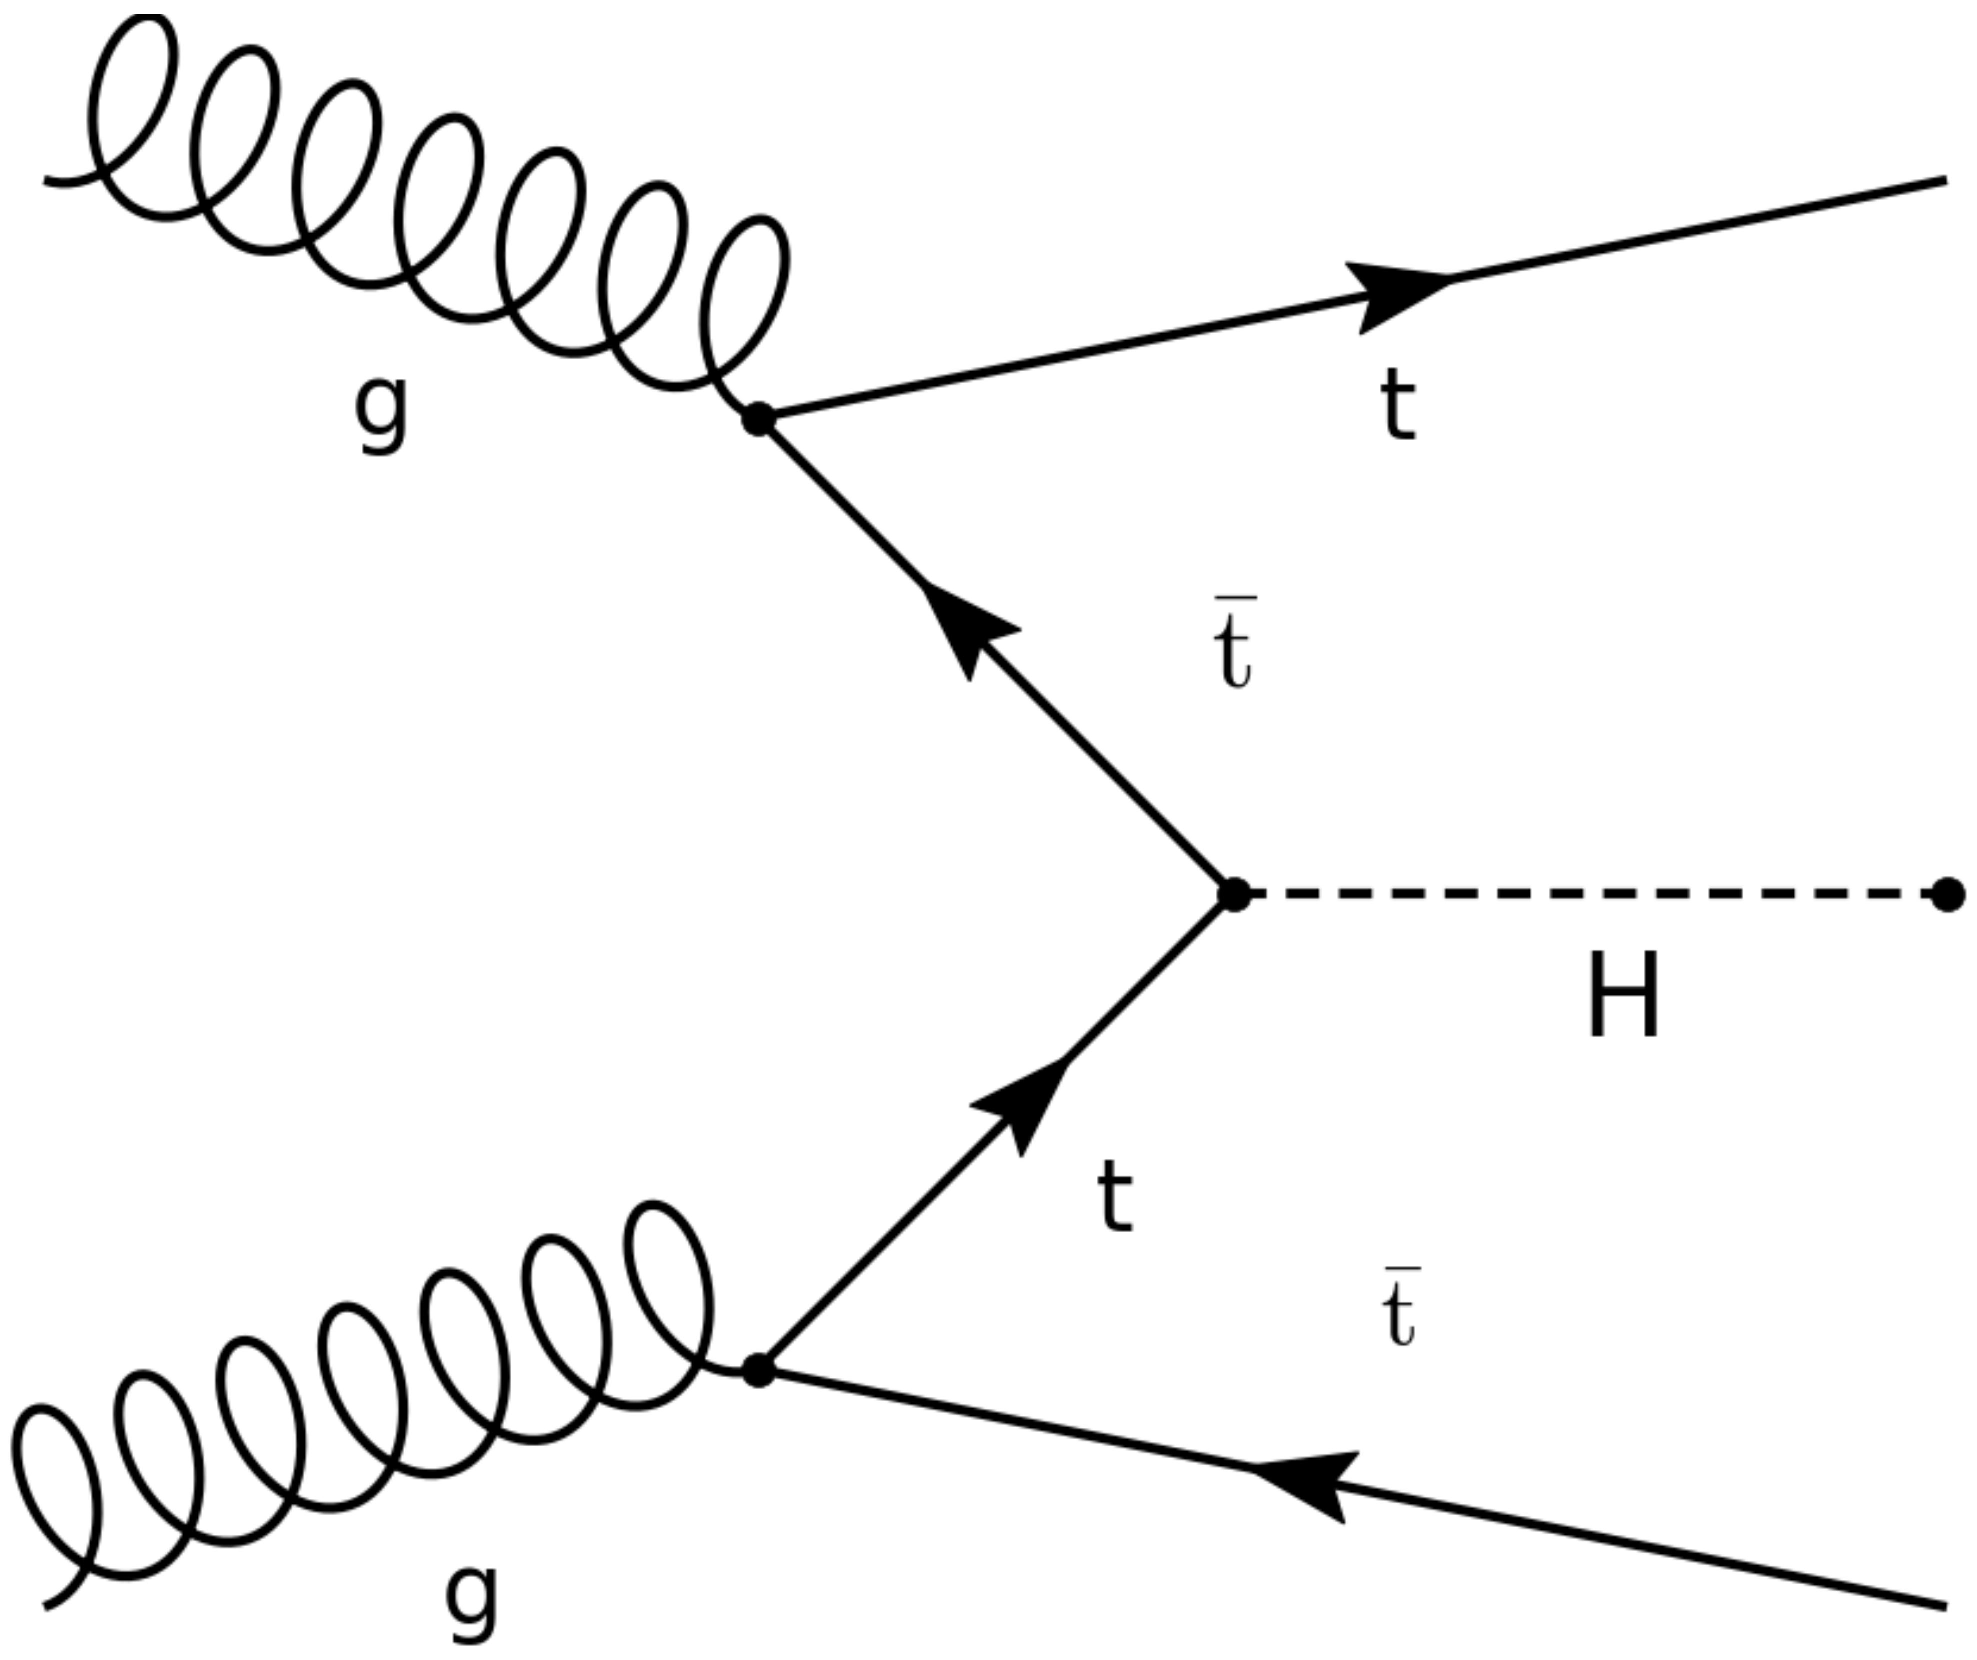

Particle Physics Simulation

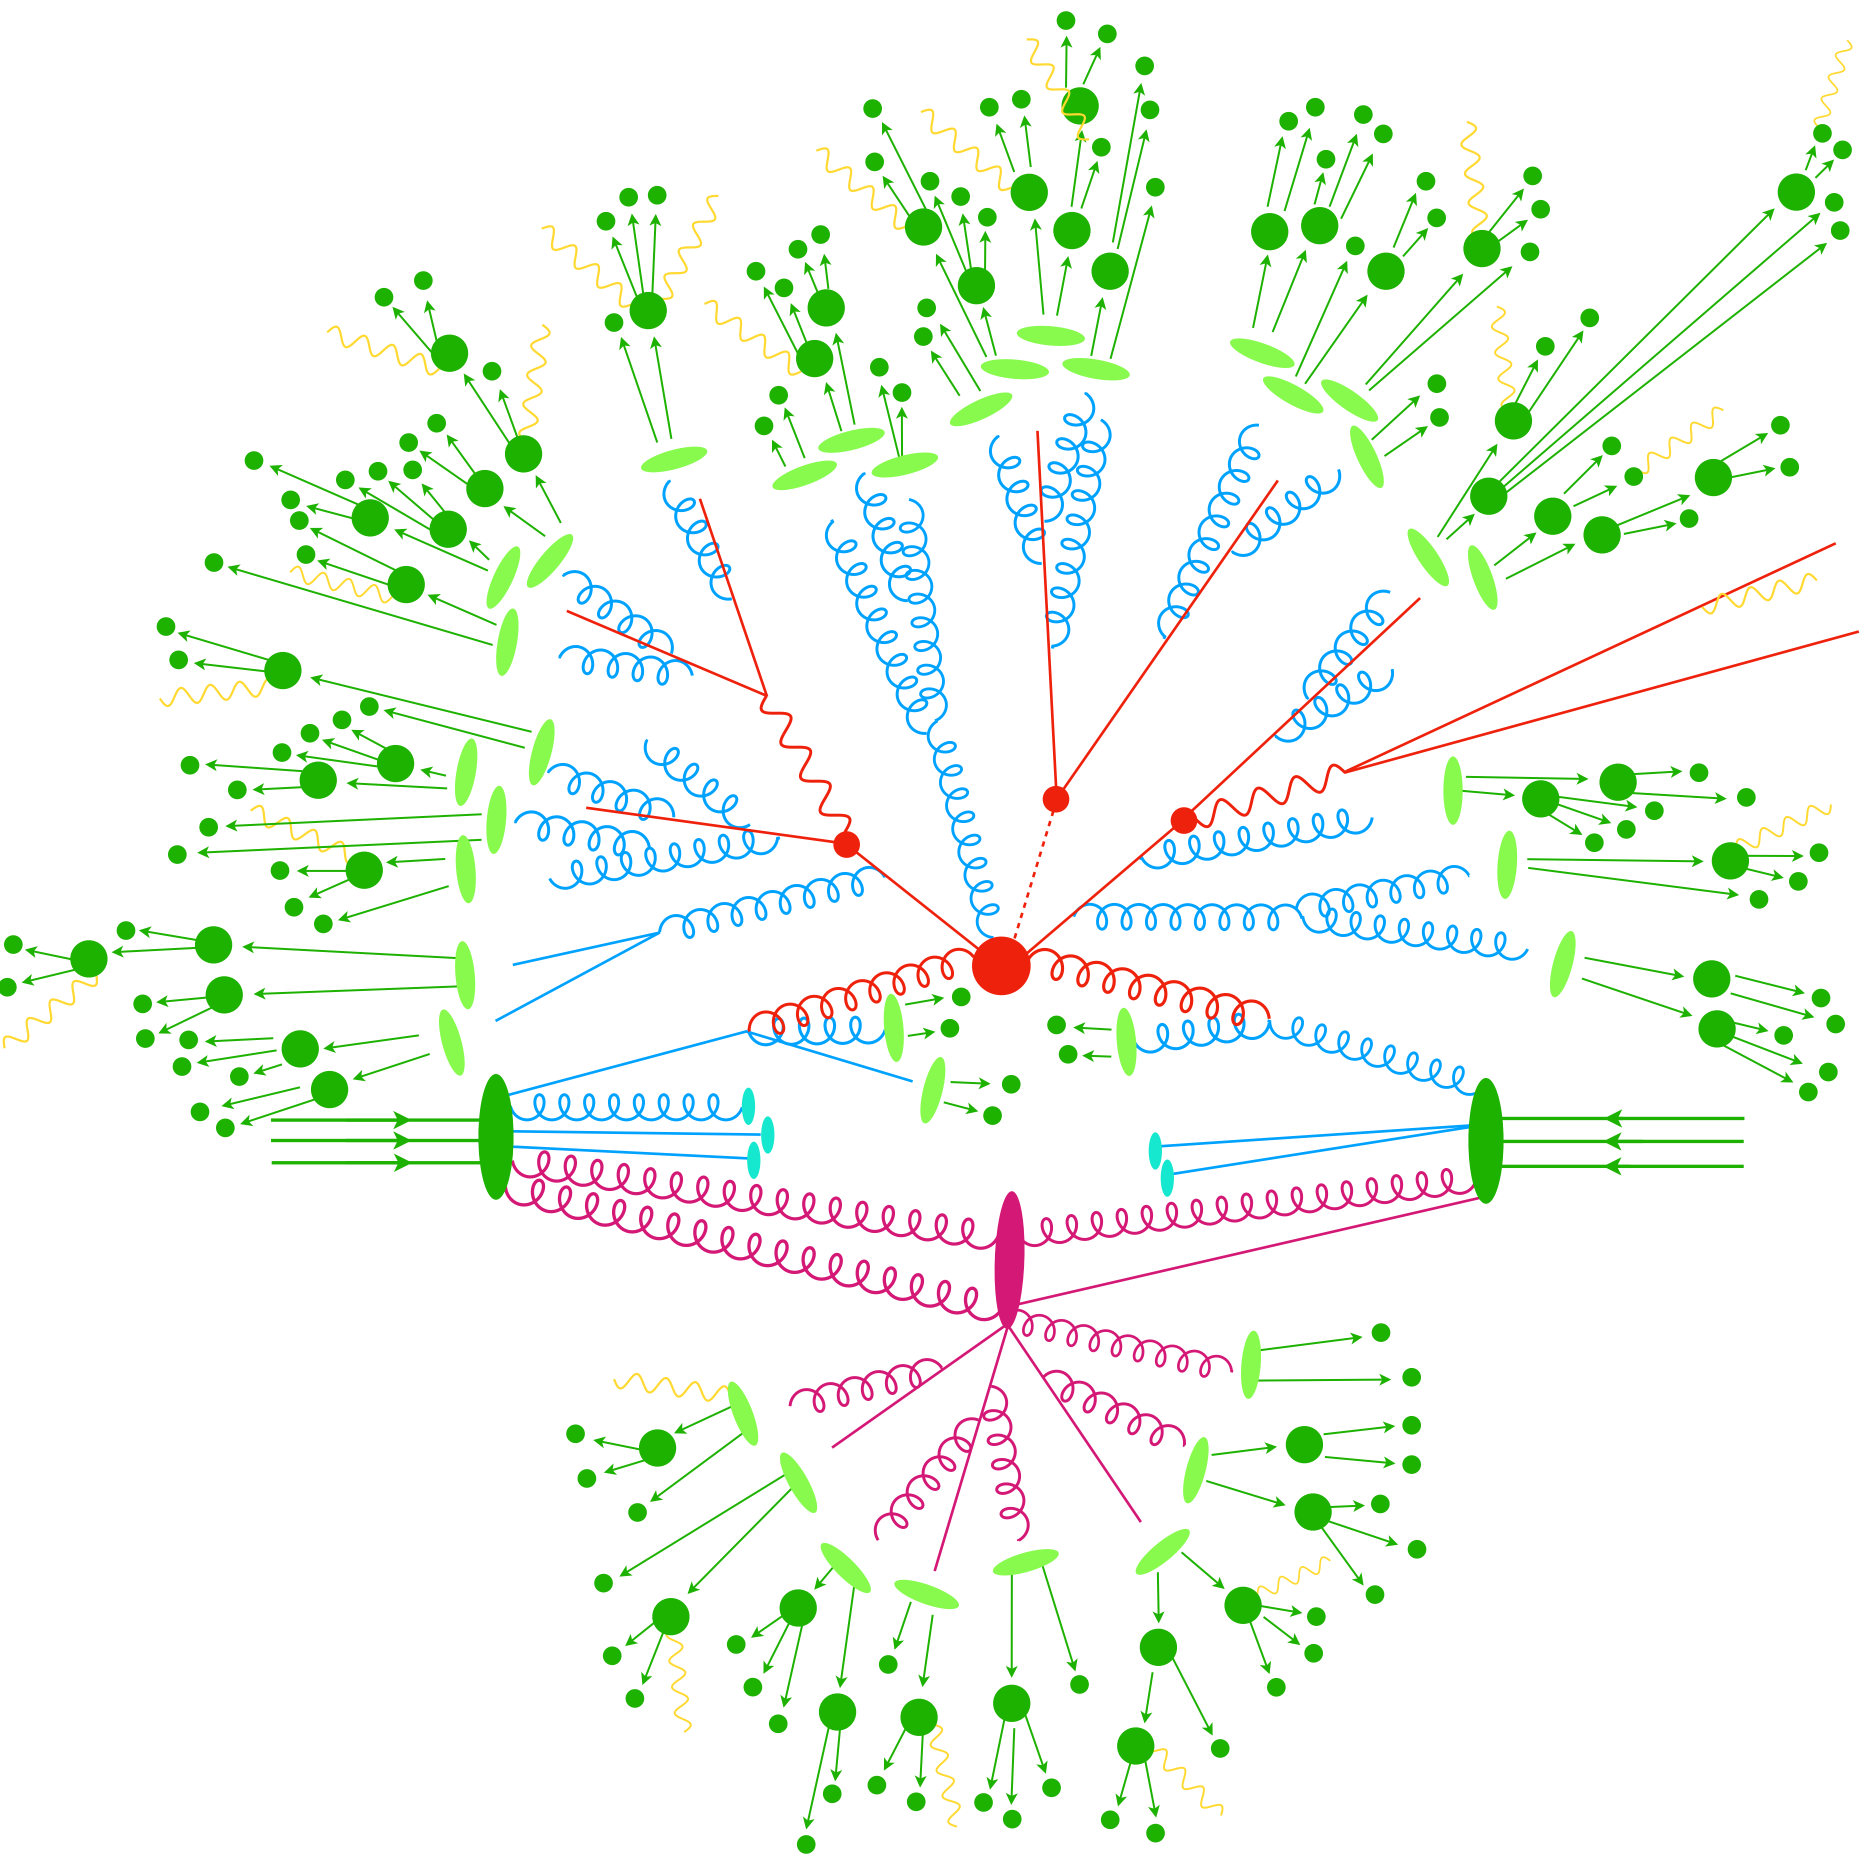

Showering/Hadronization

- Quarks only exist in bound states
- Any quark produced forms Hadron
- Propagating Hadrons emit more particles
- Collimated sprays of particles called jets
- Current models imperfect, difficult to simulate

Particle Physics Simulation

Detector Simulation

- Particles interact with detector material
- Ionization: leaves signature in detector
- Bremsstrahlung, Pair-production, Nuclear Interaction, Spallation → more particles
- New particles continue to interact
- All need to be simulated → expensive

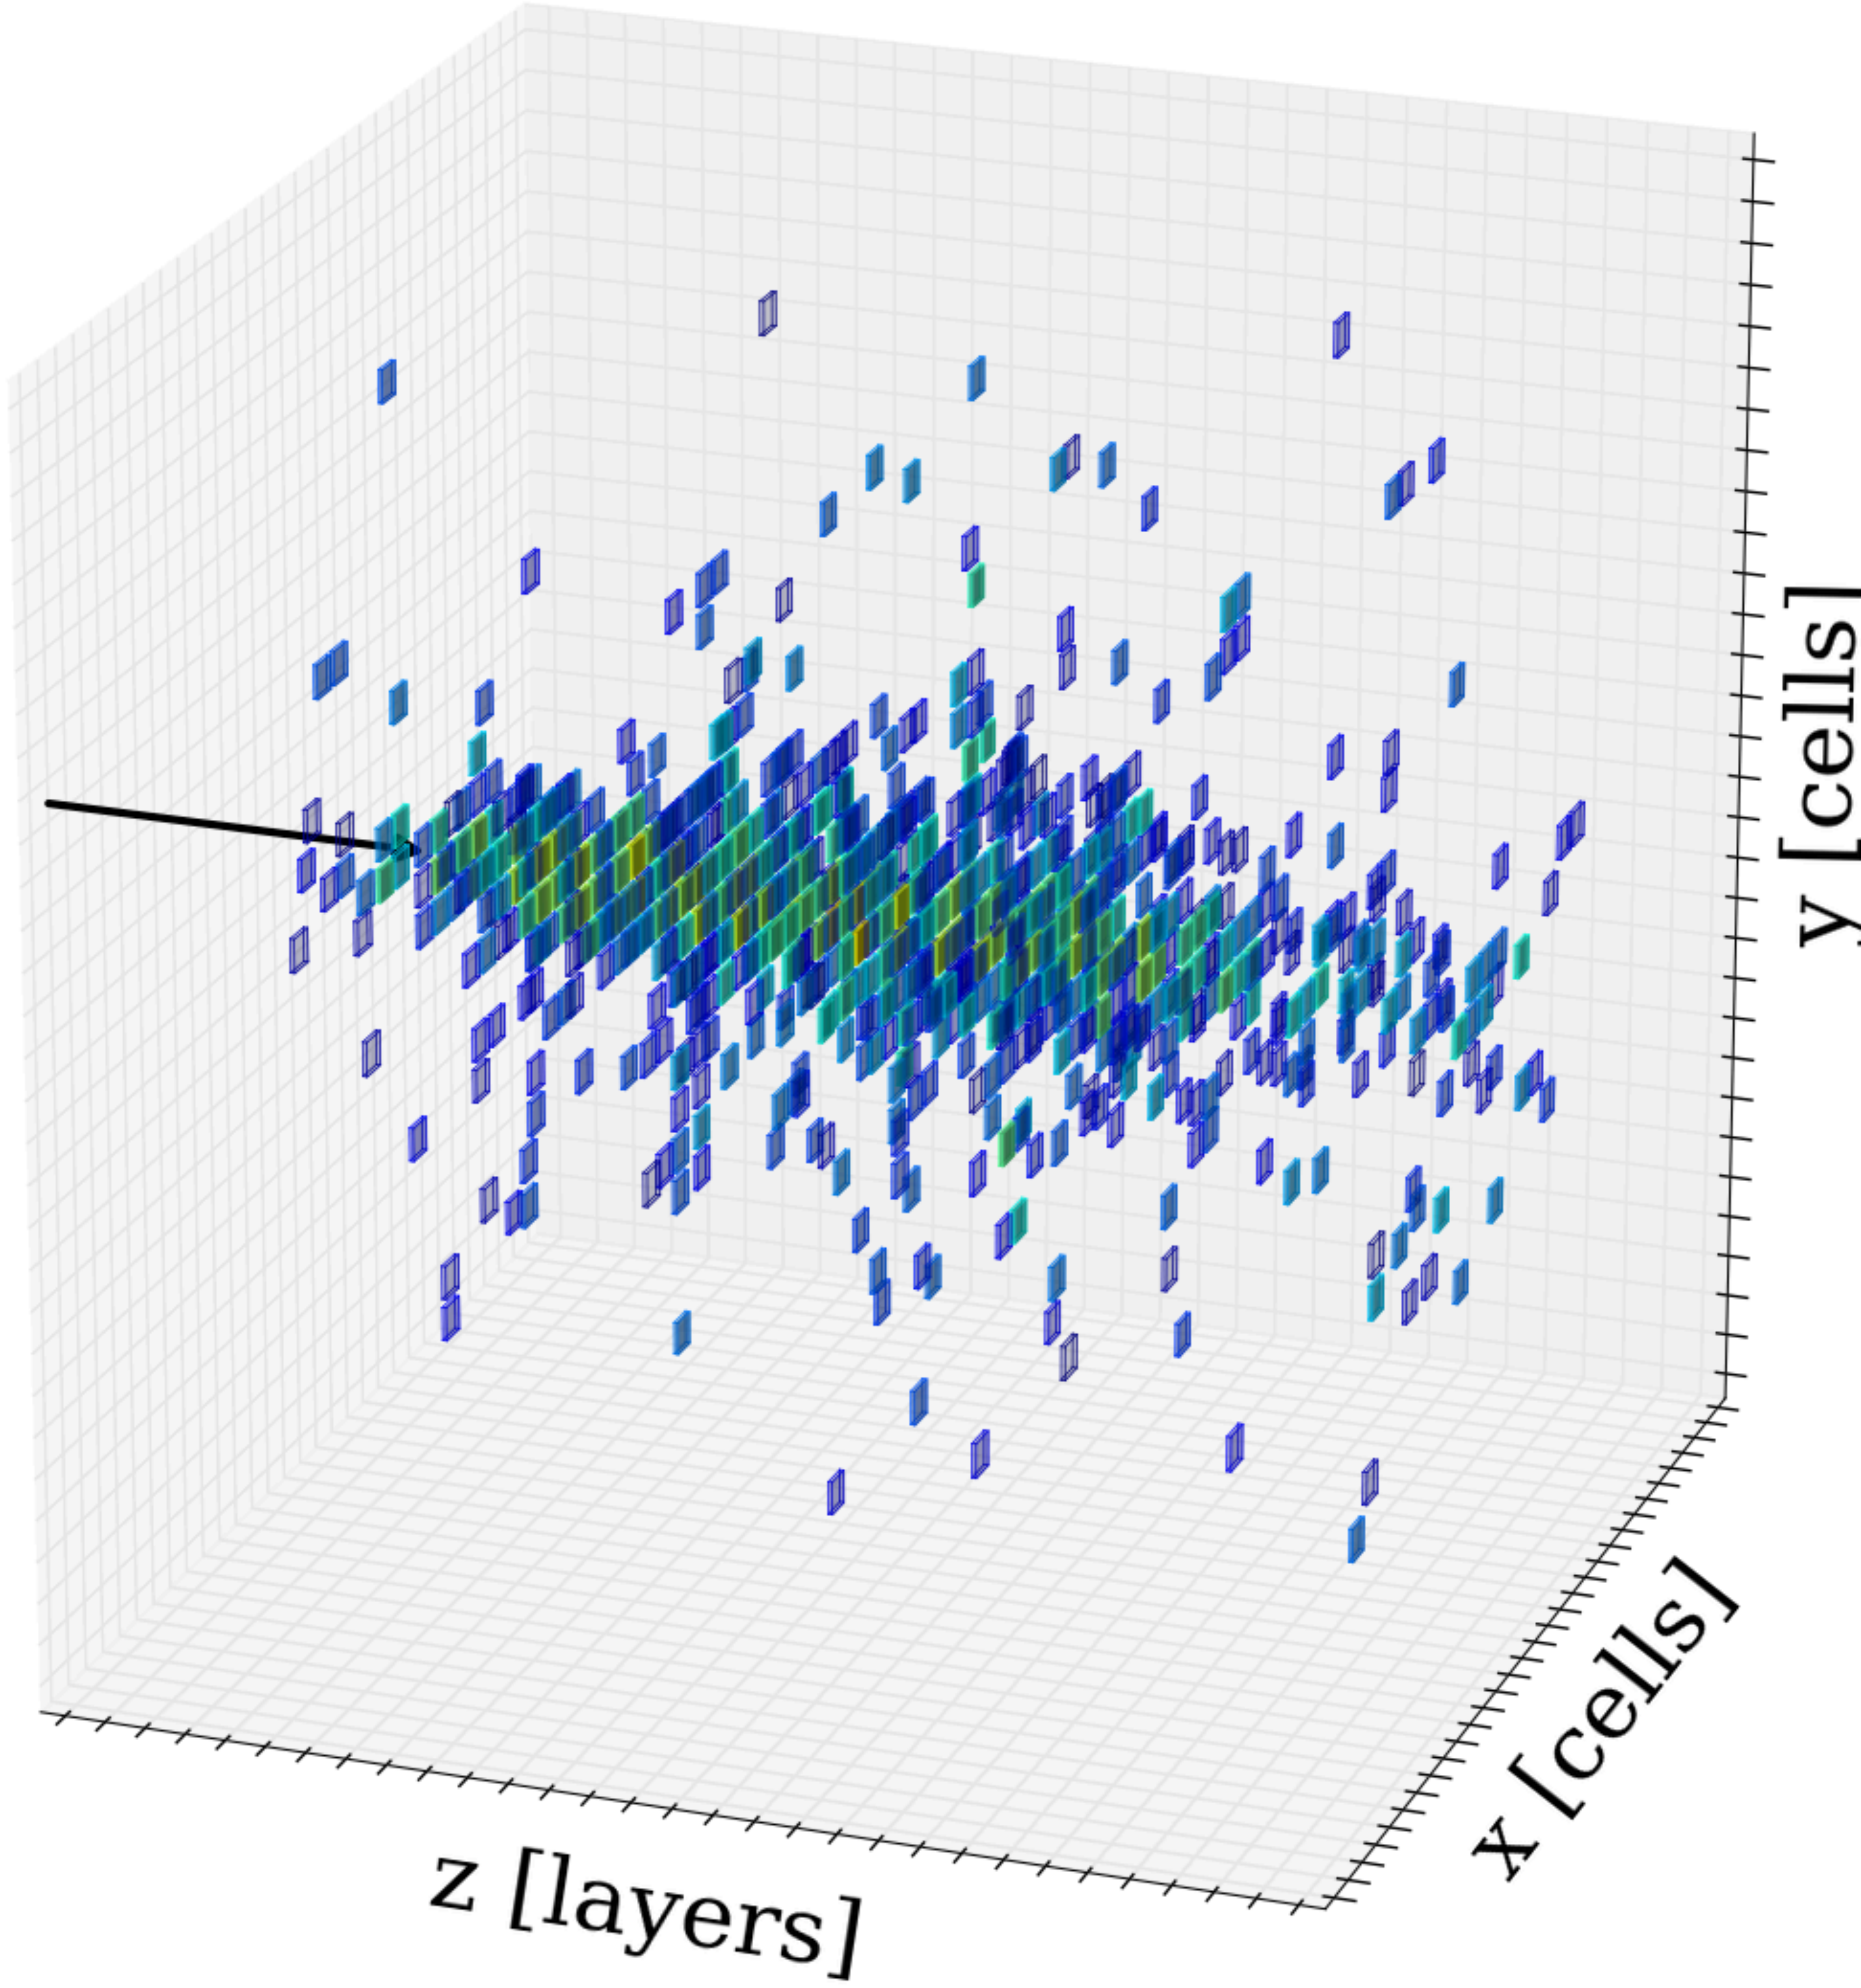

Particle Physics Simulation

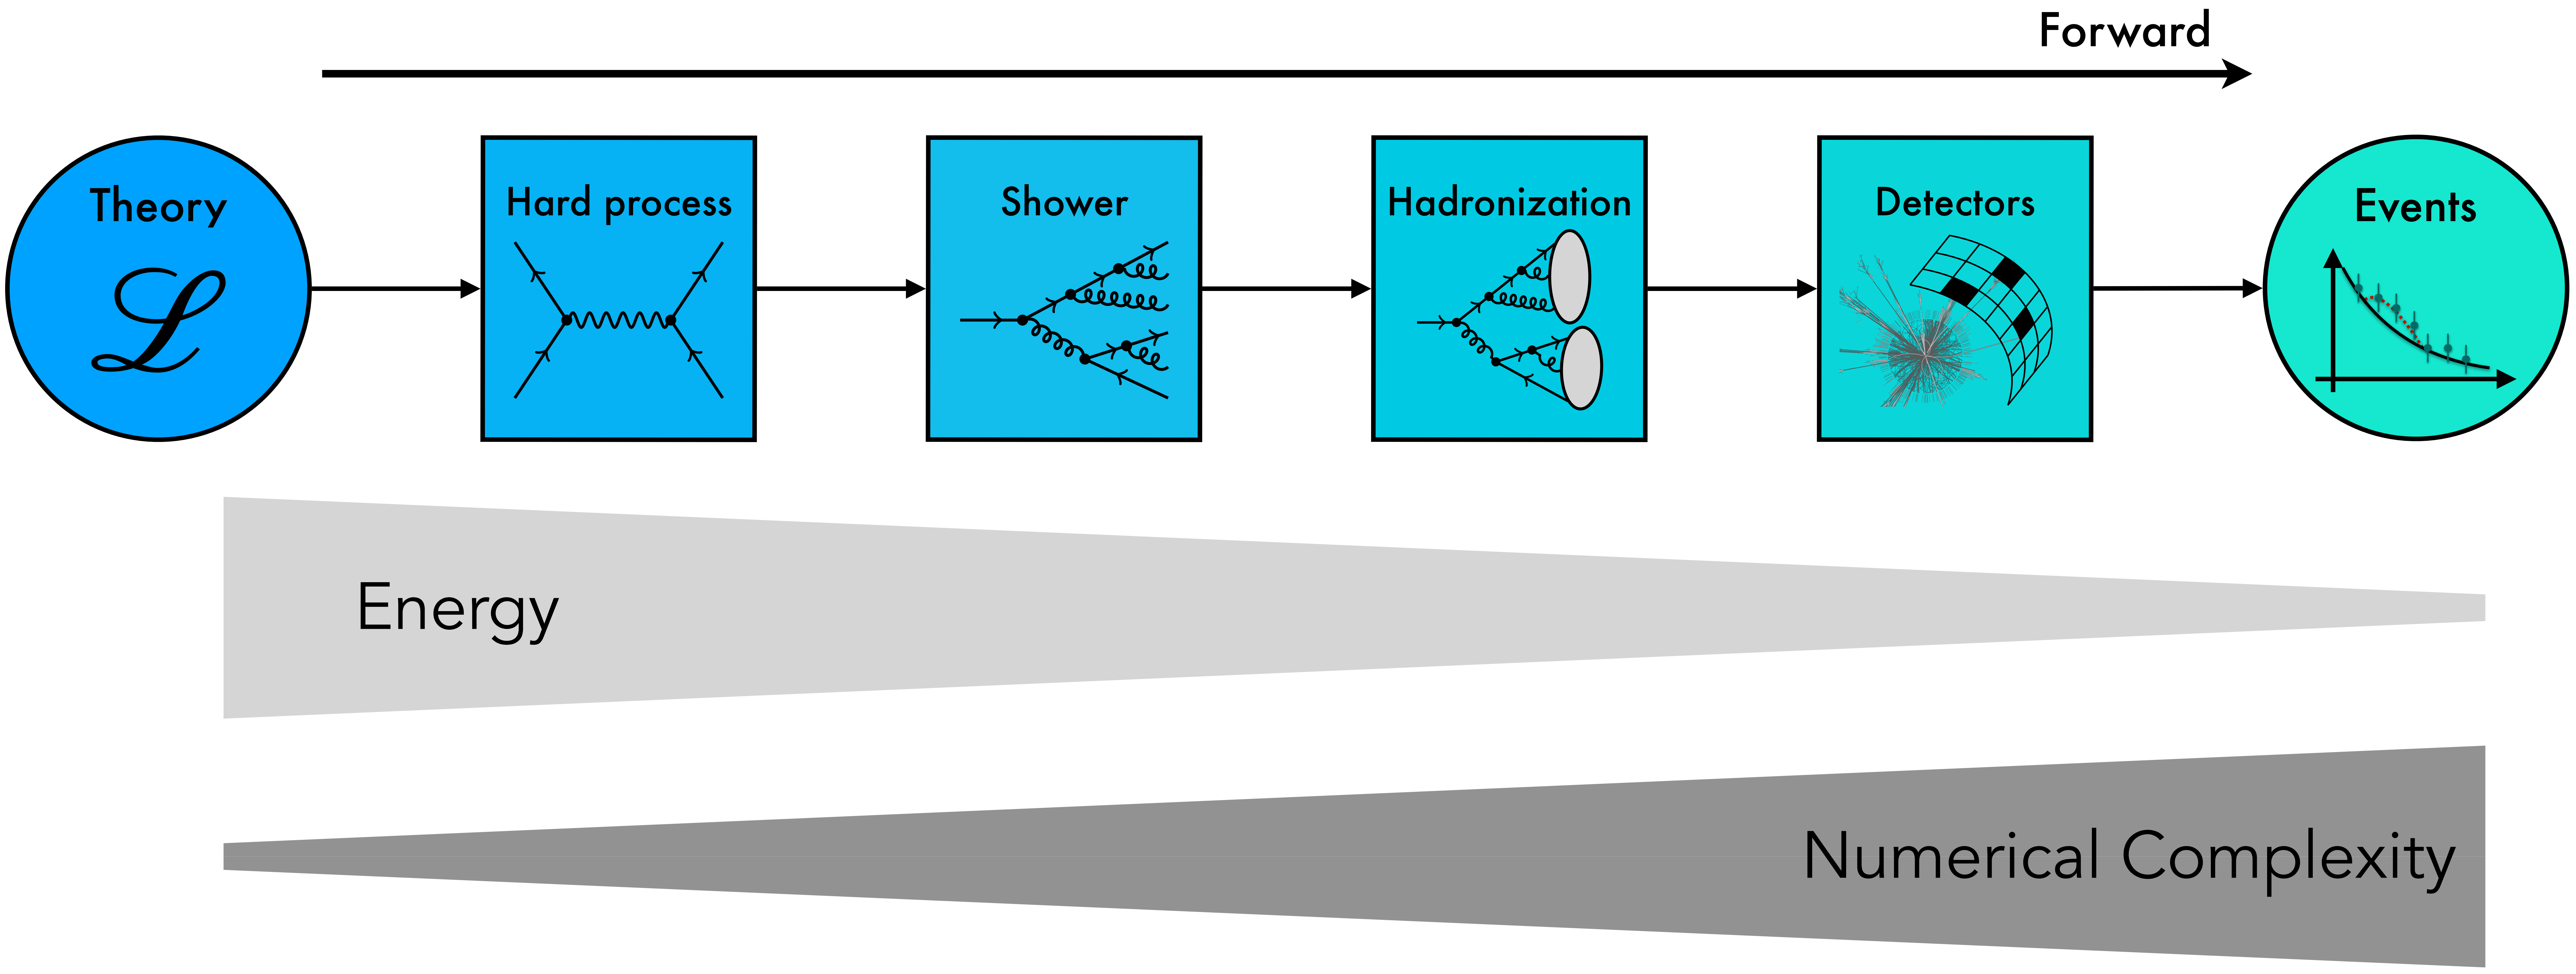

Particle Physics Simulation

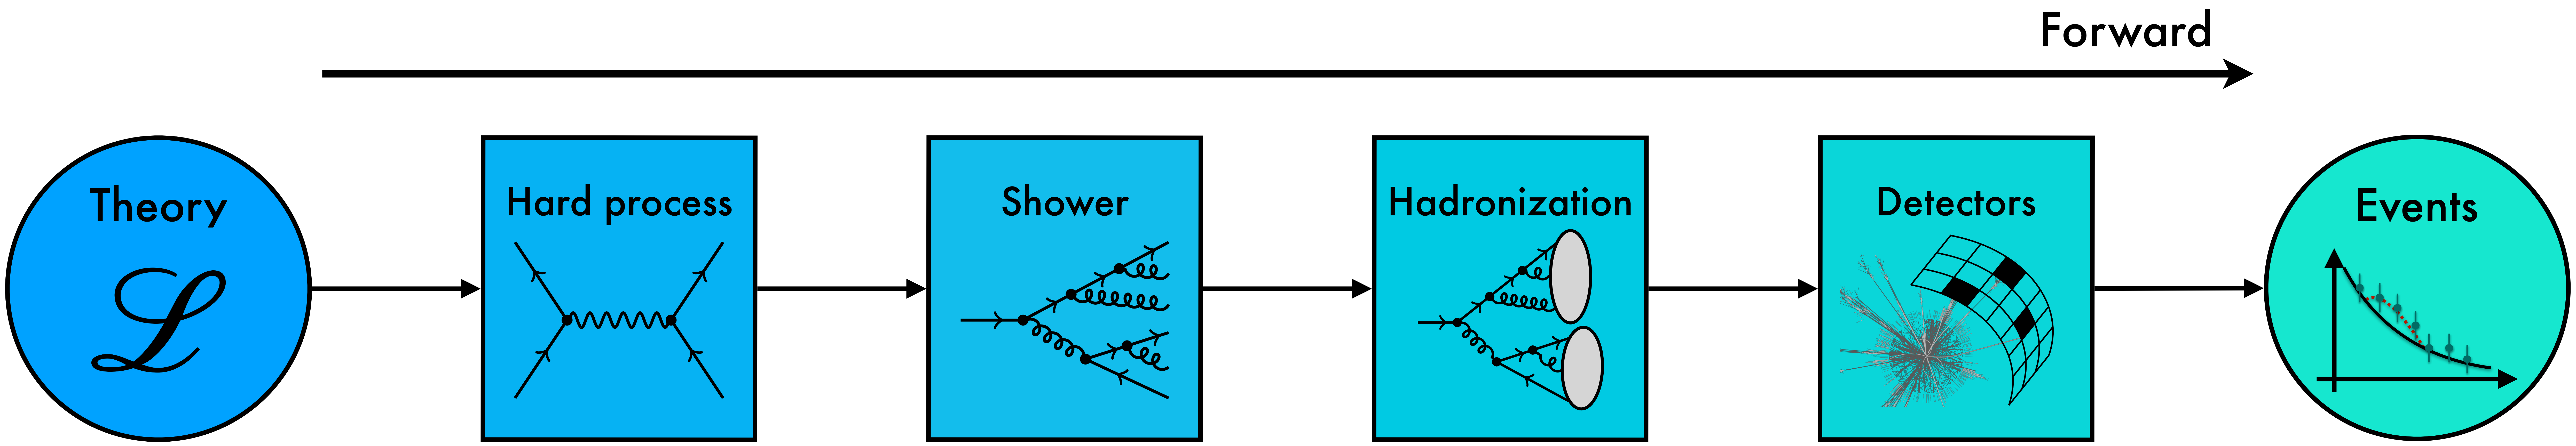

AllShowers

- First-principles precision simulations, accurate, expensive and slow
- Need more simulated data than recorded data
- Higher-rate experiments → need more and more precise simulation

Particle Physics Simulation

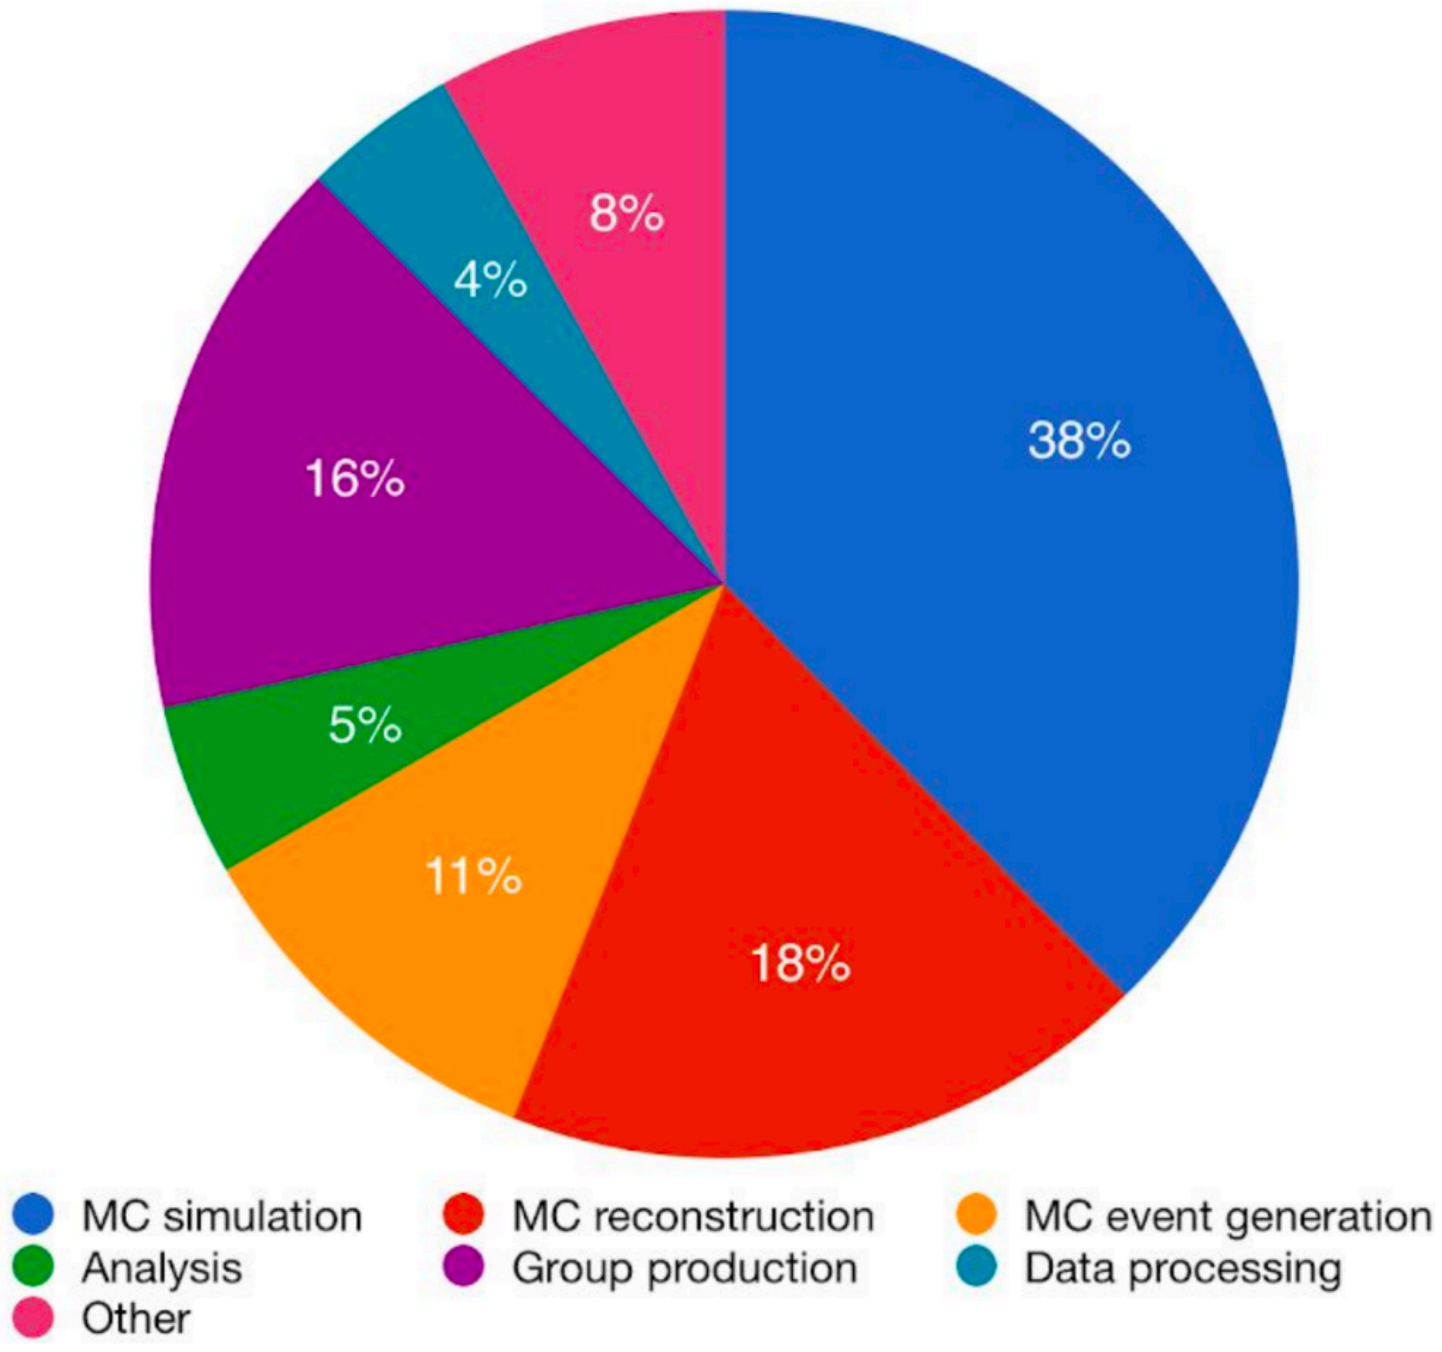

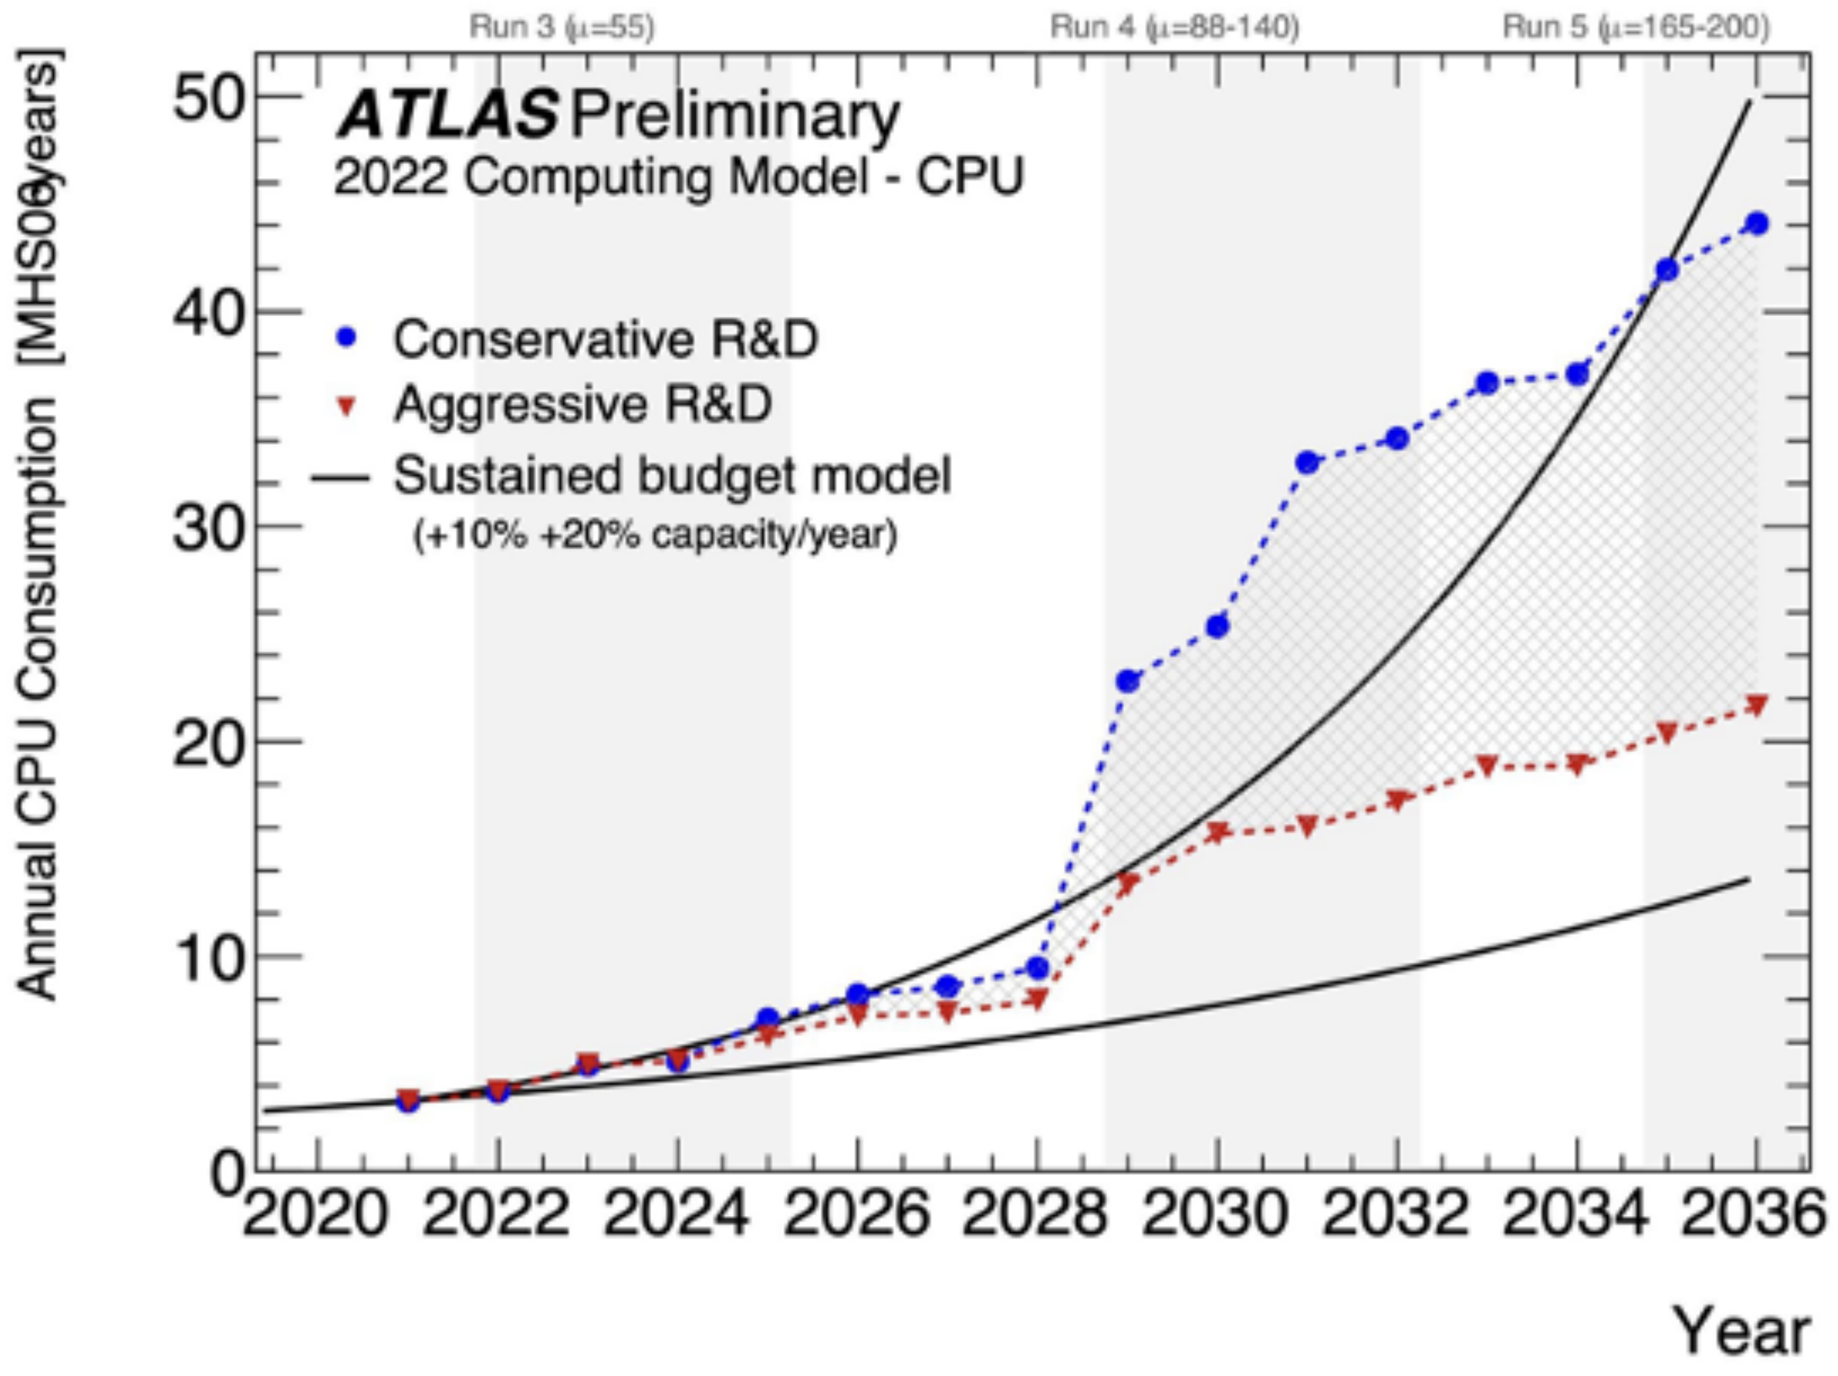

- Simulation cost is computational bottleneck for future HEP experiments
- Need accelerated methods

Al-Turany et. al. JENA Computing Initiative WP2 Report: Software and Heterogeneous Architectures, 2503.09213

Catmore et. al. ATLAS HL-LHC Computing Conceptual Design Report, CERN-LHCC-2020-015 ; LHCC-G-178

Generative Simulation

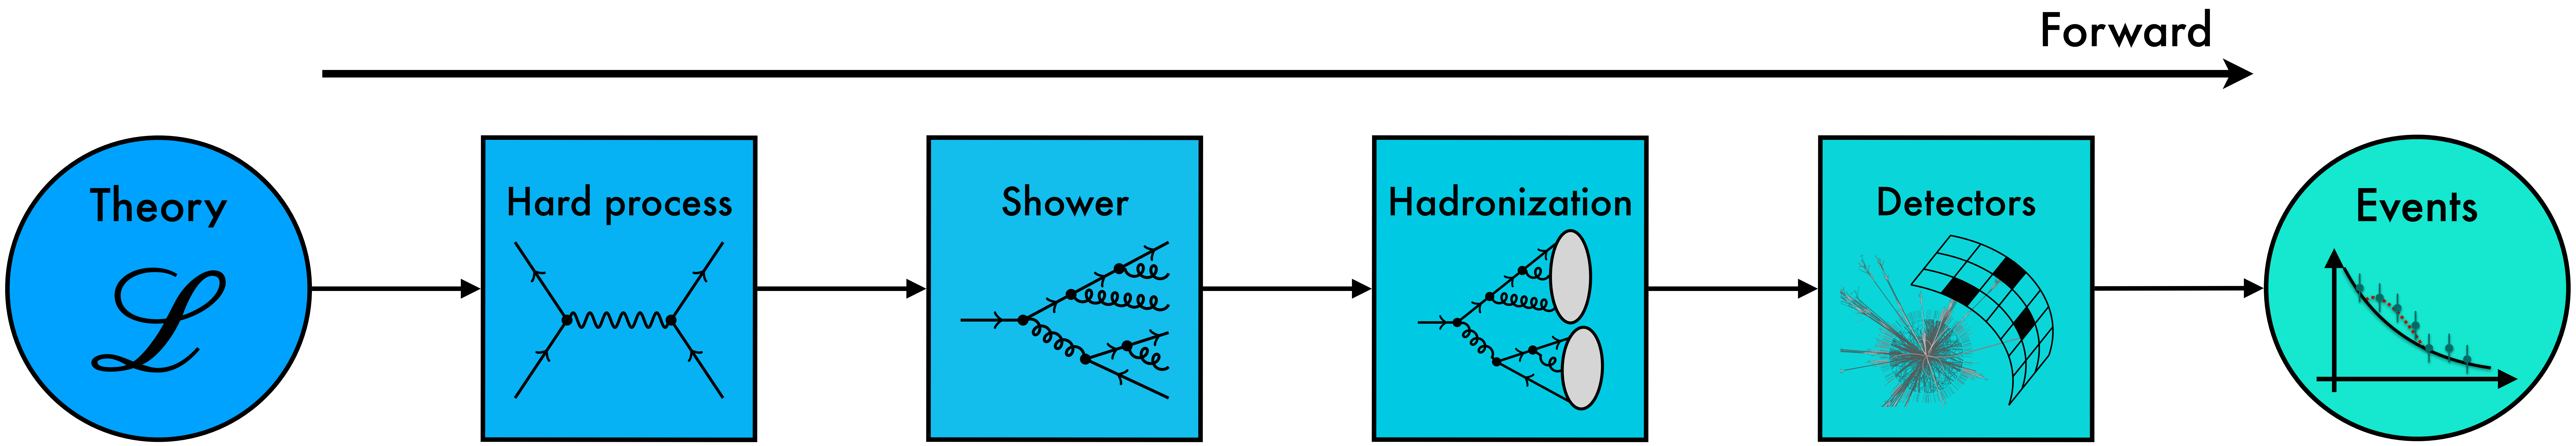

$p_{\text{scatter}}(x|\mathscr{L})$ $\quad p_{\text{shower}}(x|x_{\text{scatter}})$ $\quad p_{\text{hadron}}(x|x_{\text{shower}})$ $\quad p_{\text{detector}}(x|x_{\text{hadron}})$

$\Downarrow$

$p_{\text{gen, scatter}}(x|\mathscr{L})$ $\quad p_{\text{gen, shower}}(x|x_{\text{scatter}})$ $\quad p_{\text{gen, hadron}}(x|x_{\text{shower}})$ $\quad p_{\text{gen, detector}}(x|x_{\text{hadron}})$

- Series of conditional densities, approximate with conditional generative networks
- Faster, more efficient, leverage data

Generative Simulation

Hard Process

- Precise theory description
- Slow evaluation
- Slow sampling

- Initial approach:
  - Train Generative model on simulated data
  - Sample from model

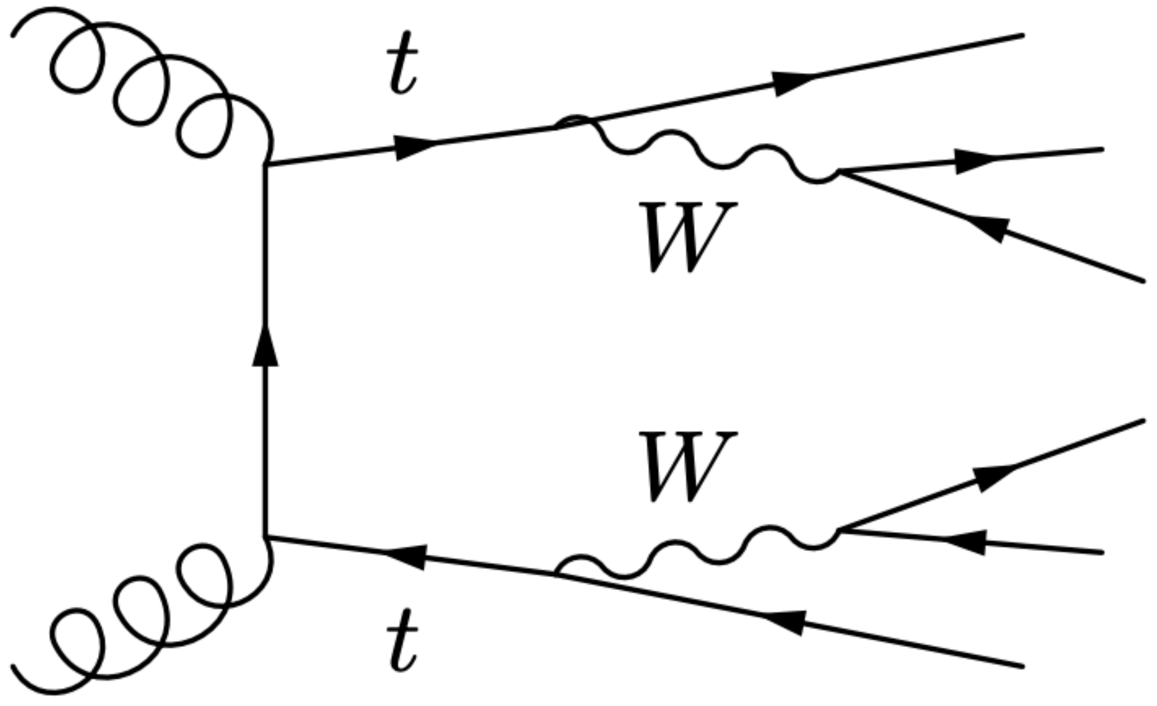

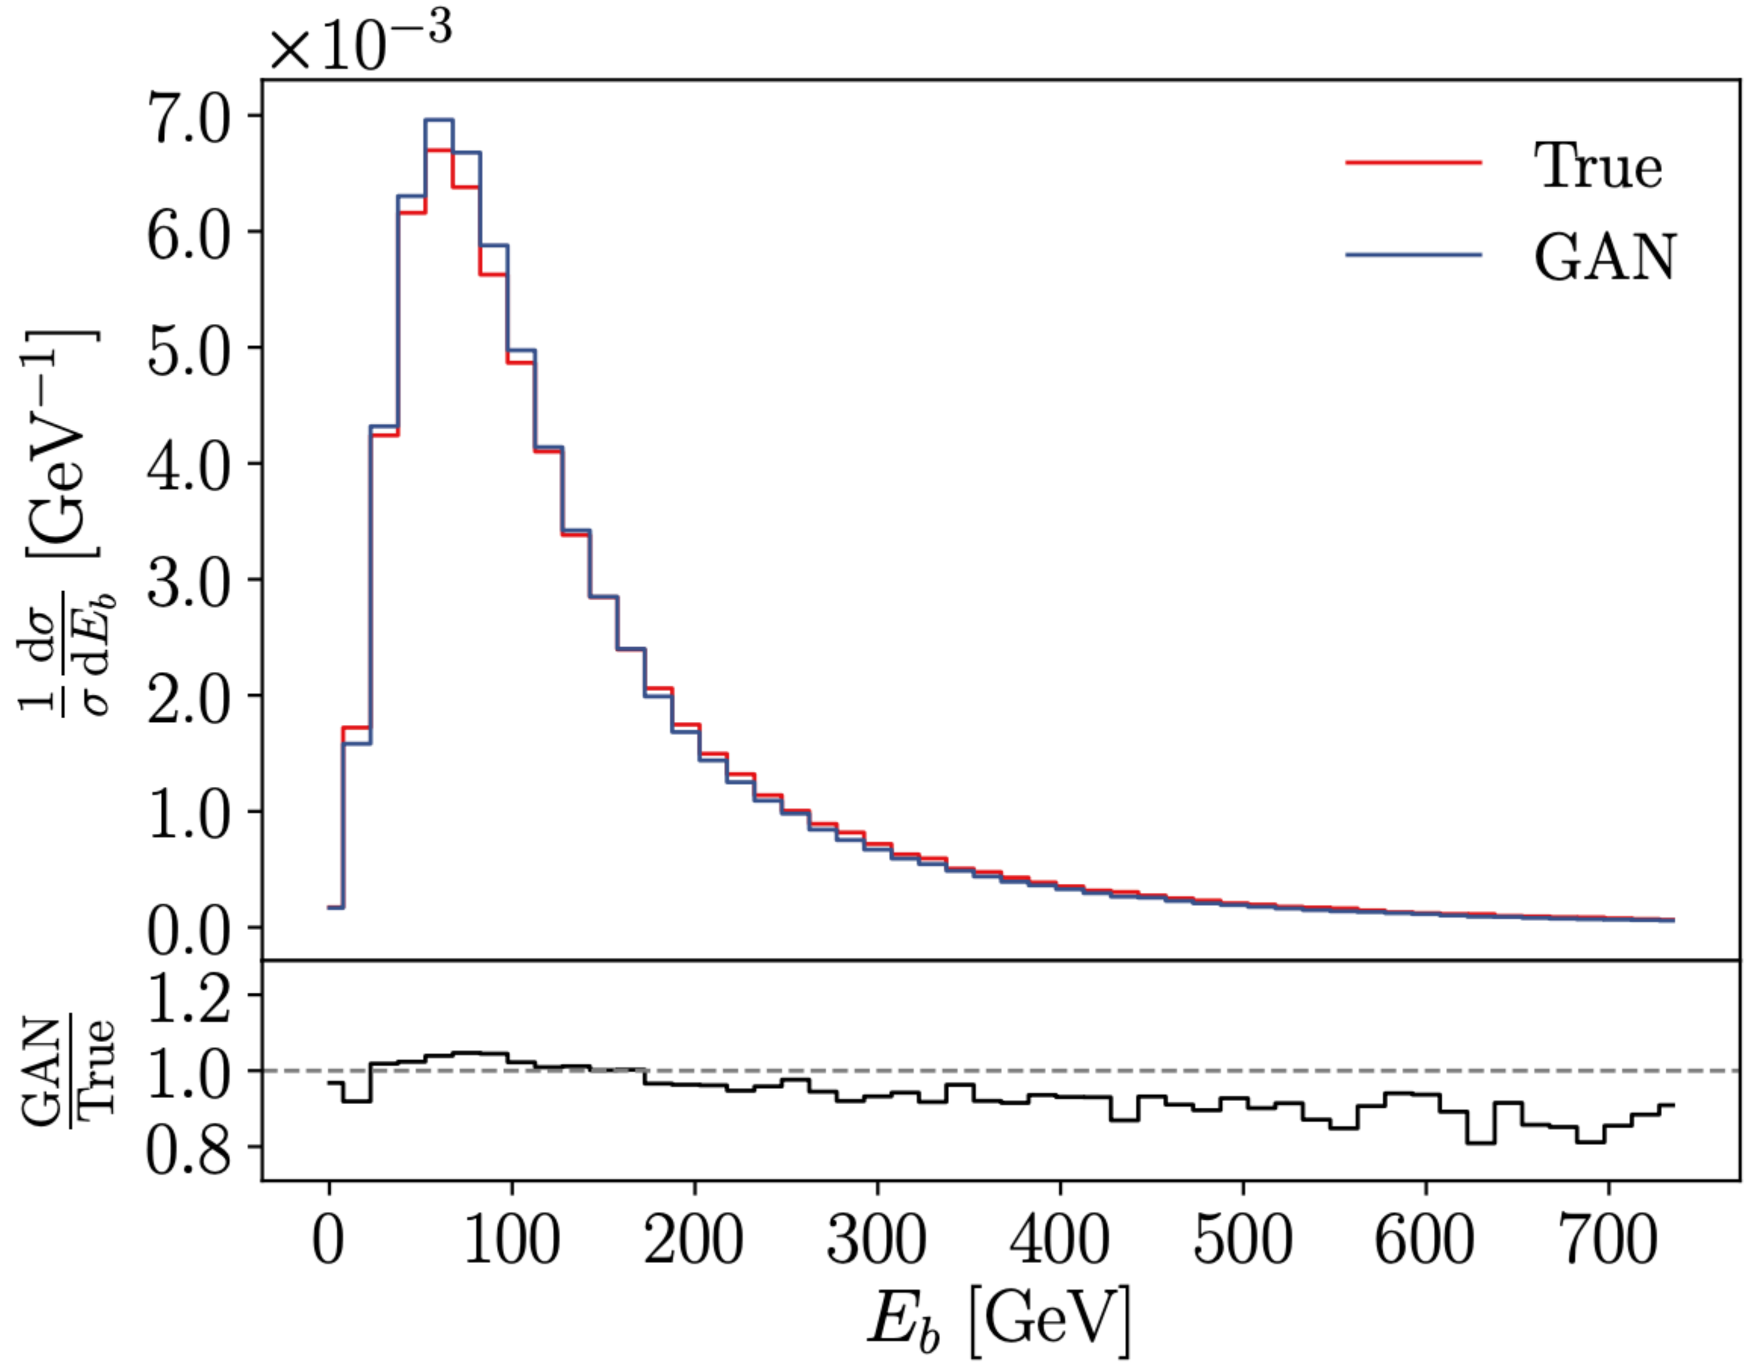

Butter, Plehn, Winterhalder How to GAN LHC Events 1907.03764

Generative Simulation

More Recent:

- Slow evaluation
  → Amplitude surrogates
- Train regression model on scattering amplitudes
- Only needs grid of points
- Evaluate model
  → Massive acceleration

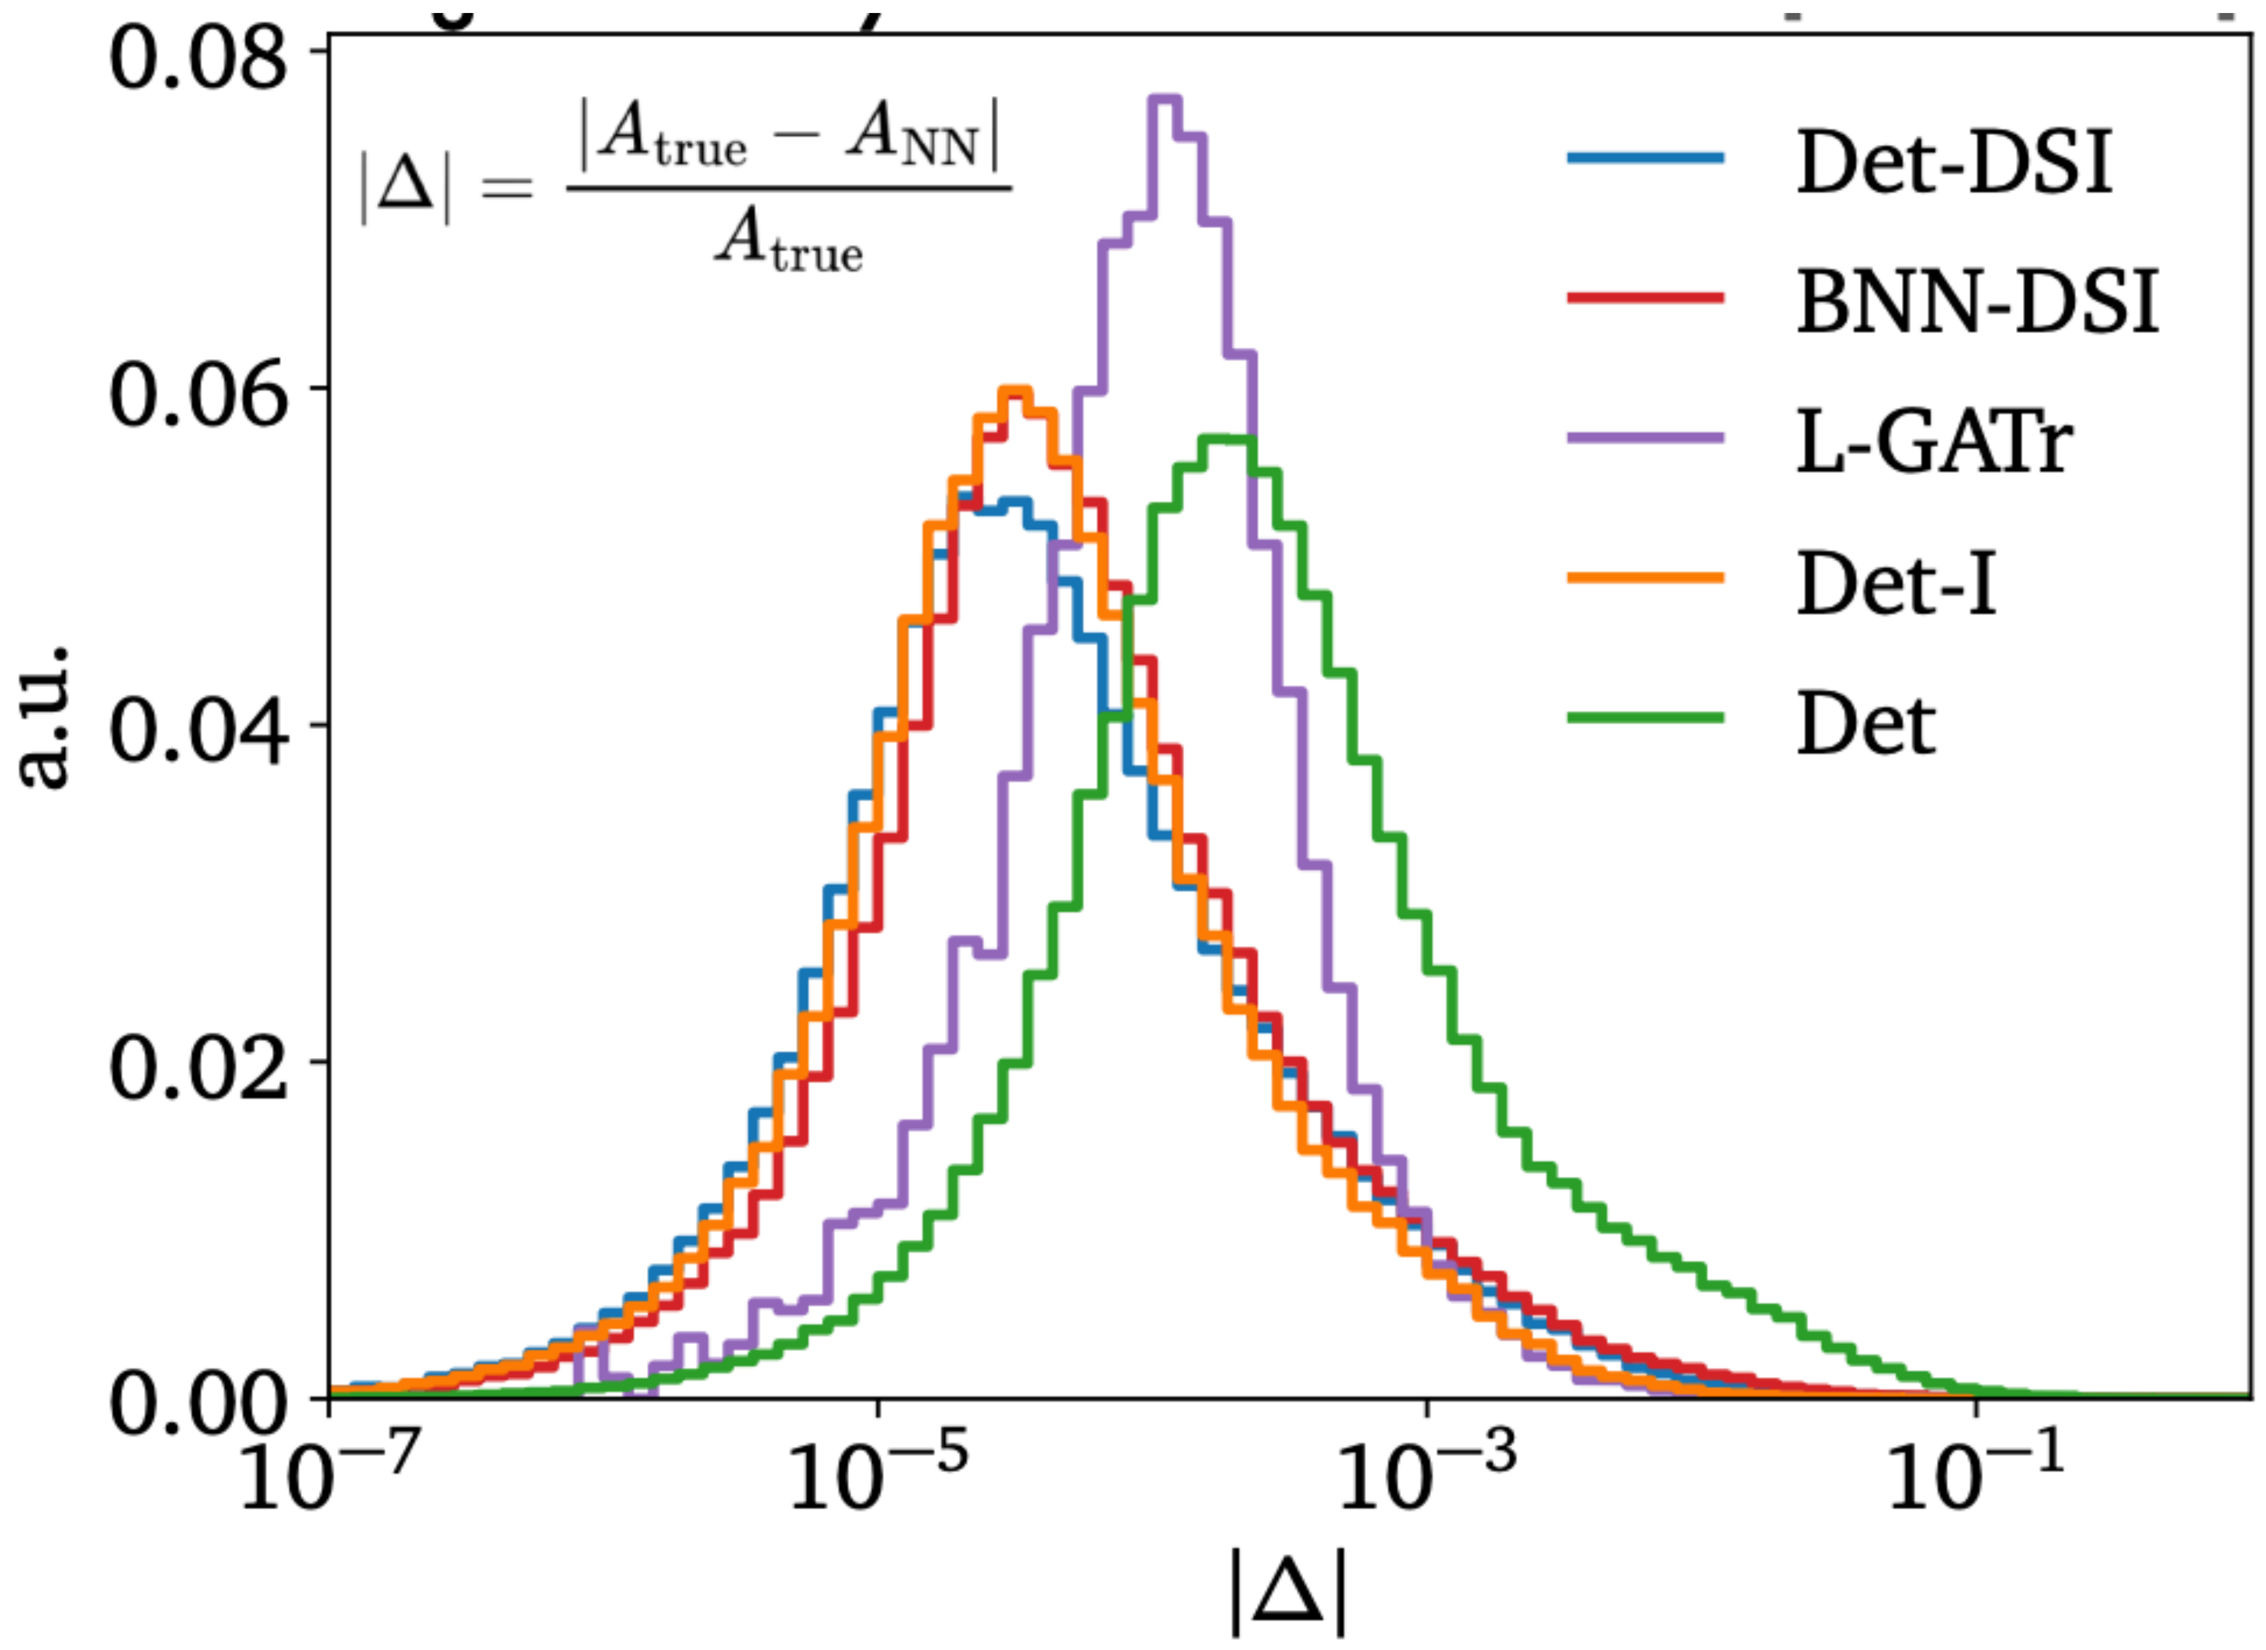

Winterhalder: Theory ex Machina EuCAIFCon 2026

Bahl et. al.: Accurate Surrogate Amplitudes with Calibrated Uncertainties 2412.12069

Generative Simulation

More Recent:

- Slow sampling

  → Neural Importance Sampling

- Integrate over phase space
- Needs numerical integration
- Speed up by sampling proposal form Normalizing flow
- Both surrogates and NIS already deployed in MadGraph

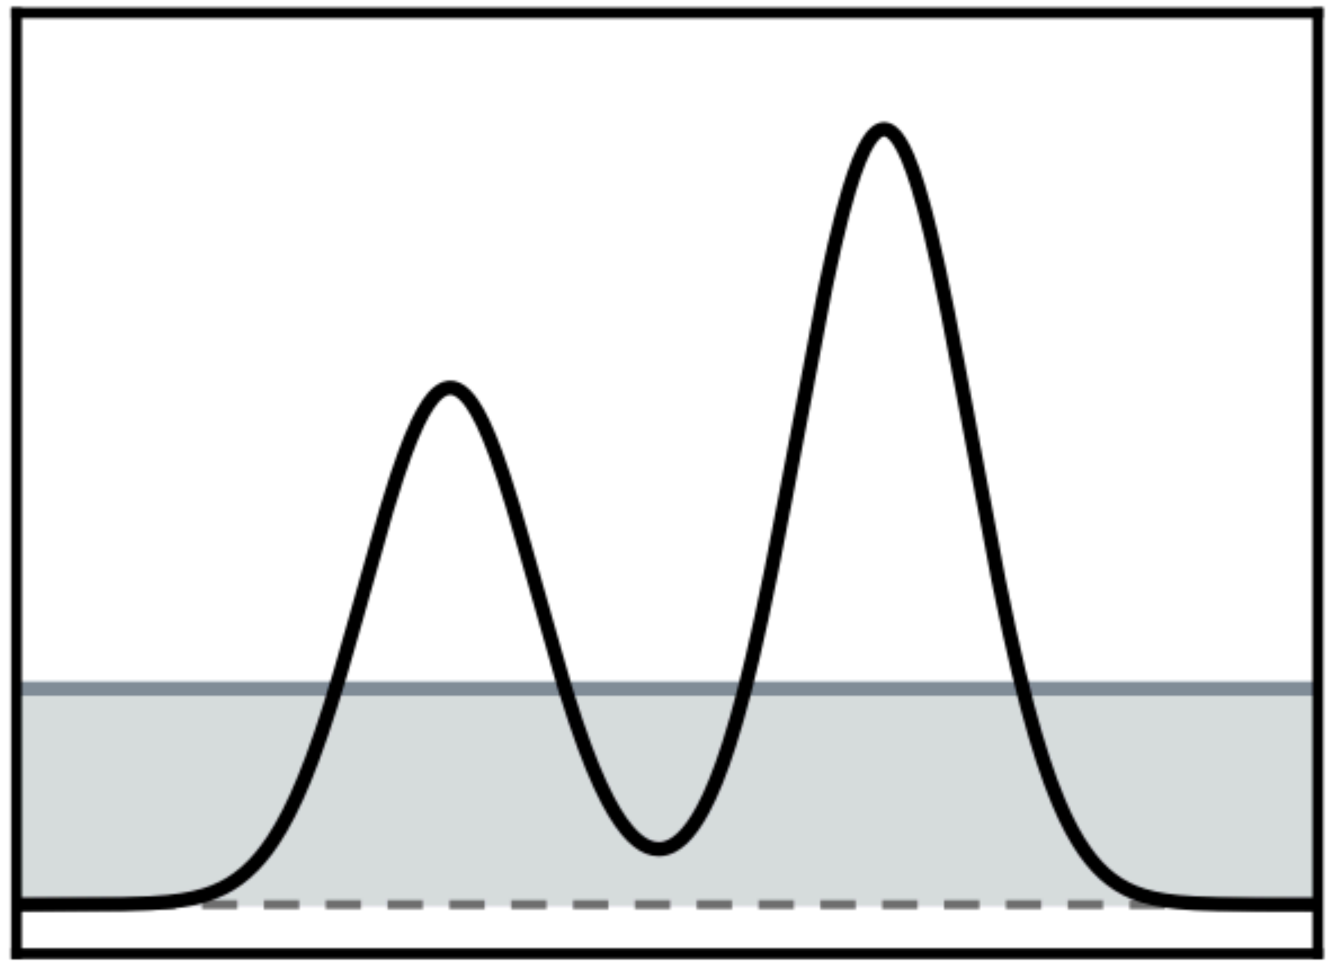

Flat sampling: inefficient

$$I = \langle f(x) \rangle_{x \sim \text{unif}}$$

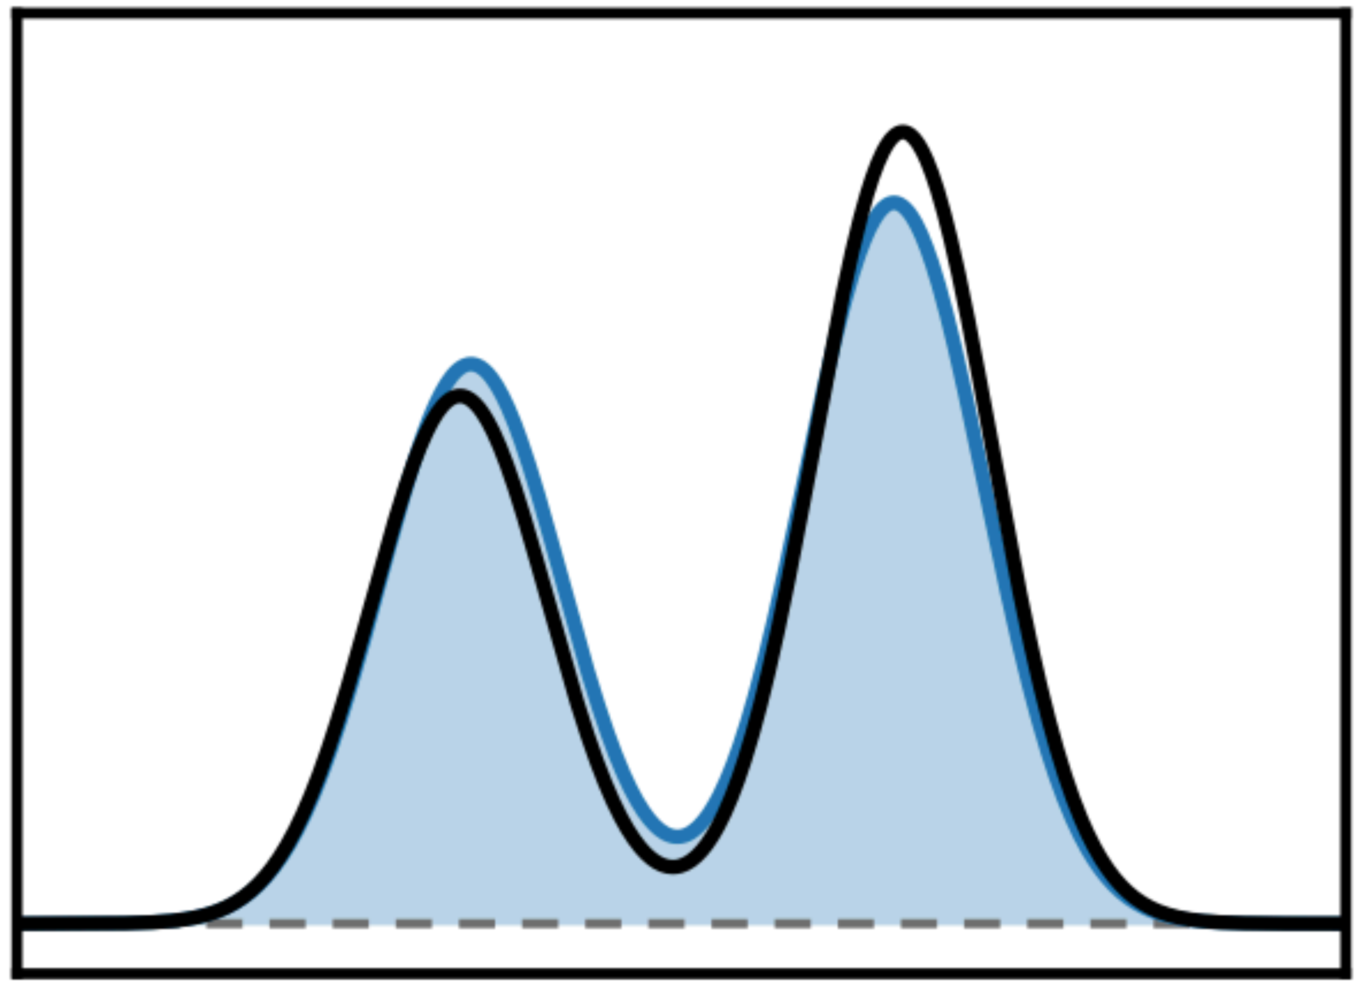

Importance sampling: find p close to w

$$I = \left\langle \frac{f(x)}{g(x)} \right\rangle_{x \sim g(x)}$$

Winterhalder: Theory ex Machina EuCAIFCon 2026

Heimel et. al.: MadNIS -- Neural Multi-Channel Importance Sampling 2212.06172

Generative Simulation

Calorimeter Simulation

- Stochastic understanding
- Very slow individual modeling
- High dimensional data

- Initial approach:
  - Skip all intermediate steps
  - Train model on end result

  → Learn full calorimeter shower

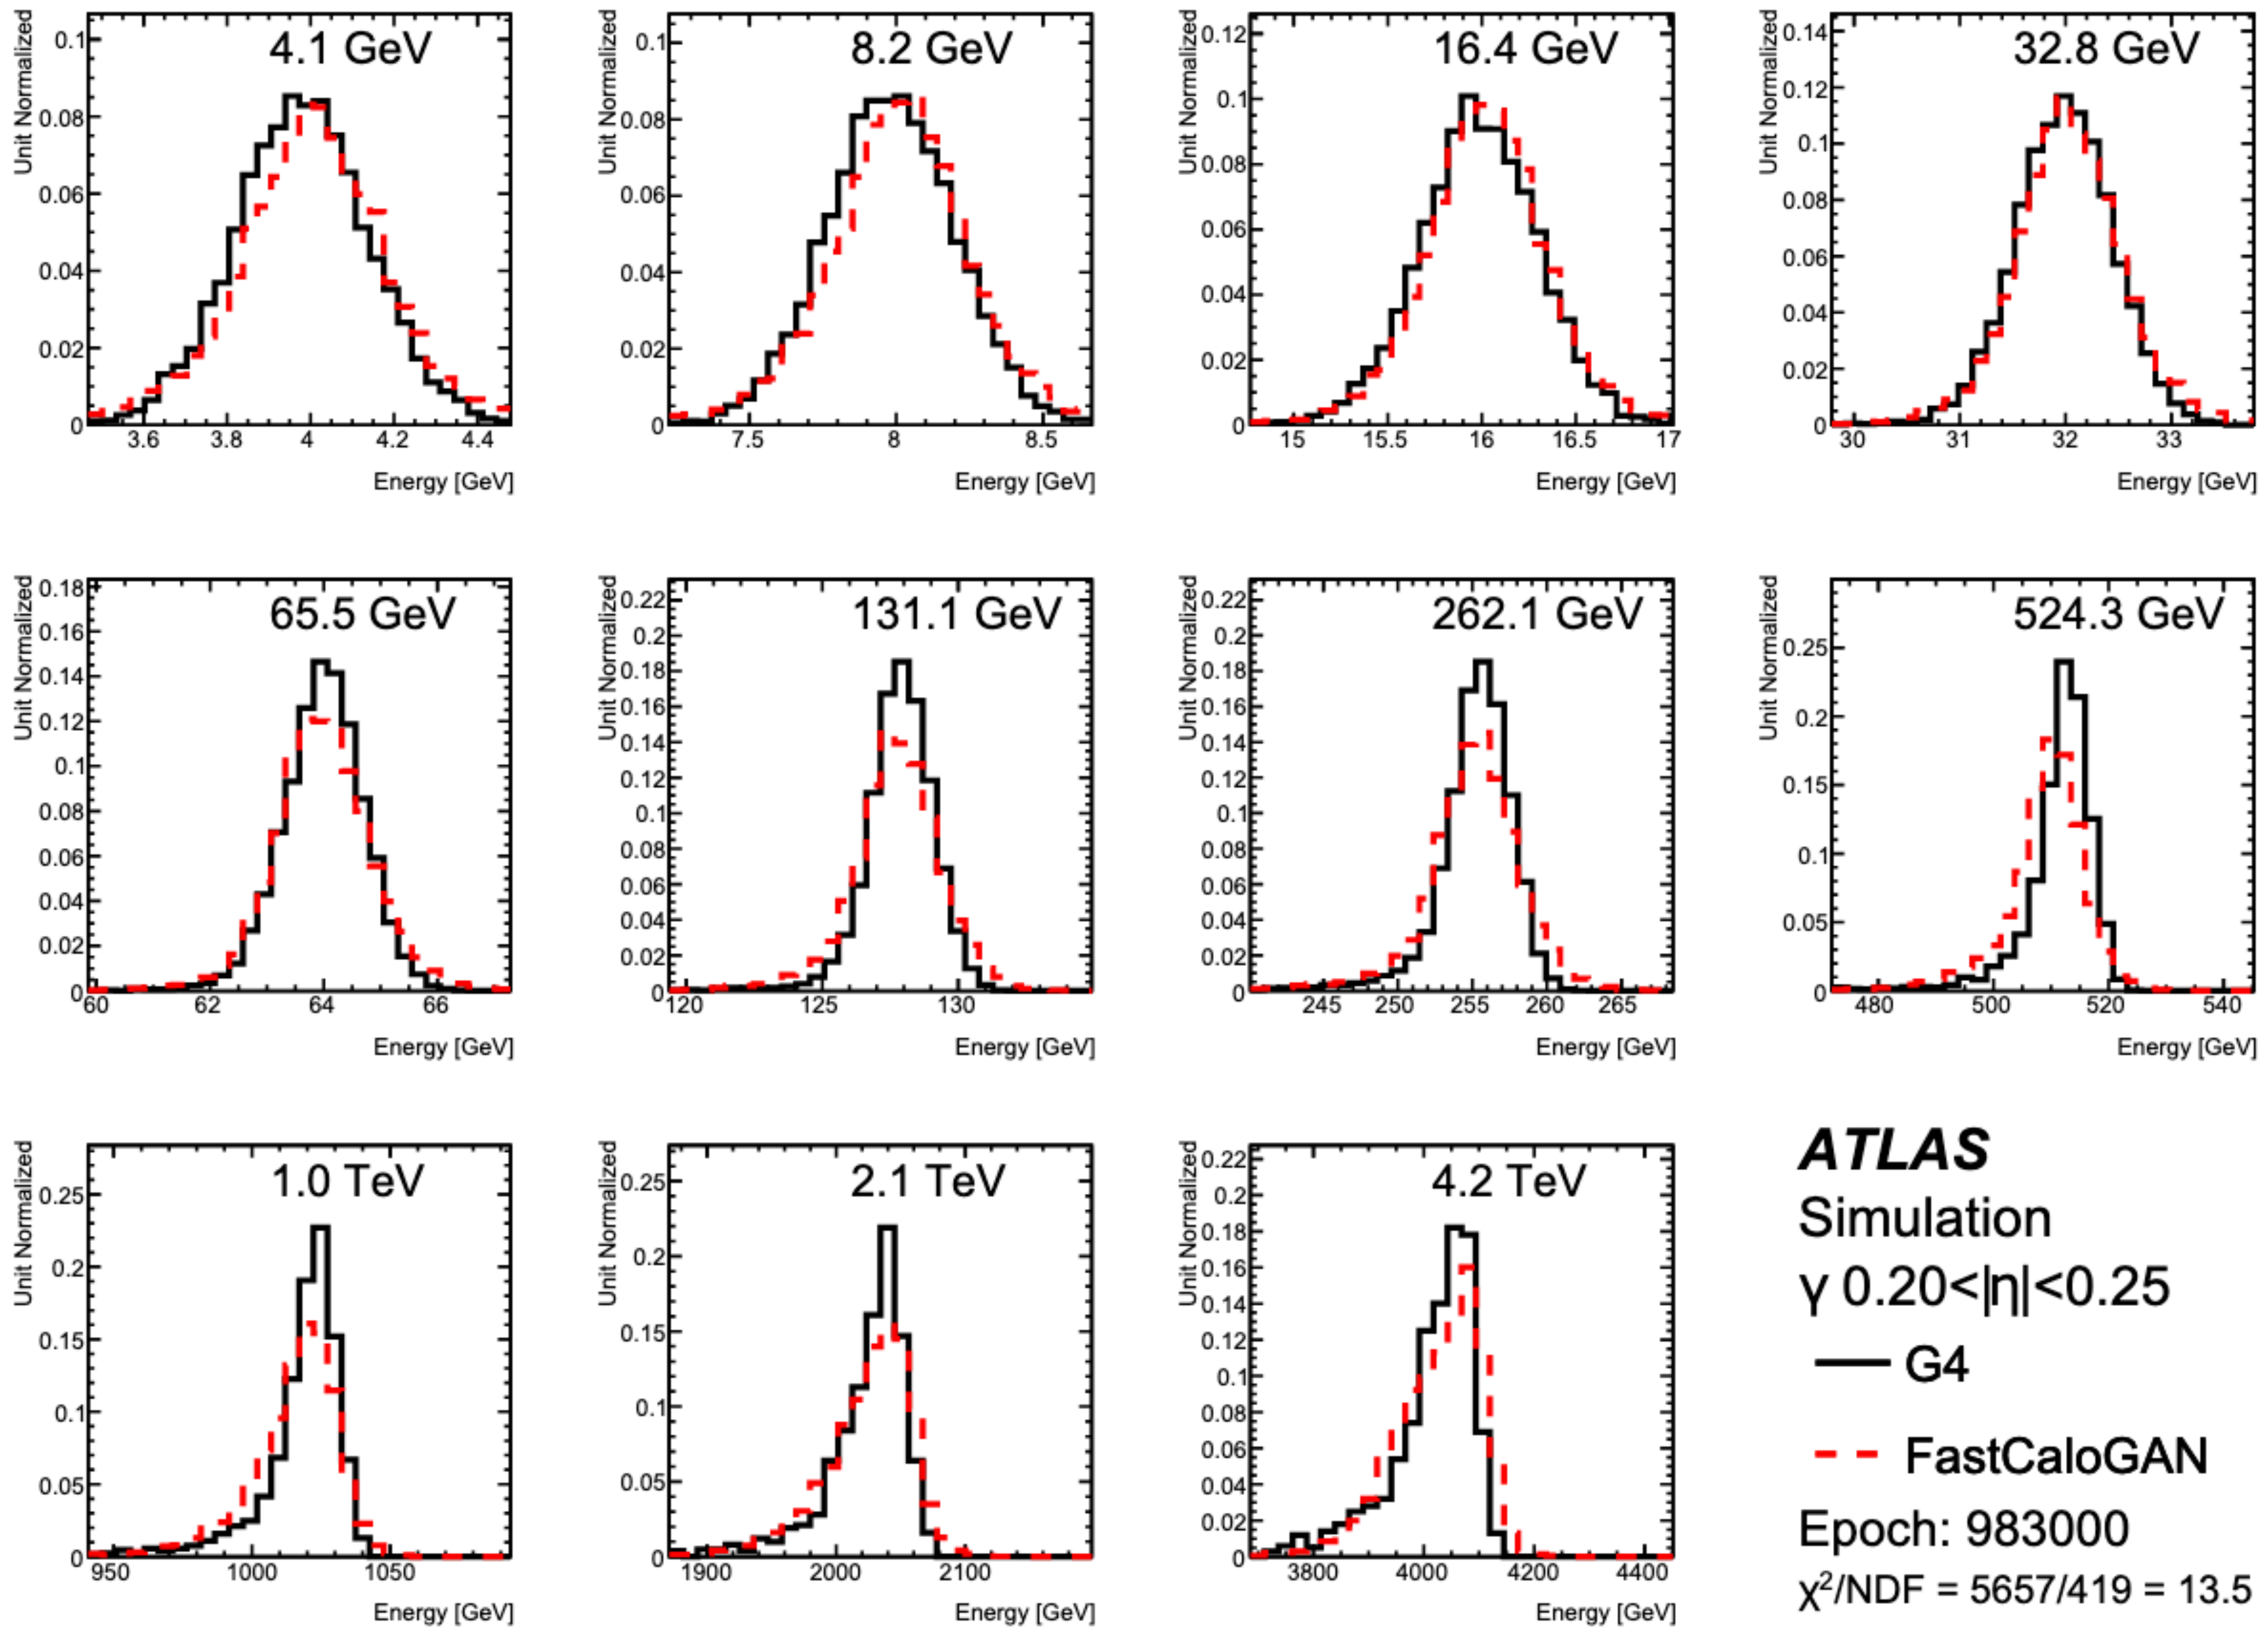

ATLAS Collaboration, AtlFast3: the next generation of fast simulation in ATLAS 2109.02551

Calorimeter Simulation

- Problem: detector are big, shower are localized

  → Learn only region around shower

- Next Problem: shower position is not constant

  → Pointcloud model, project back to detector geometry

  Universal for position and angle

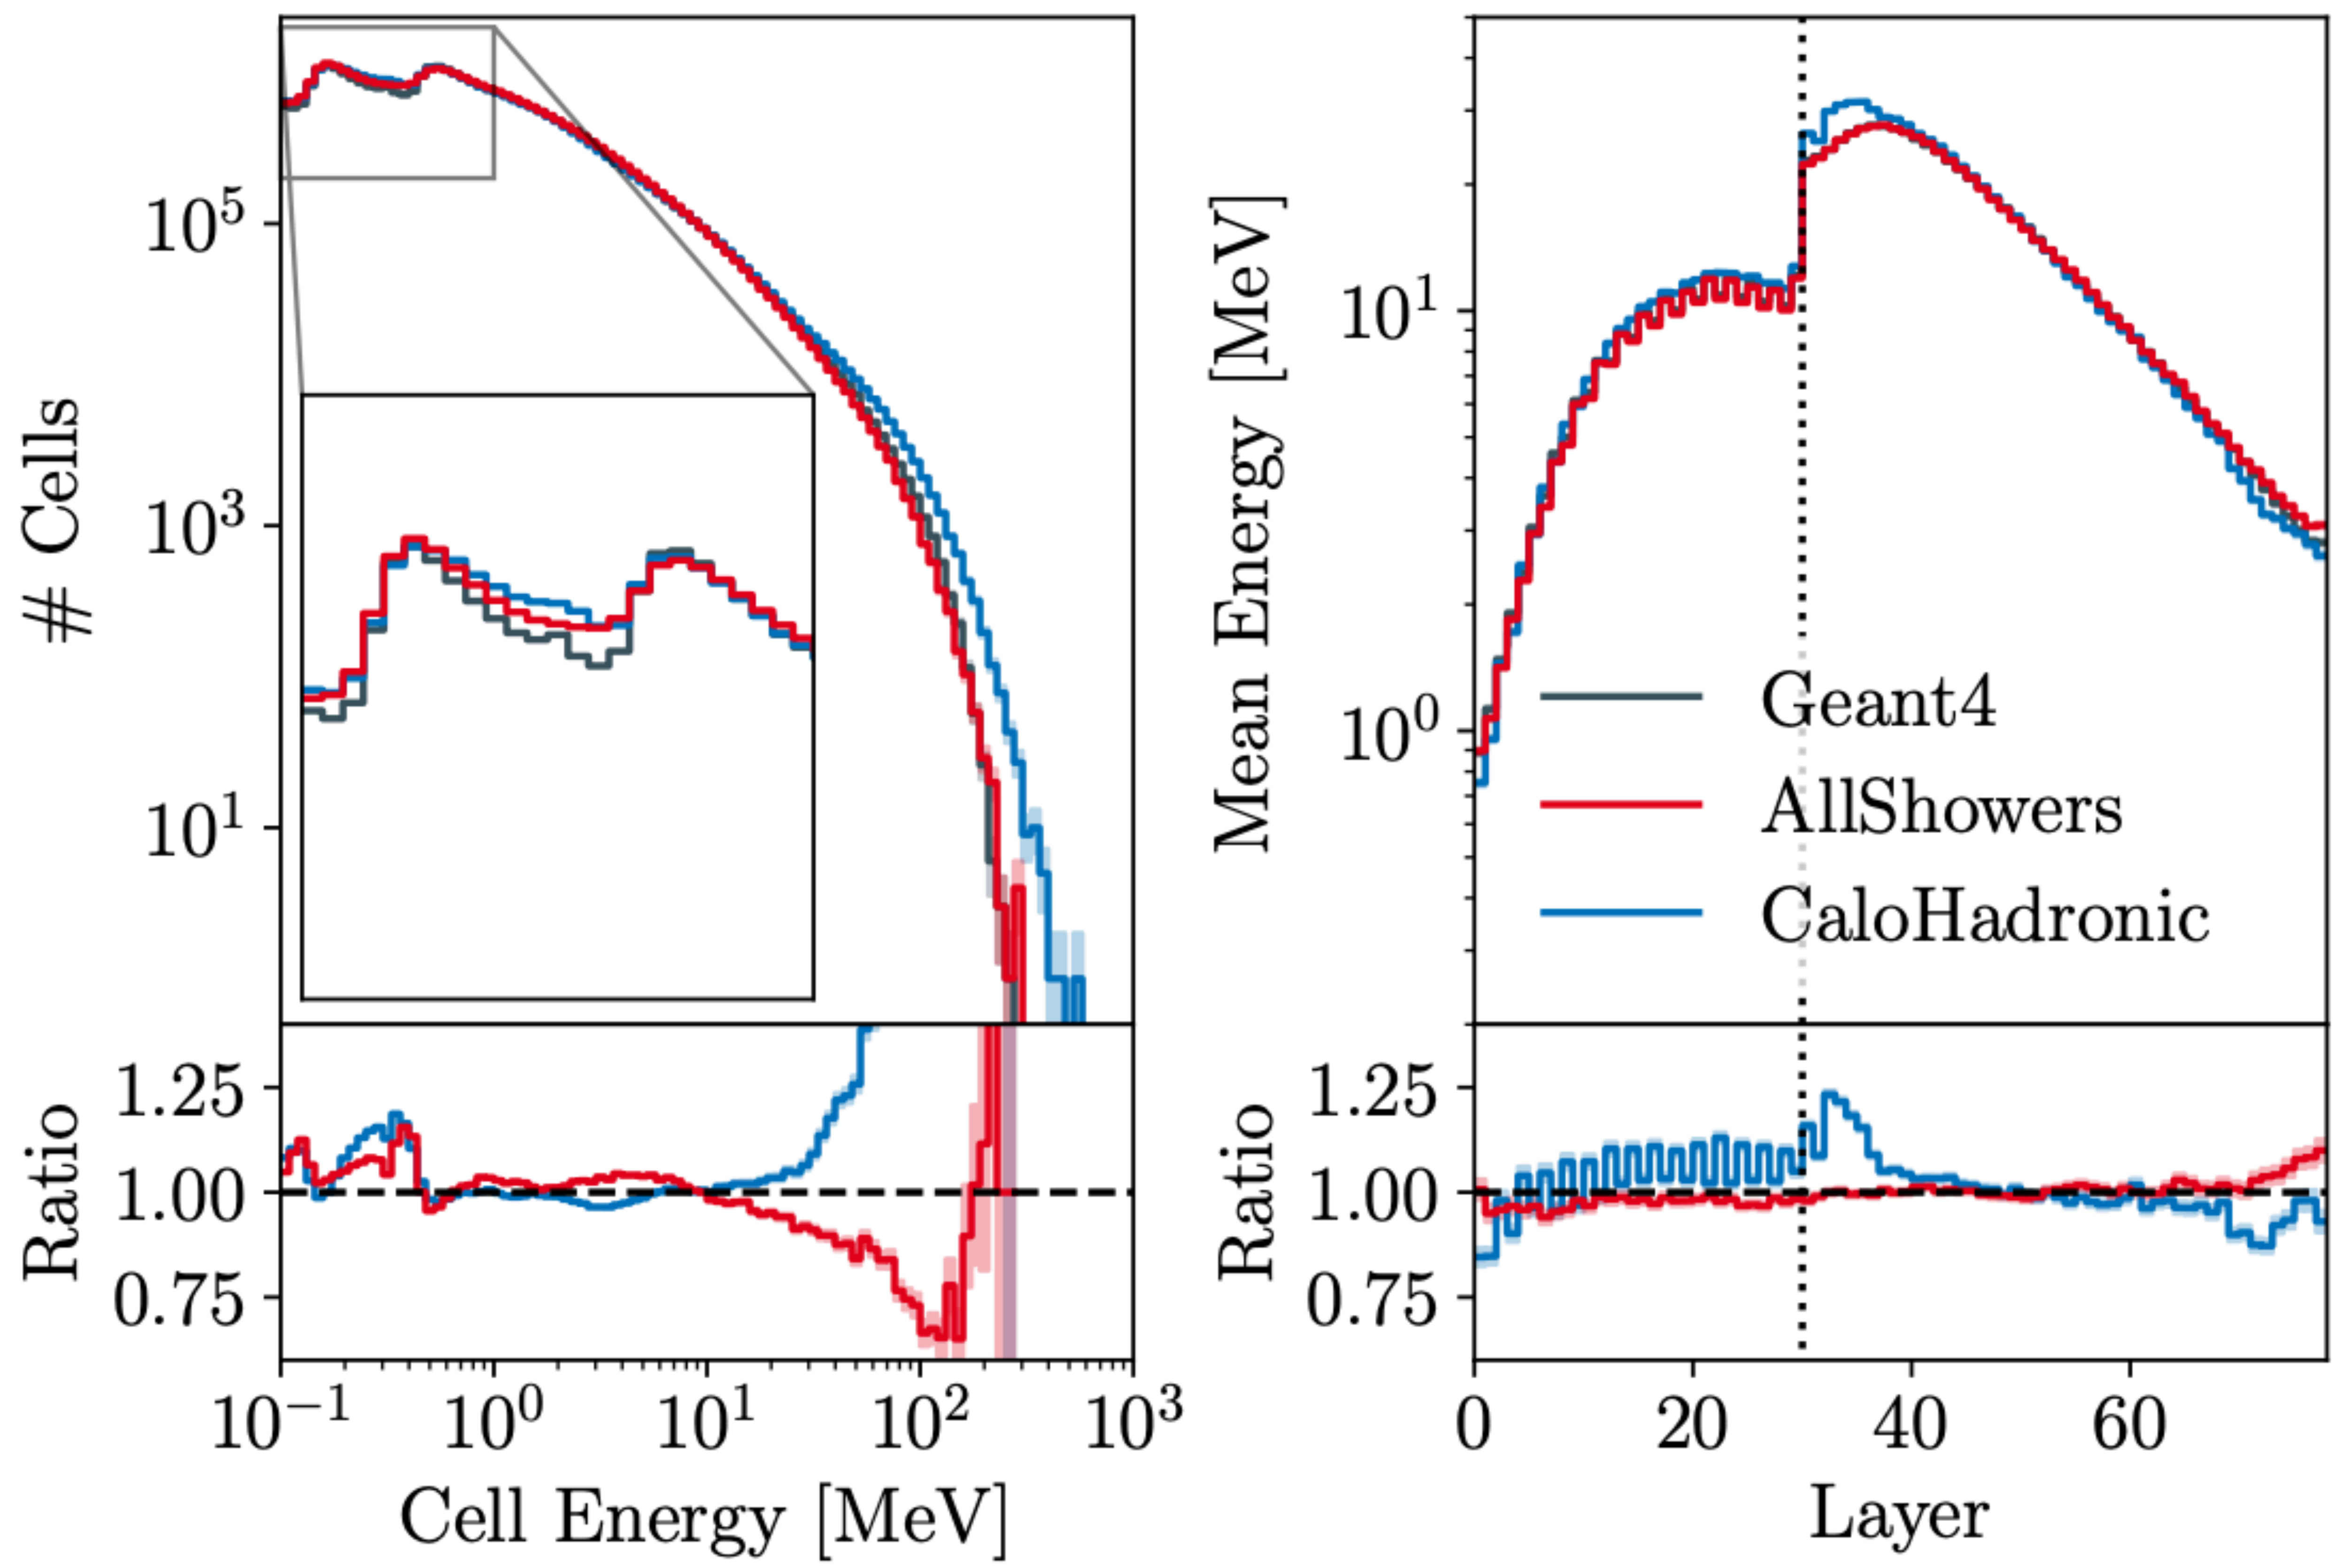

Buss et al.:AllShowers: One model for all calorimeter showers 2601.11716

Generative Simulation

Calorimeter Simulation

- Highly active field of research

- Massive comparison challenge

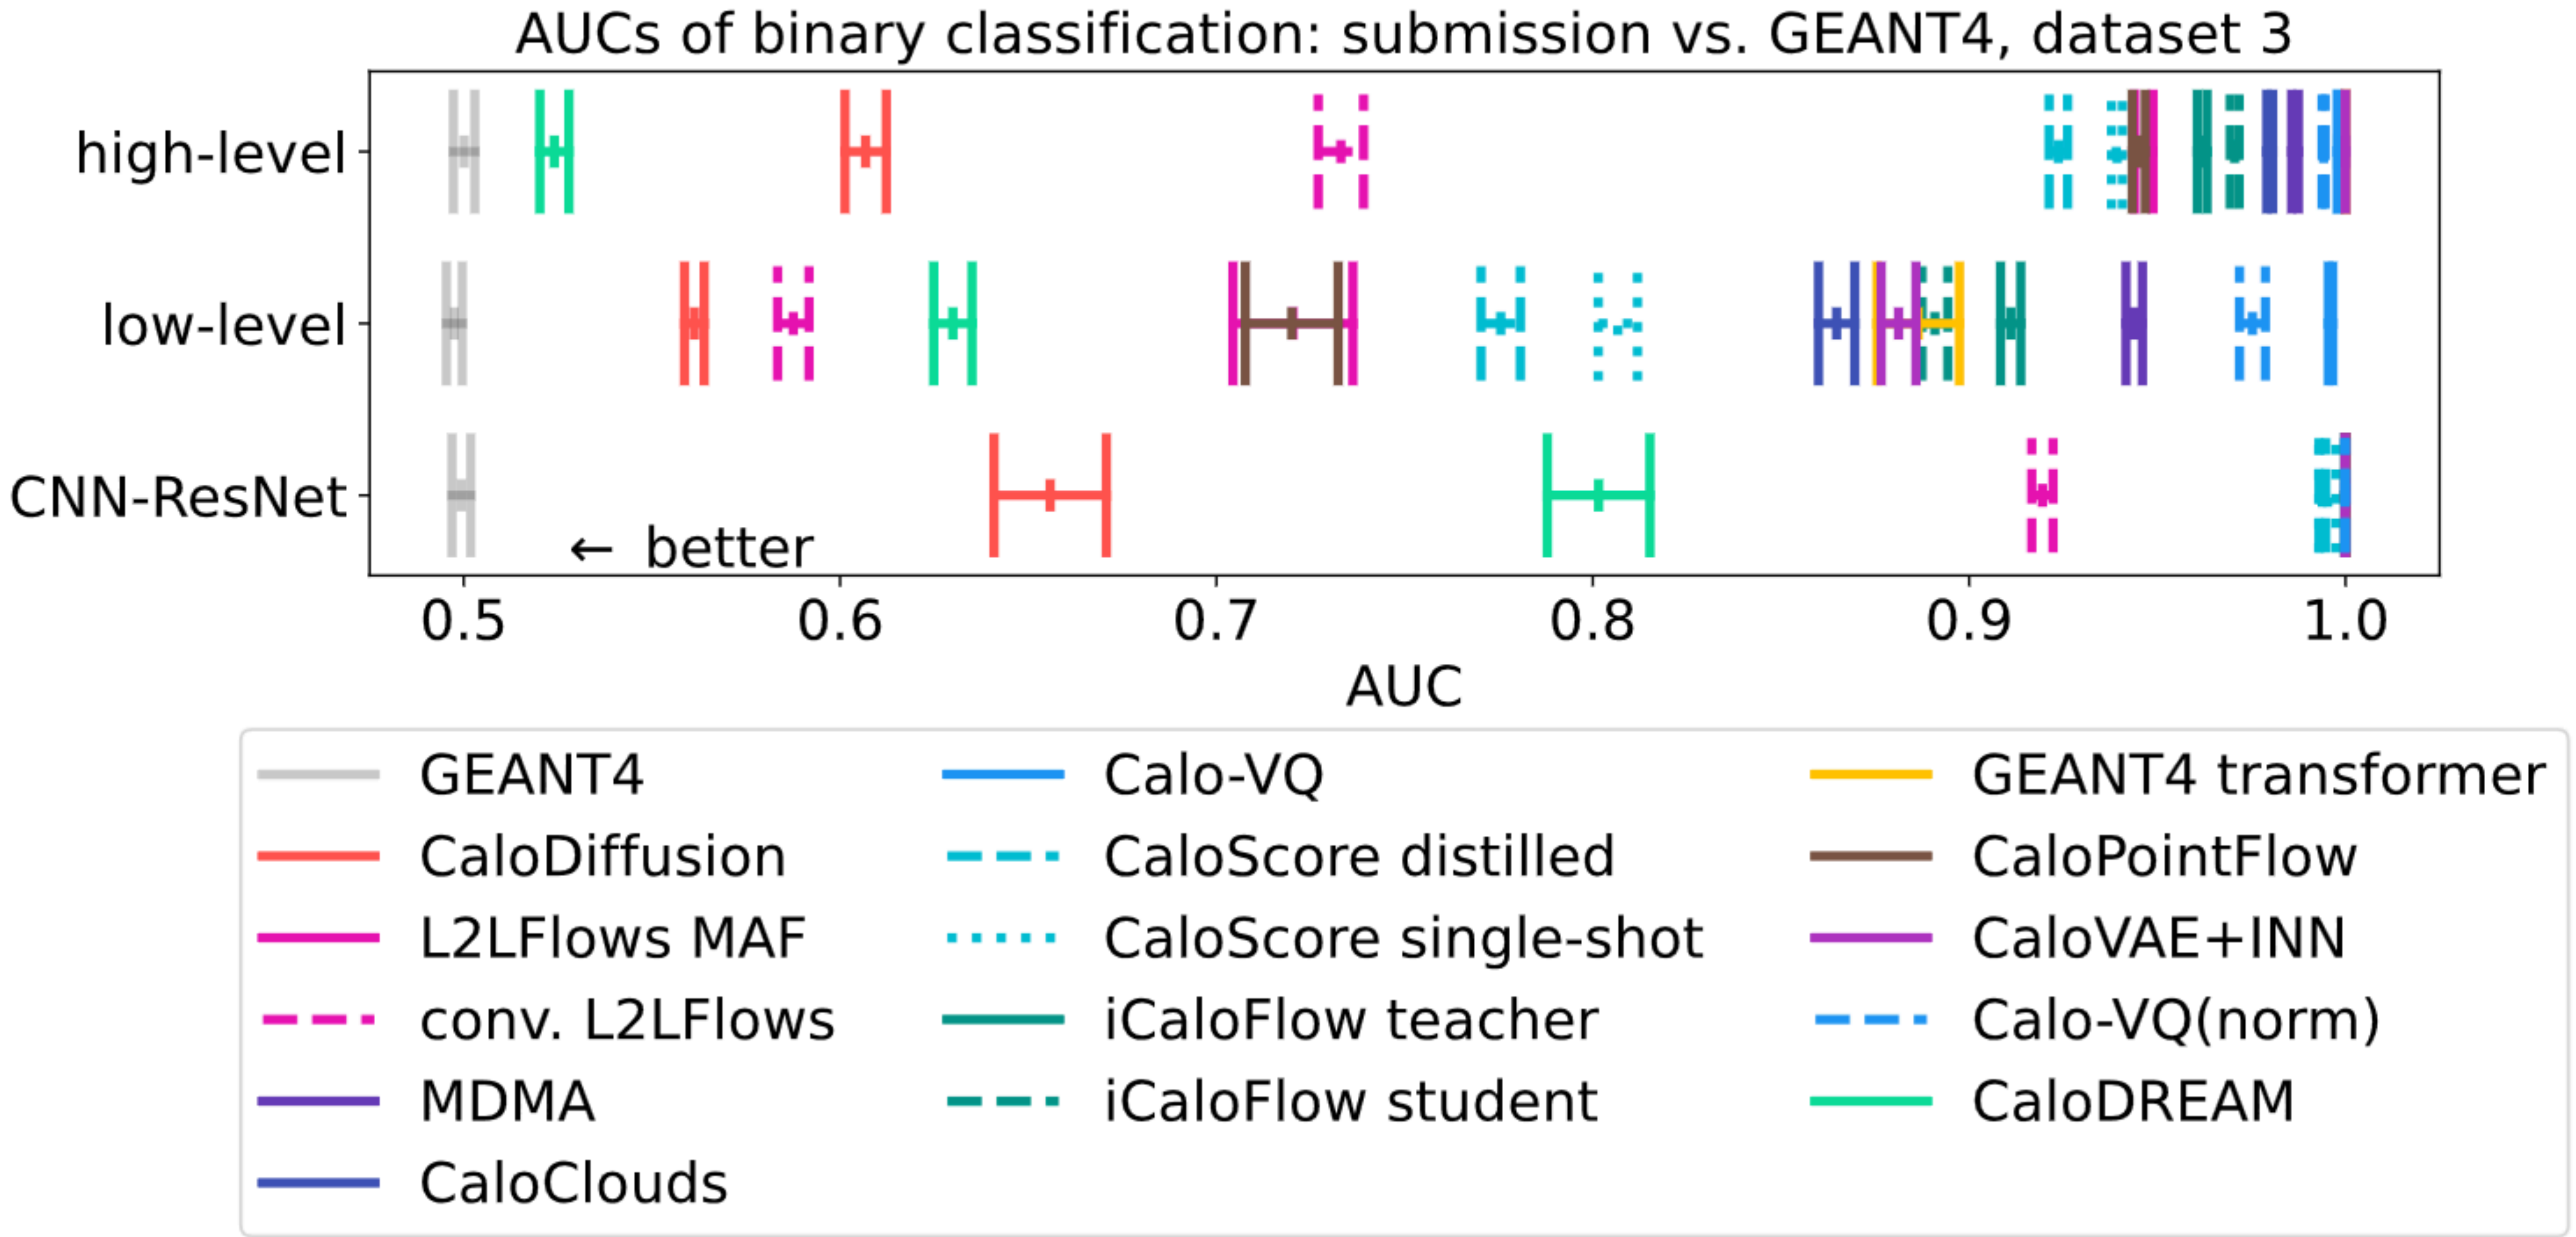

Krause et al.:CaloChallenge 2022: a community challenge for fast calorimeter simulation 2410.21611

Generative Simulation

Fast simulation

- Train on simulated data
- Sample from generative network
- Faster than classical simulation
  - GAN
  - Normalizing flow
  - Diffusion models
  - Autoregressive transformers

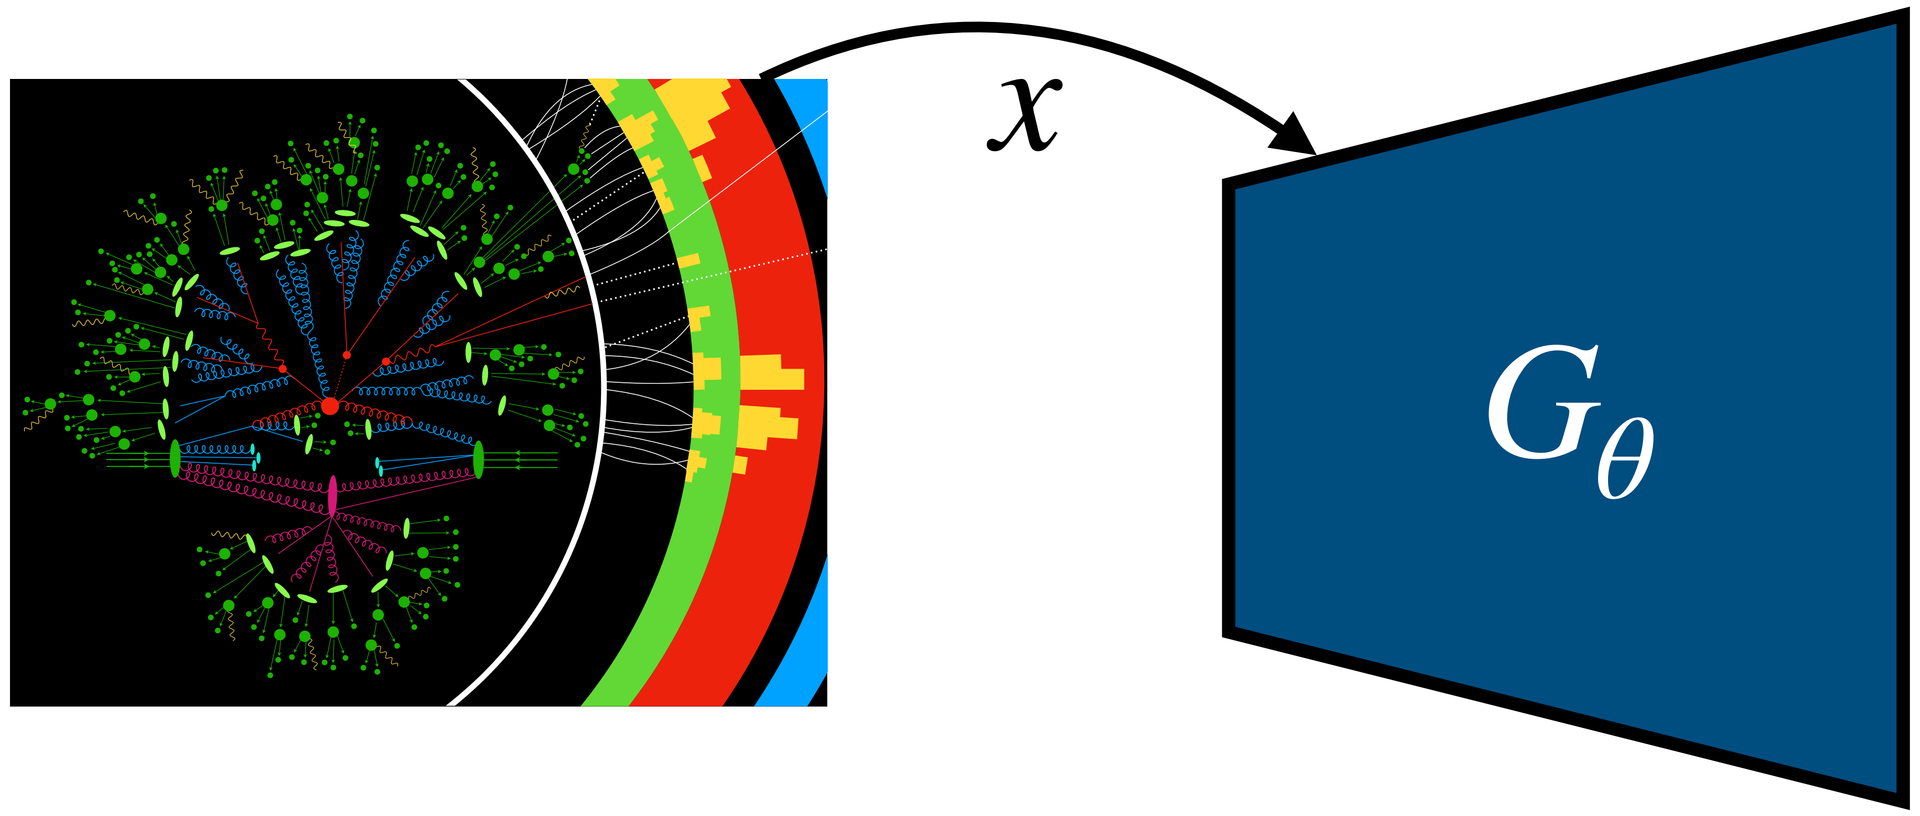

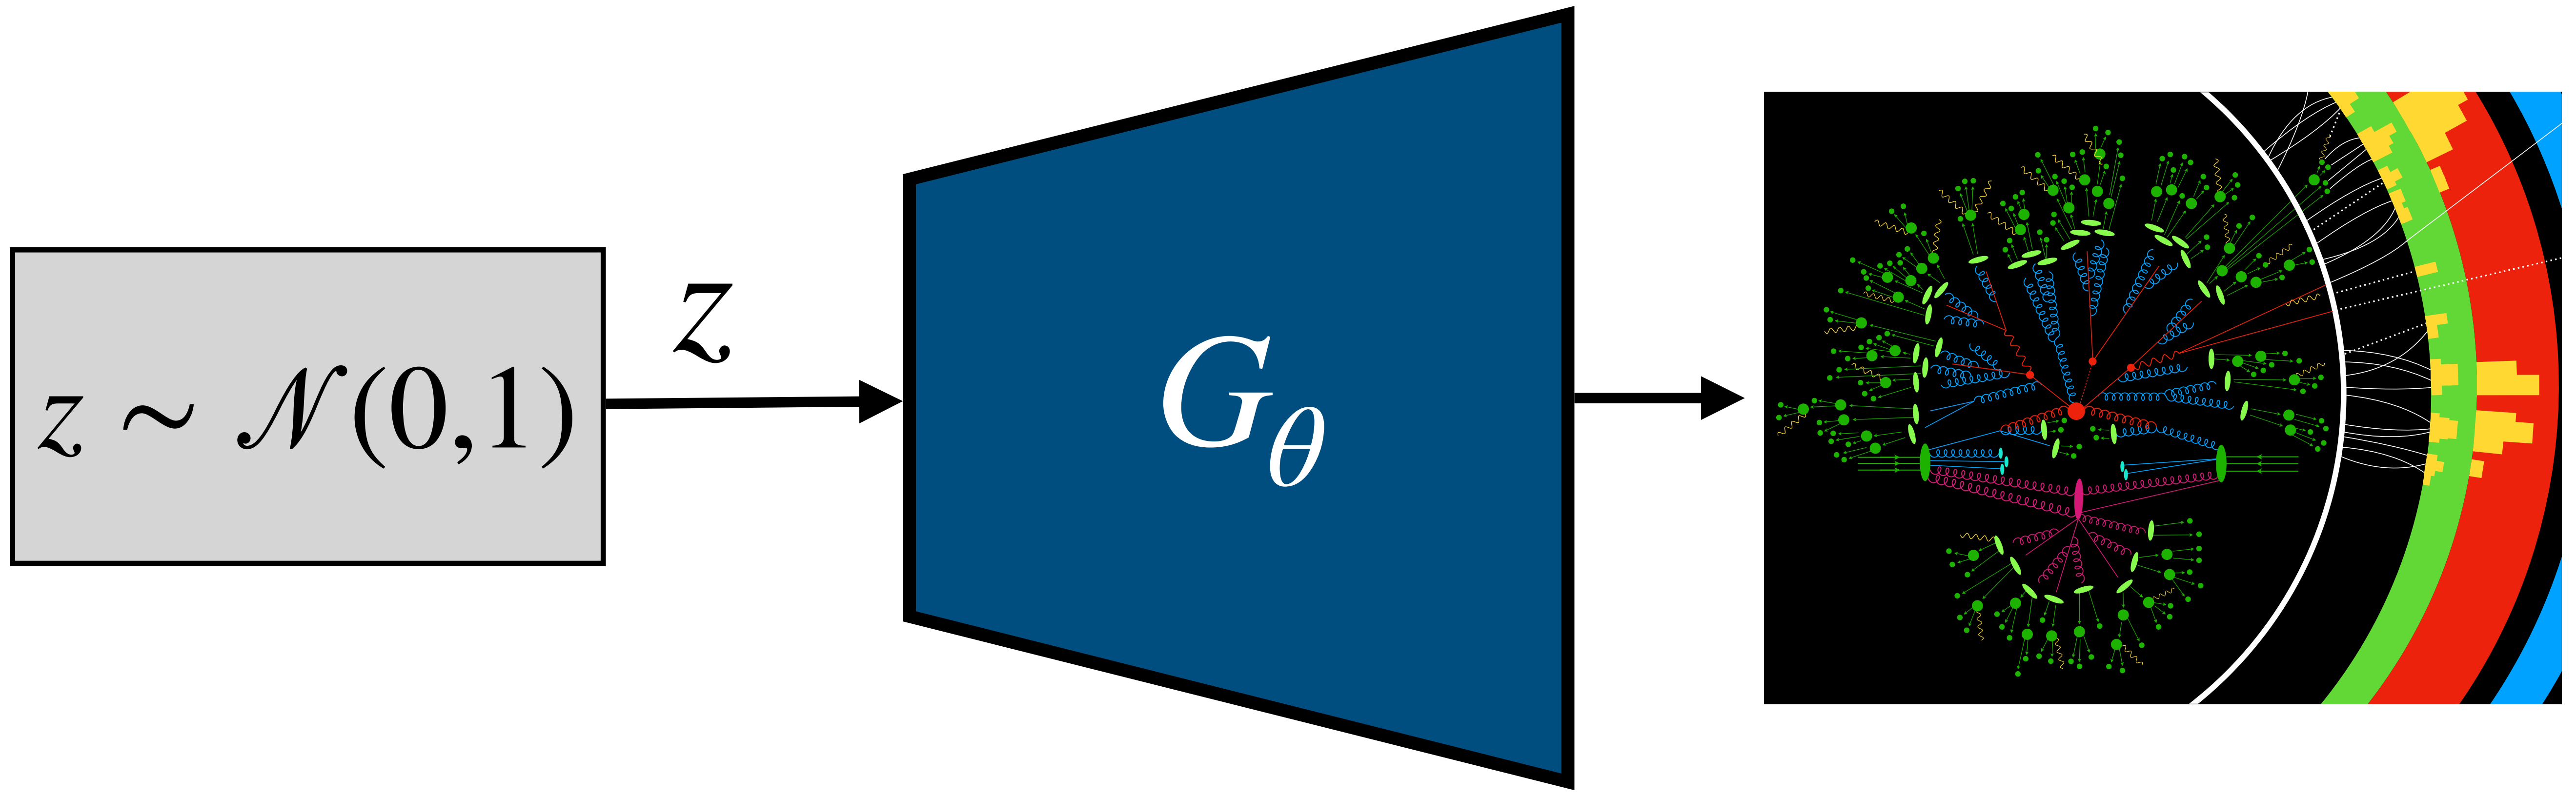

AI Uncertainties

Physics needs uncertainty estimates

ML uncertainty quantification:

- Heteroscedastic losses, repulsive ensemble, **bayesian neural networks**
- Transferable to generative models with likelihood losses
- Bayesian normalizing flows
- Uncertainty on **density estimation**

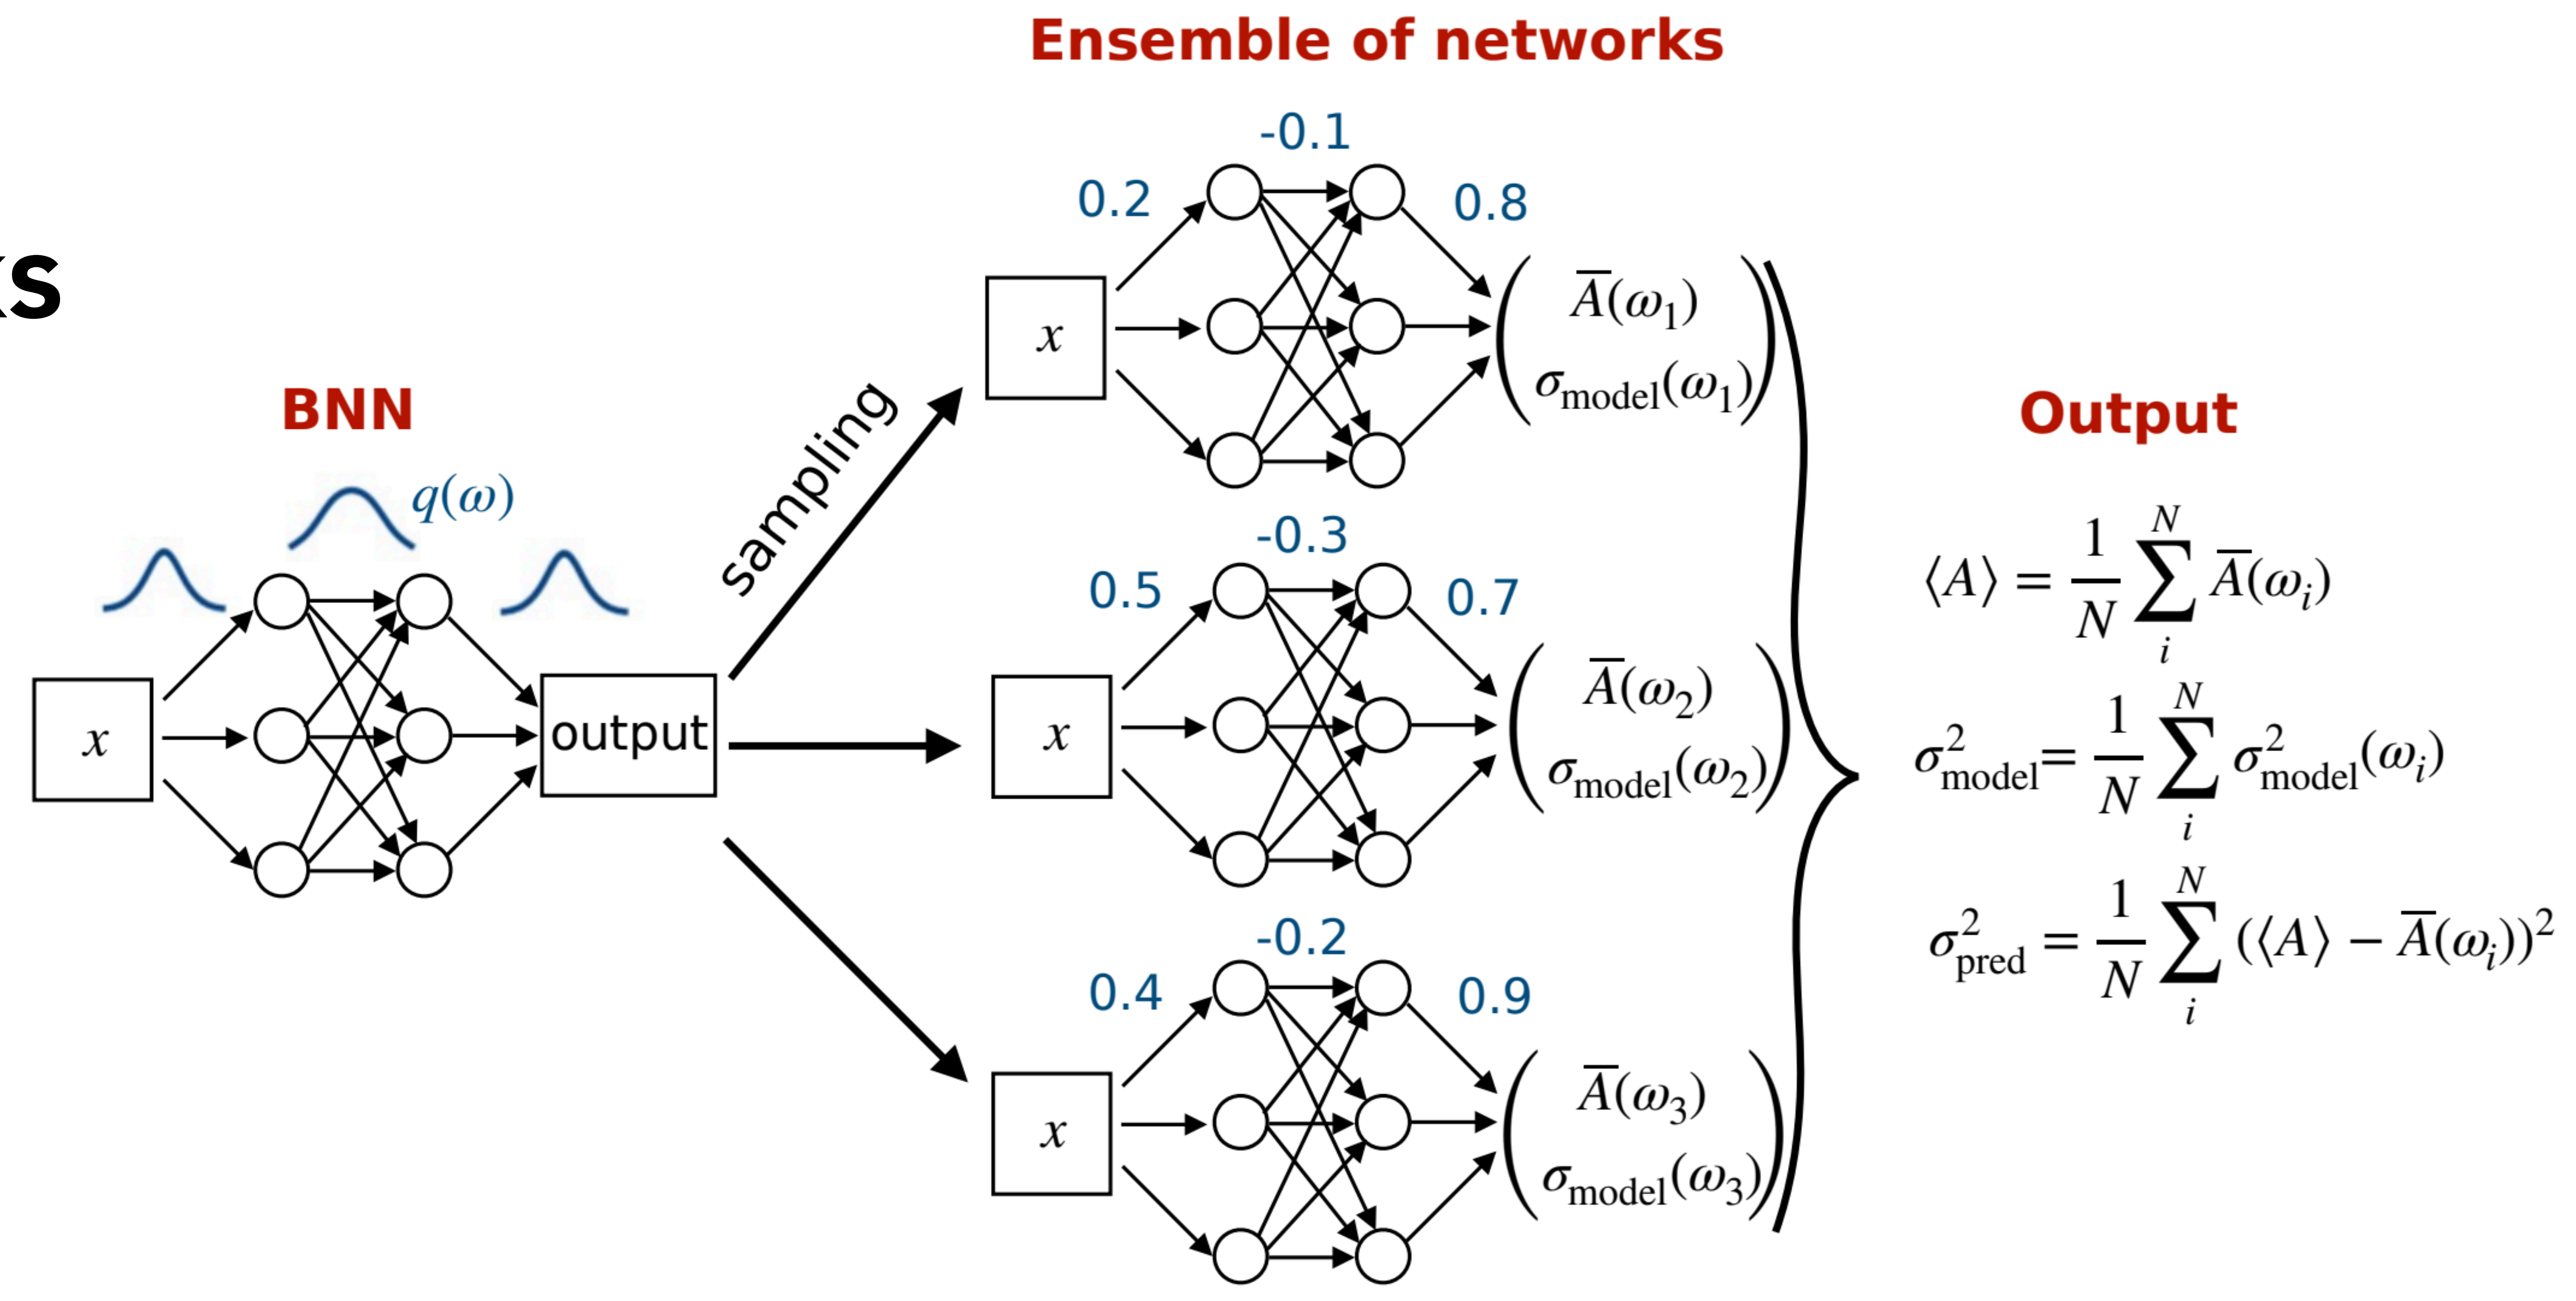

AI Uncertainties

What uncertainty is learned?

- For Generative models:
- **Statistical uncertainty** σ_{stat}
  - Scales $\dfrac{a}{\sqrt{n_{\text{gen}}}}$
  - Infinitely reducible
- **Generalization uncertainty** σ_{model}
  - Scales $\dfrac{b}{\sqrt{n_{\text{train}}}}$
  - Irreducible by sampling

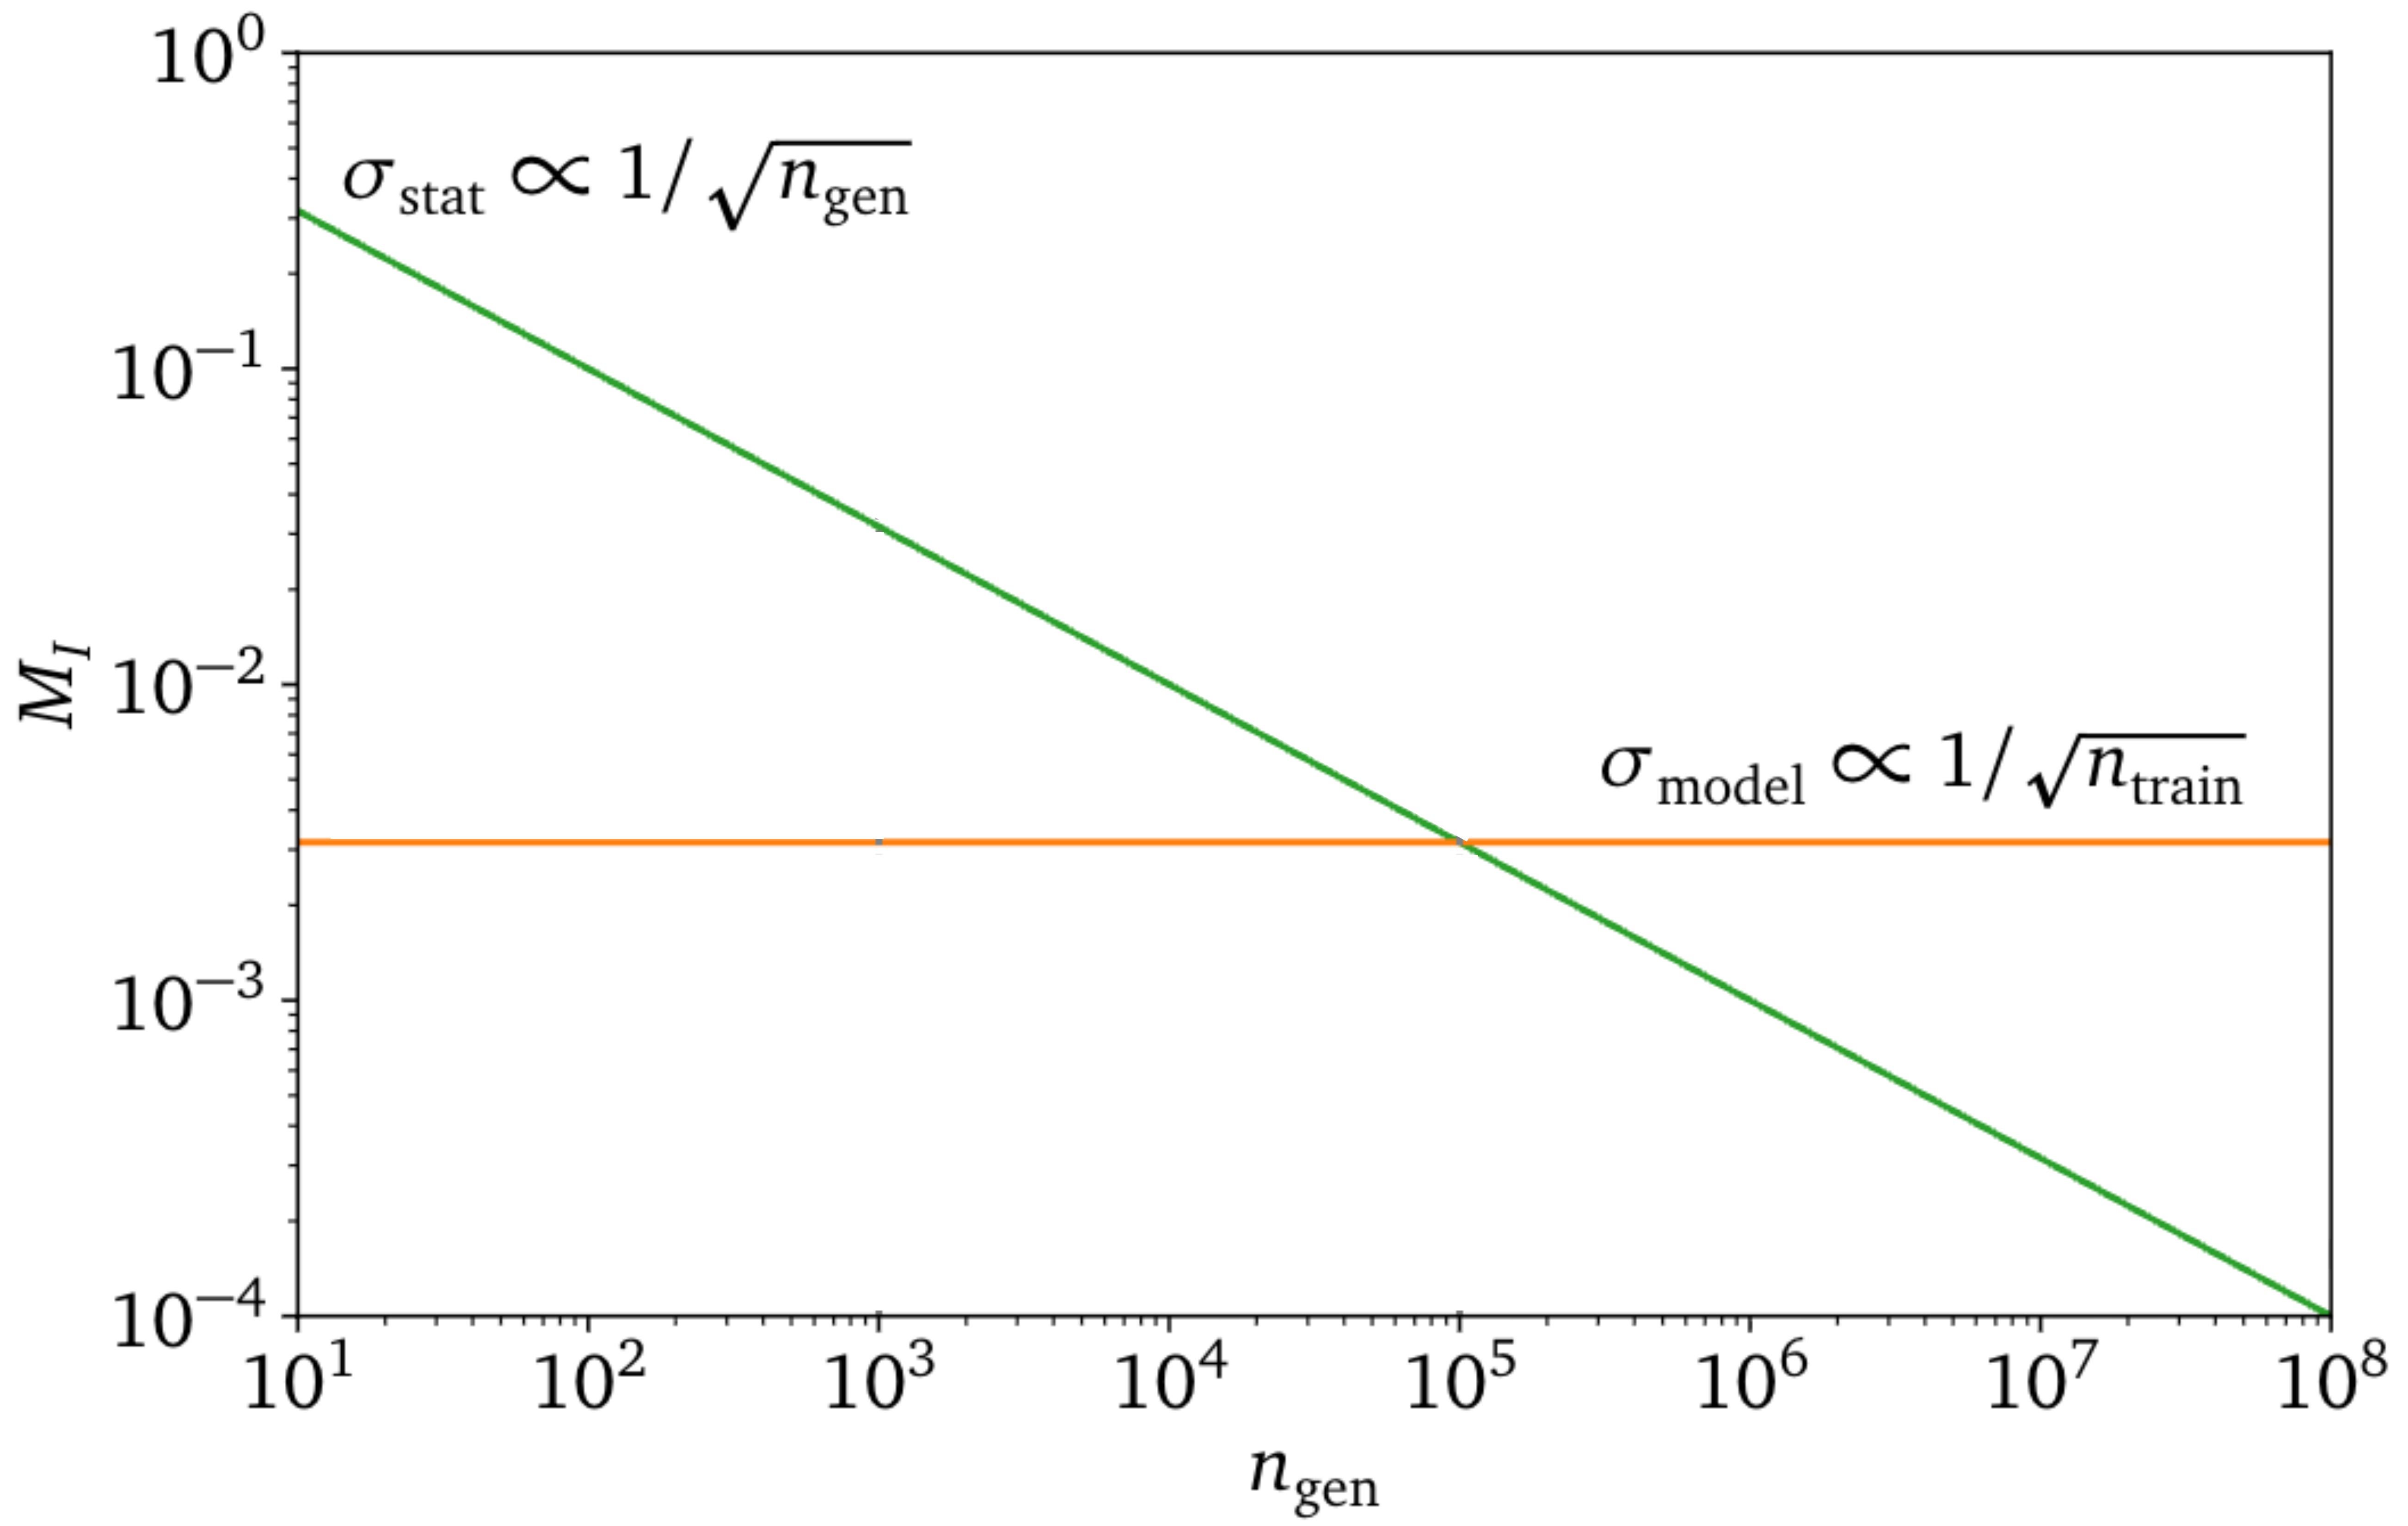

AI Uncertainties

Generalization uncertainty

- Most applications:
  - Inference data ≠ training data
  - Extrapolation outside training

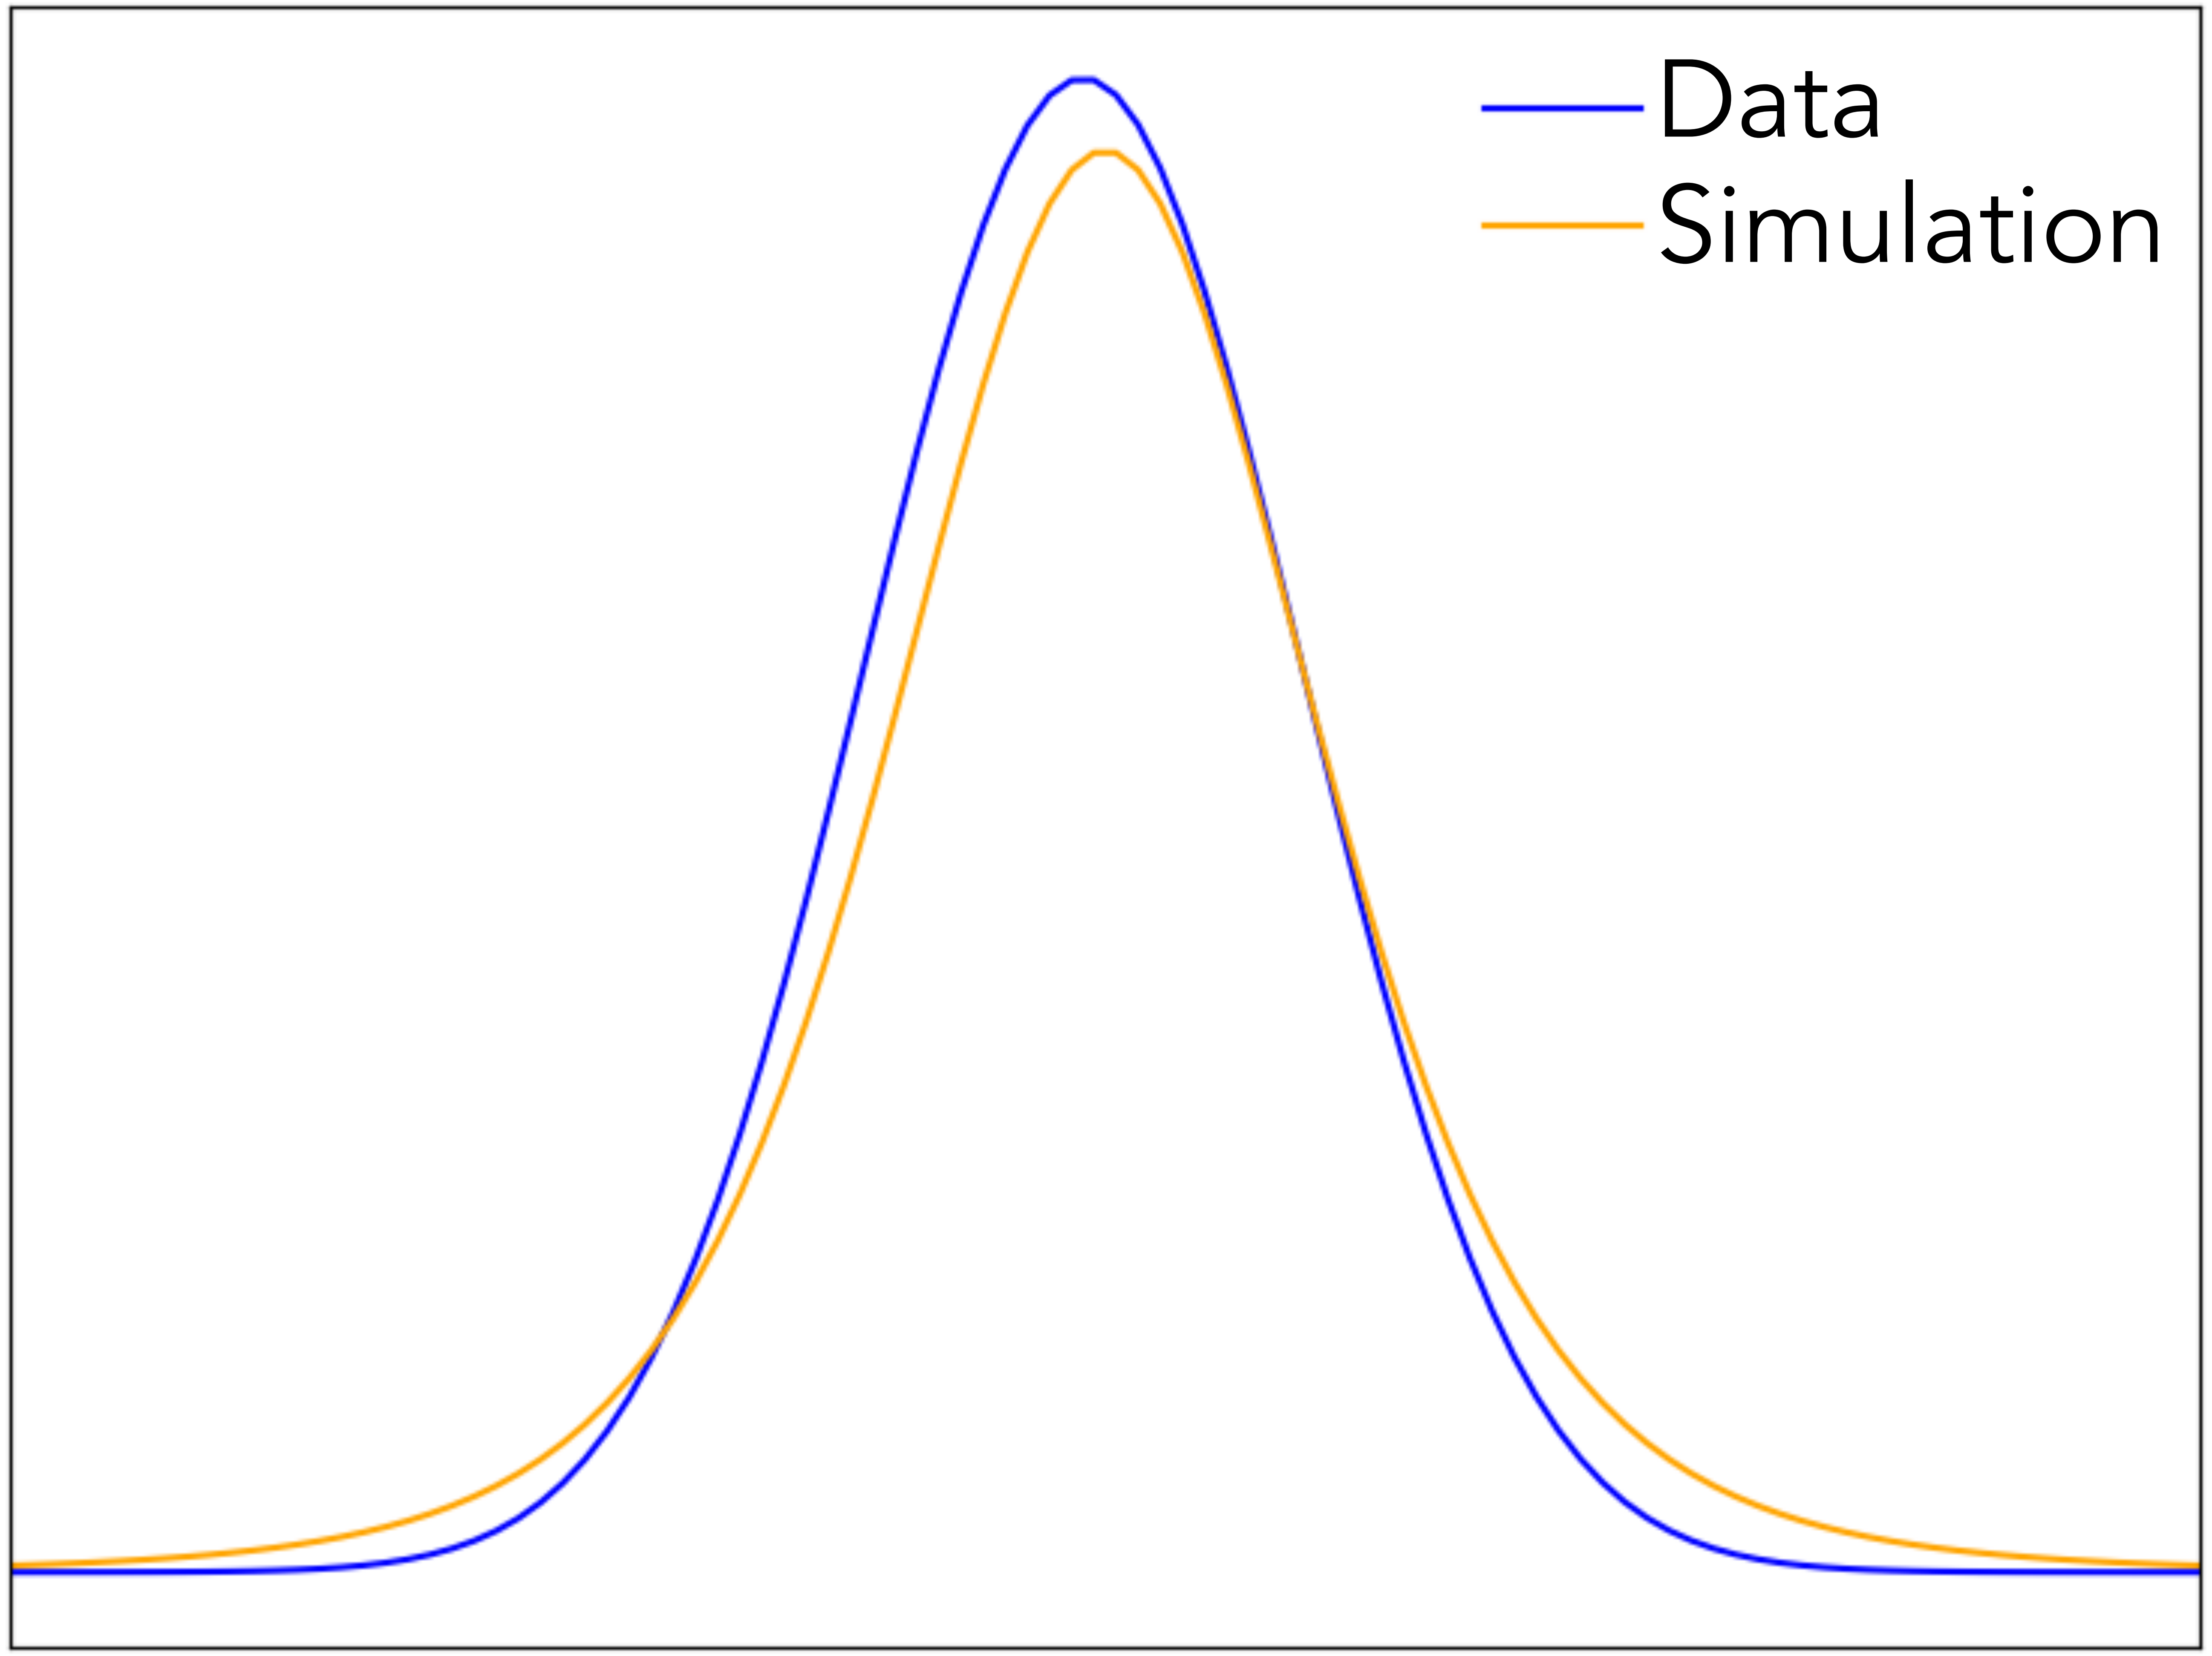

AI Uncertainties

Generalization uncertainty

- Most applications:
  - Inference data ≠ training data
  - Extrapolation outside training
- Generative modeling:
  - Extrapolation in resolution

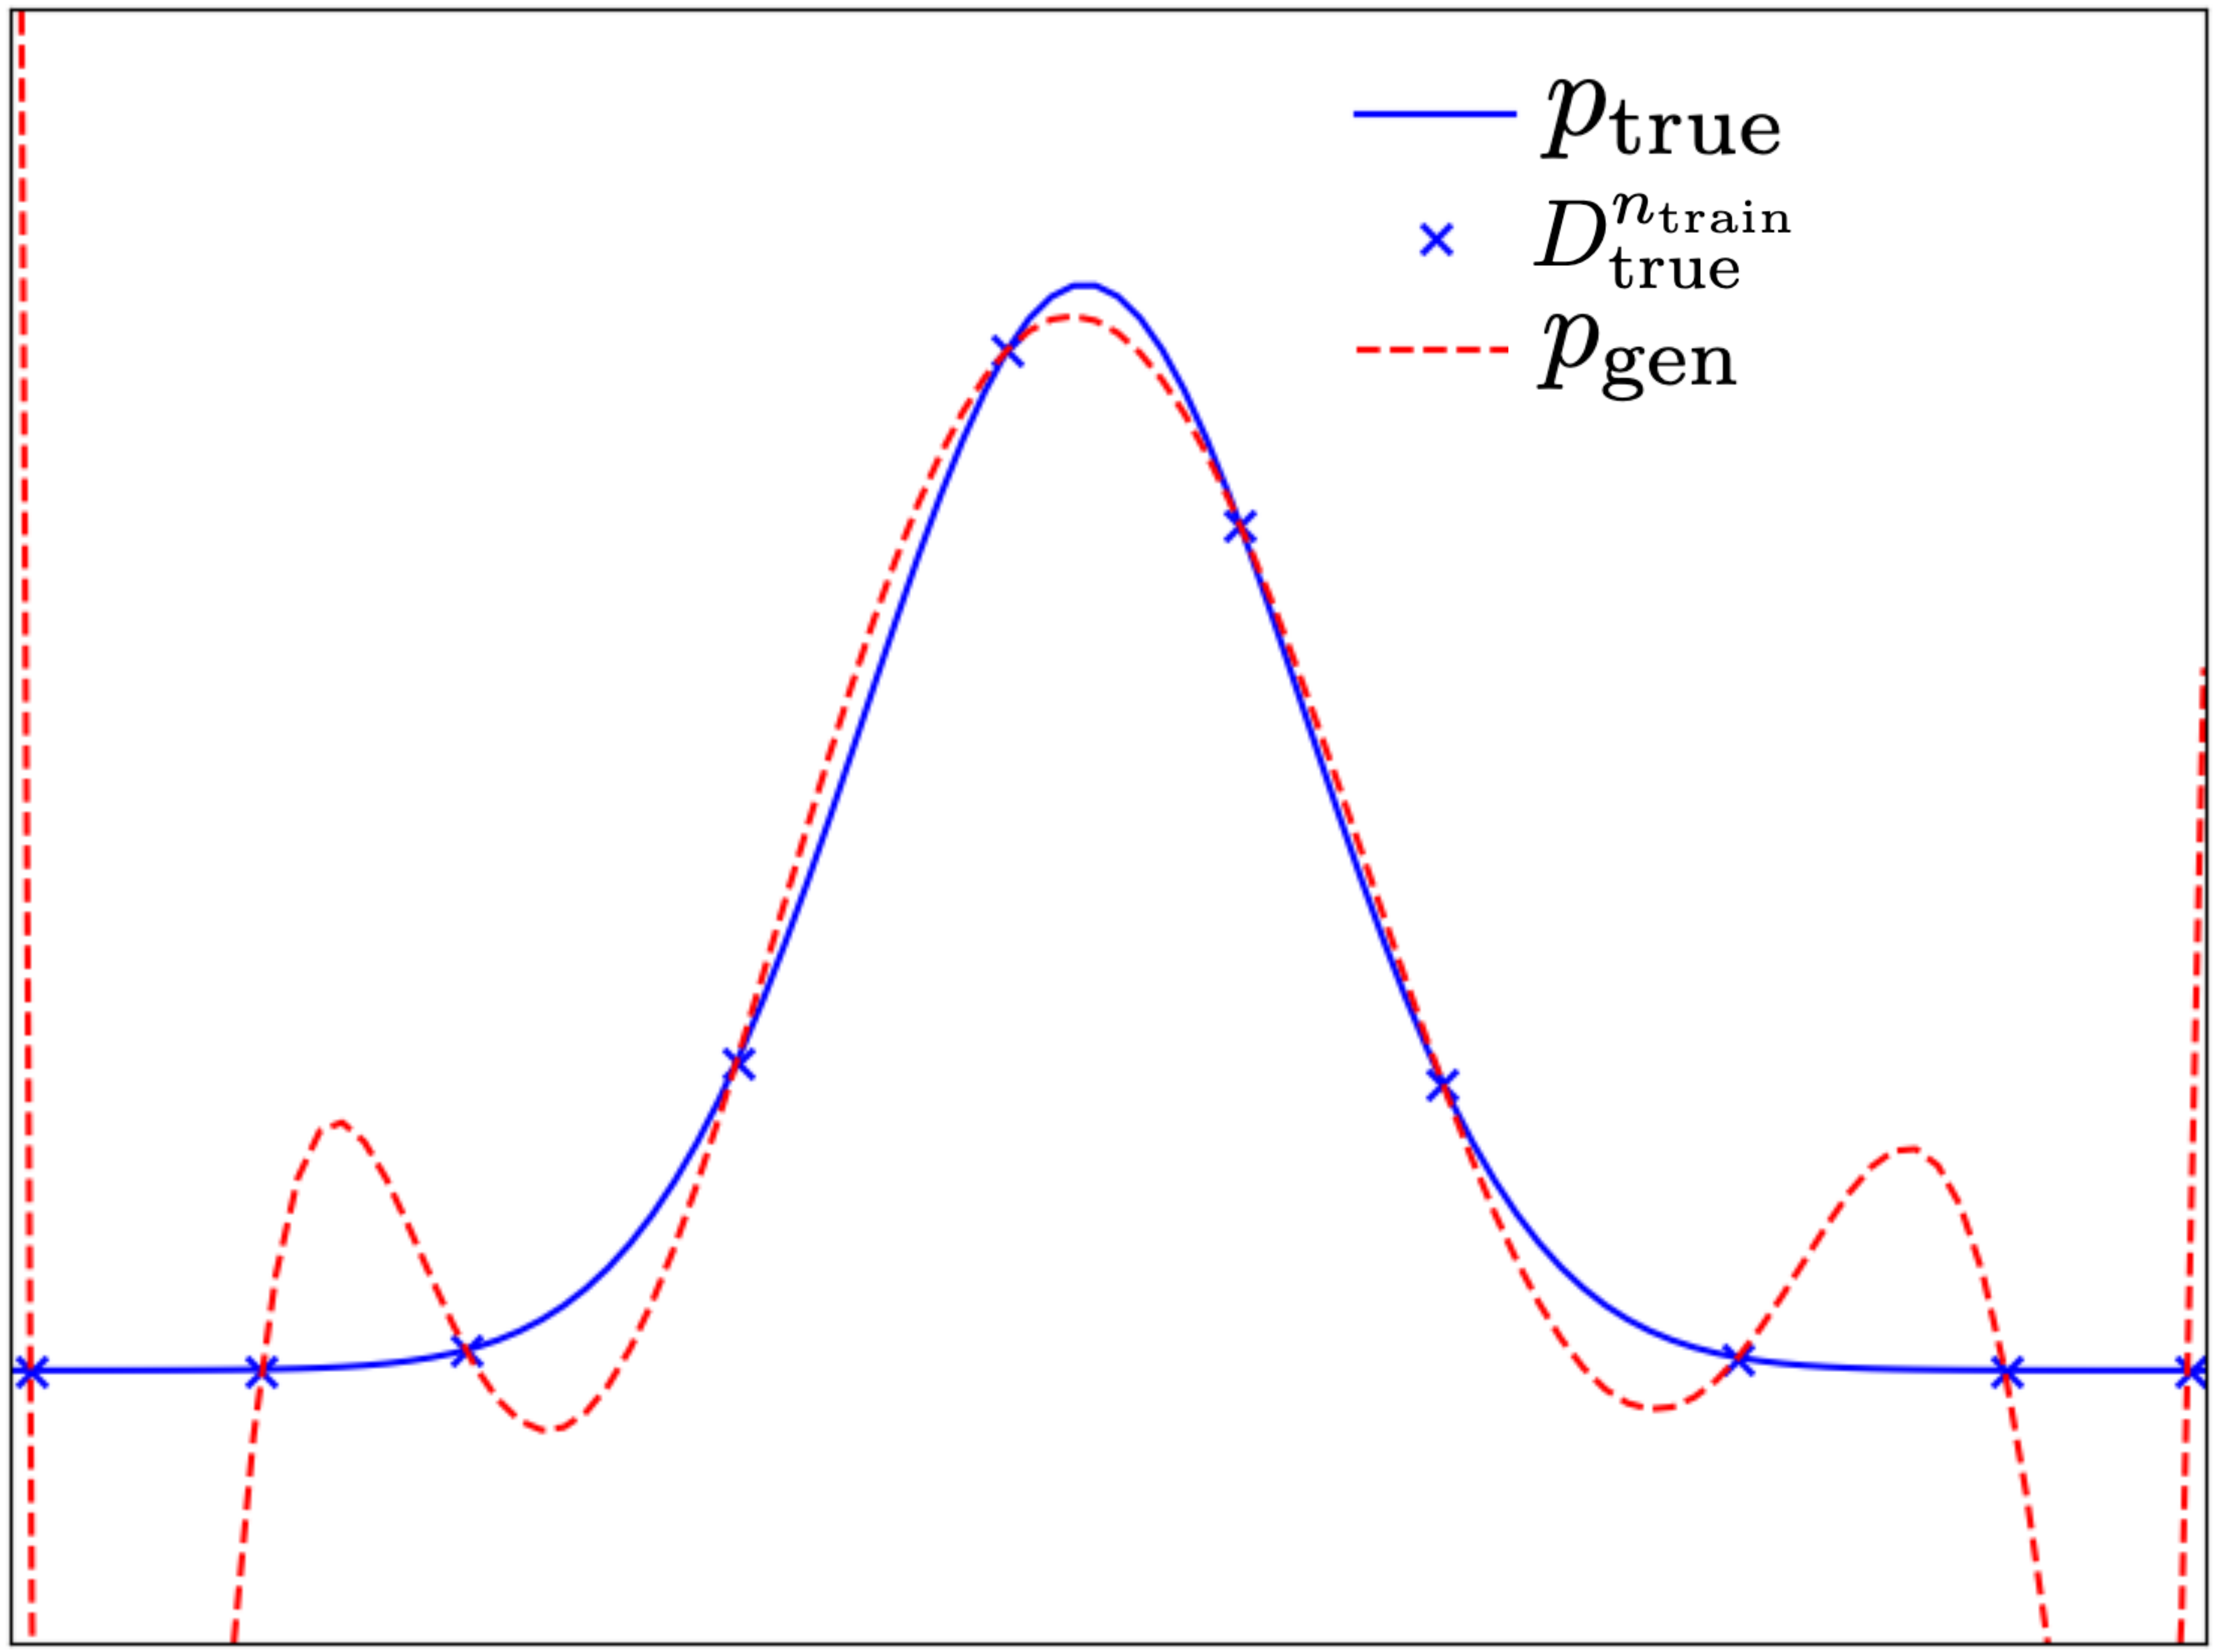

AI Uncertainties

Fundamental question

- What happens if I exceed the training resolution?
- How many points n_{gen} can I draw?
- Generative simulation assumes $n_{\mathrm{gen}} > n_{\mathrm{train}}$
- Is this really true?

“If I train a model on 1 million data points, how many can I draw from that model?”
—Louis Lyons

→Amplification

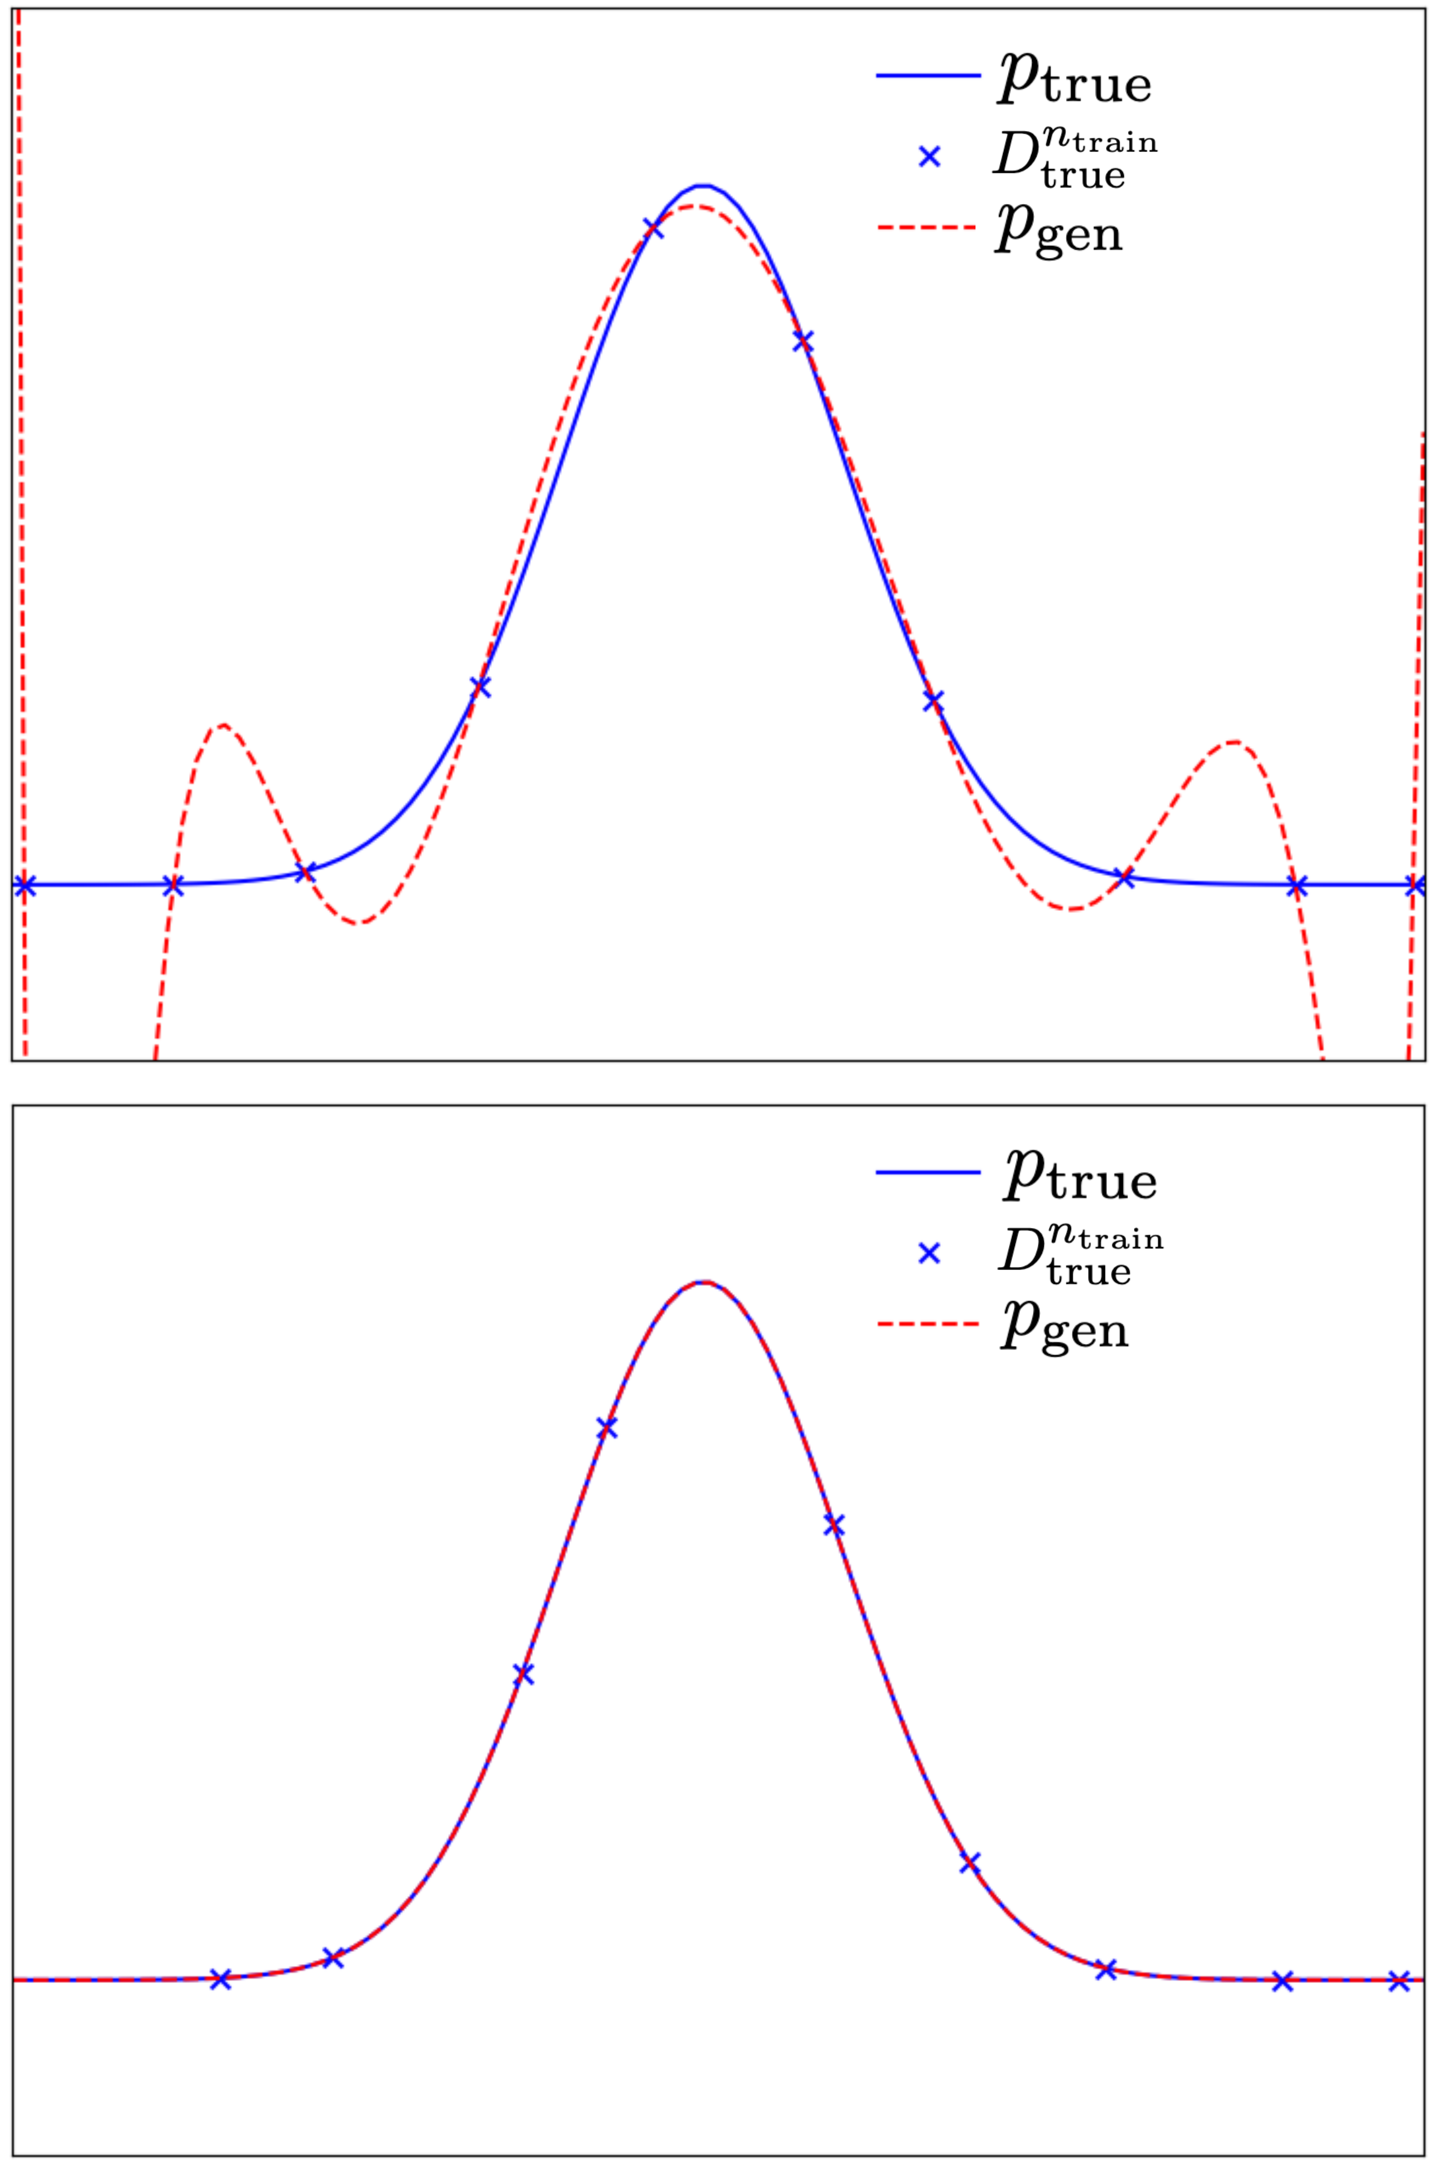

Amplification

How many generative points can I sample from the model?

$$\sigma^2_{\text{gen}} = \sigma^2_{\text{model}} + \sigma^2_{\text{stat}}$$

$\sigma_{\text{model}} > \sigma_{\text{stat}}$: new points stop being useful

$$\sigma_{\text{stat}}(n)\Big|_{n=n_{\text{equiv}}} \overset{!}{=} \sigma_{\text{model}}$$

n_{equiv}: size of true data set with same stat. uncertainty as generalization uncertainty

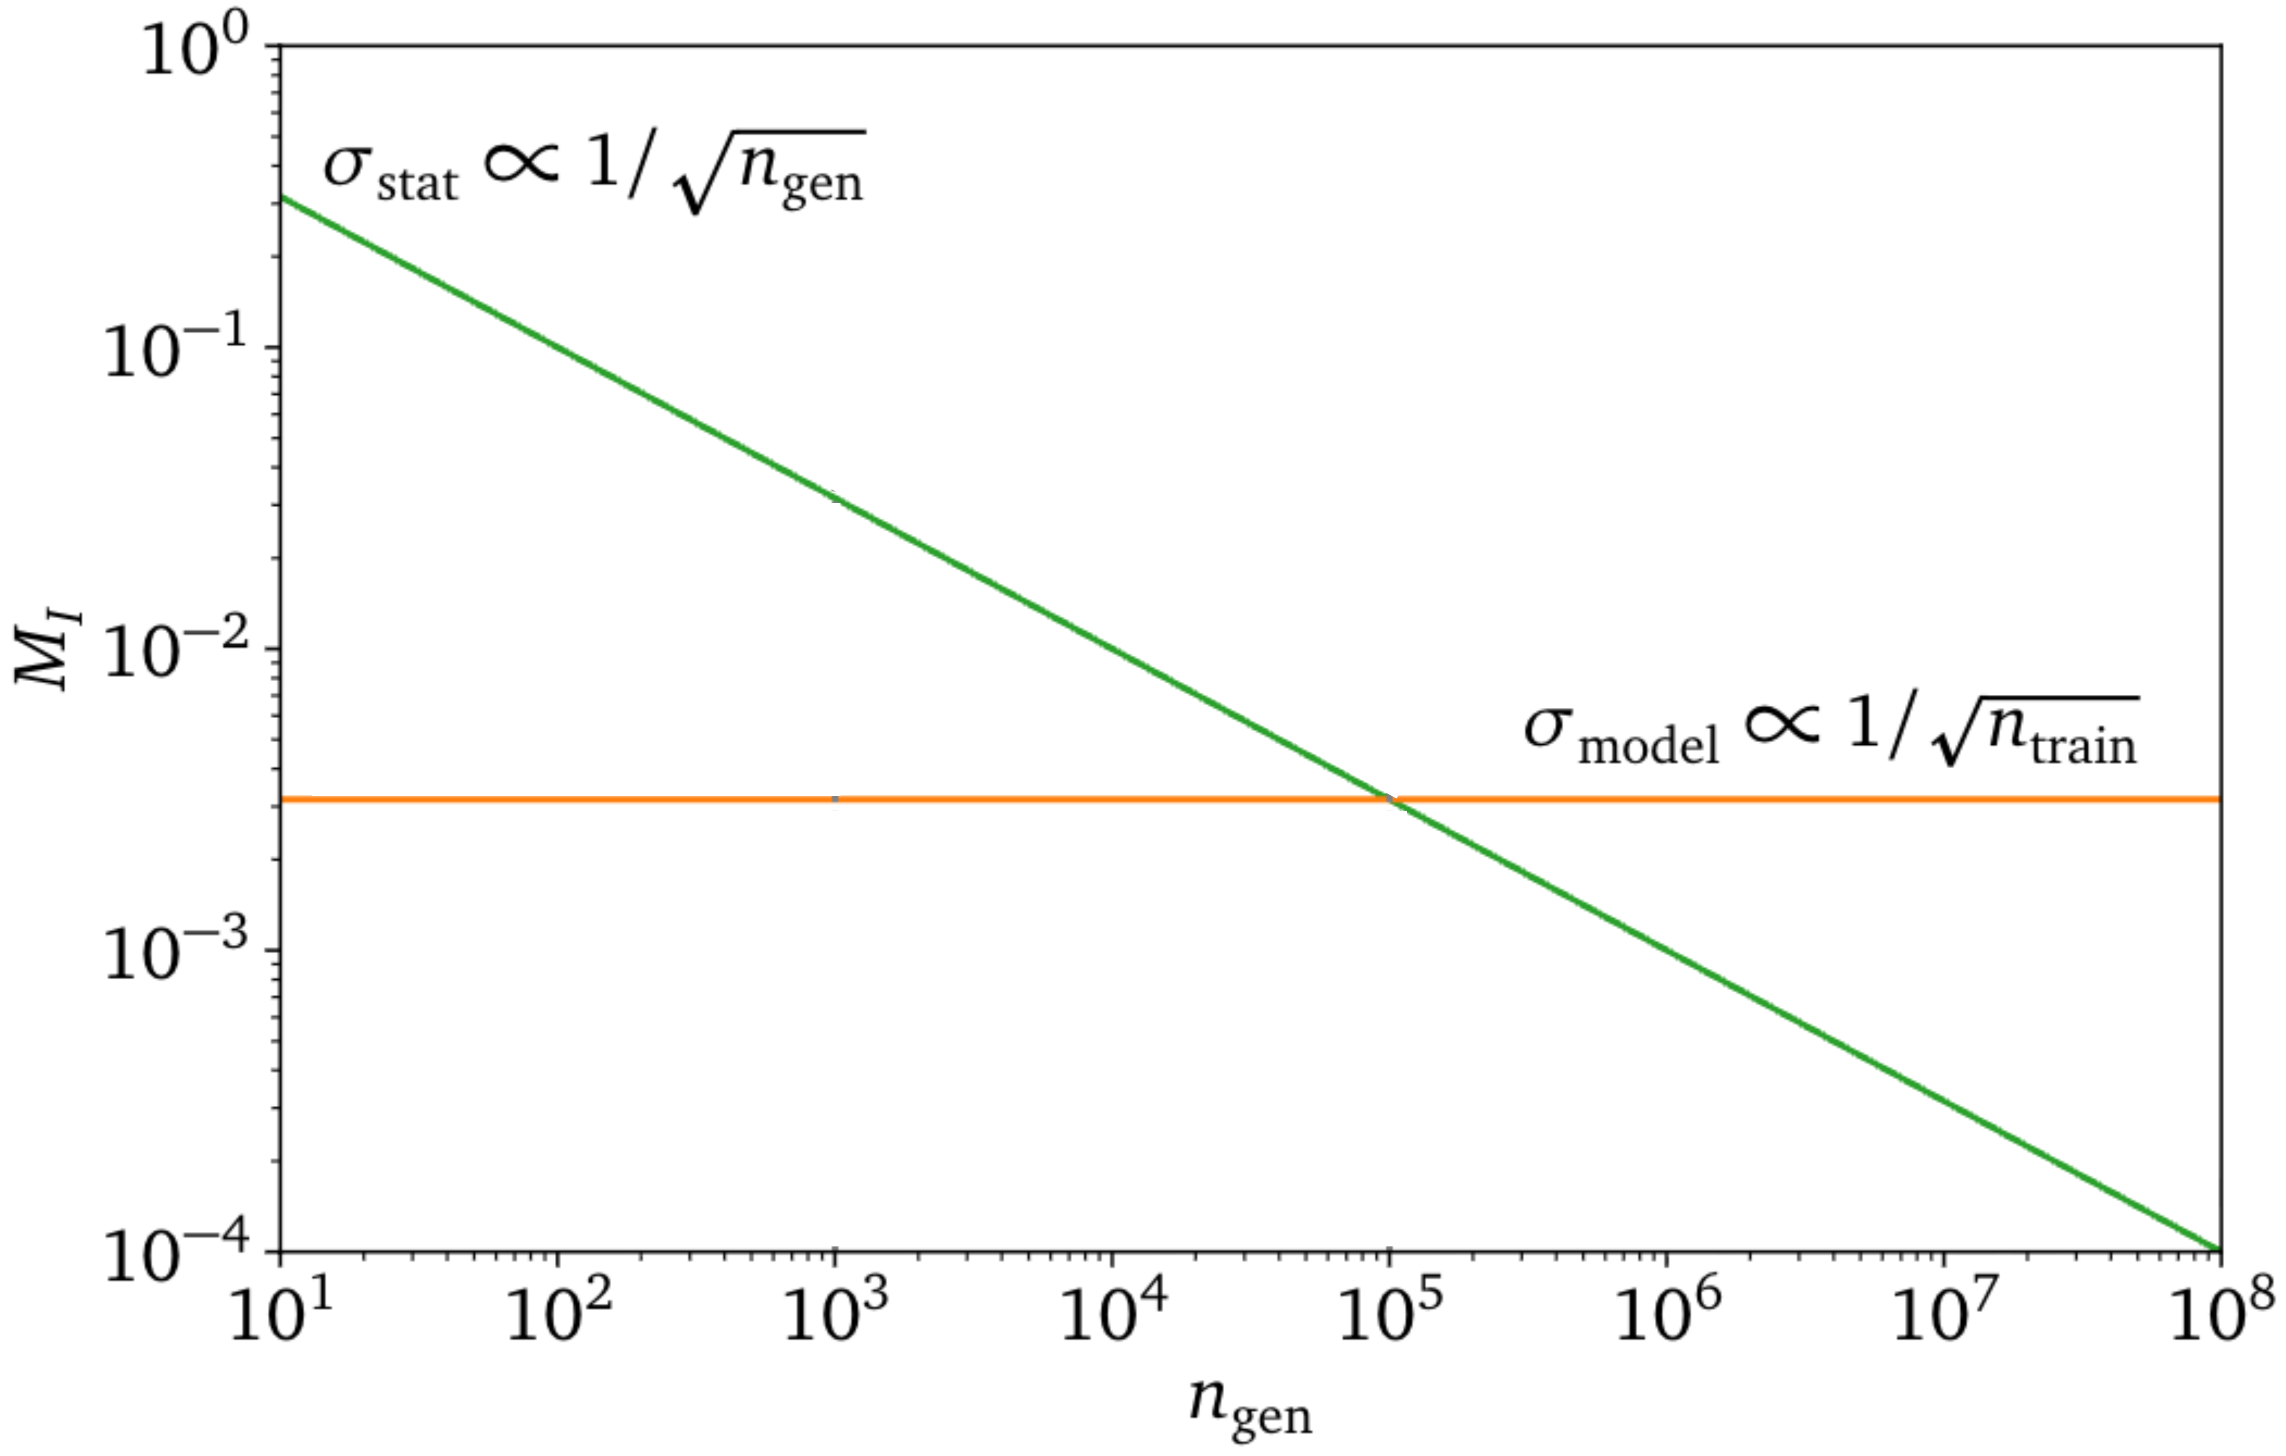

Amplification

How many generative points can I sample from the model?

$$\sigma_{\text{gen}}^2 = \sigma_{\text{model}}^2 + \sigma_{\text{stat}}^2$$

$\sigma_{\text{model}} > \sigma_{\text{stat}}$: new points stop being useful

$$\sigma_{\text{stat}}(n)\Big|_{n=n_{\text{equiv}}} \stackrel{!}{=} \sigma_{\text{model}}$$

n_{equiv} : size of true data set with same stat. uncertainty as generalization uncertainty

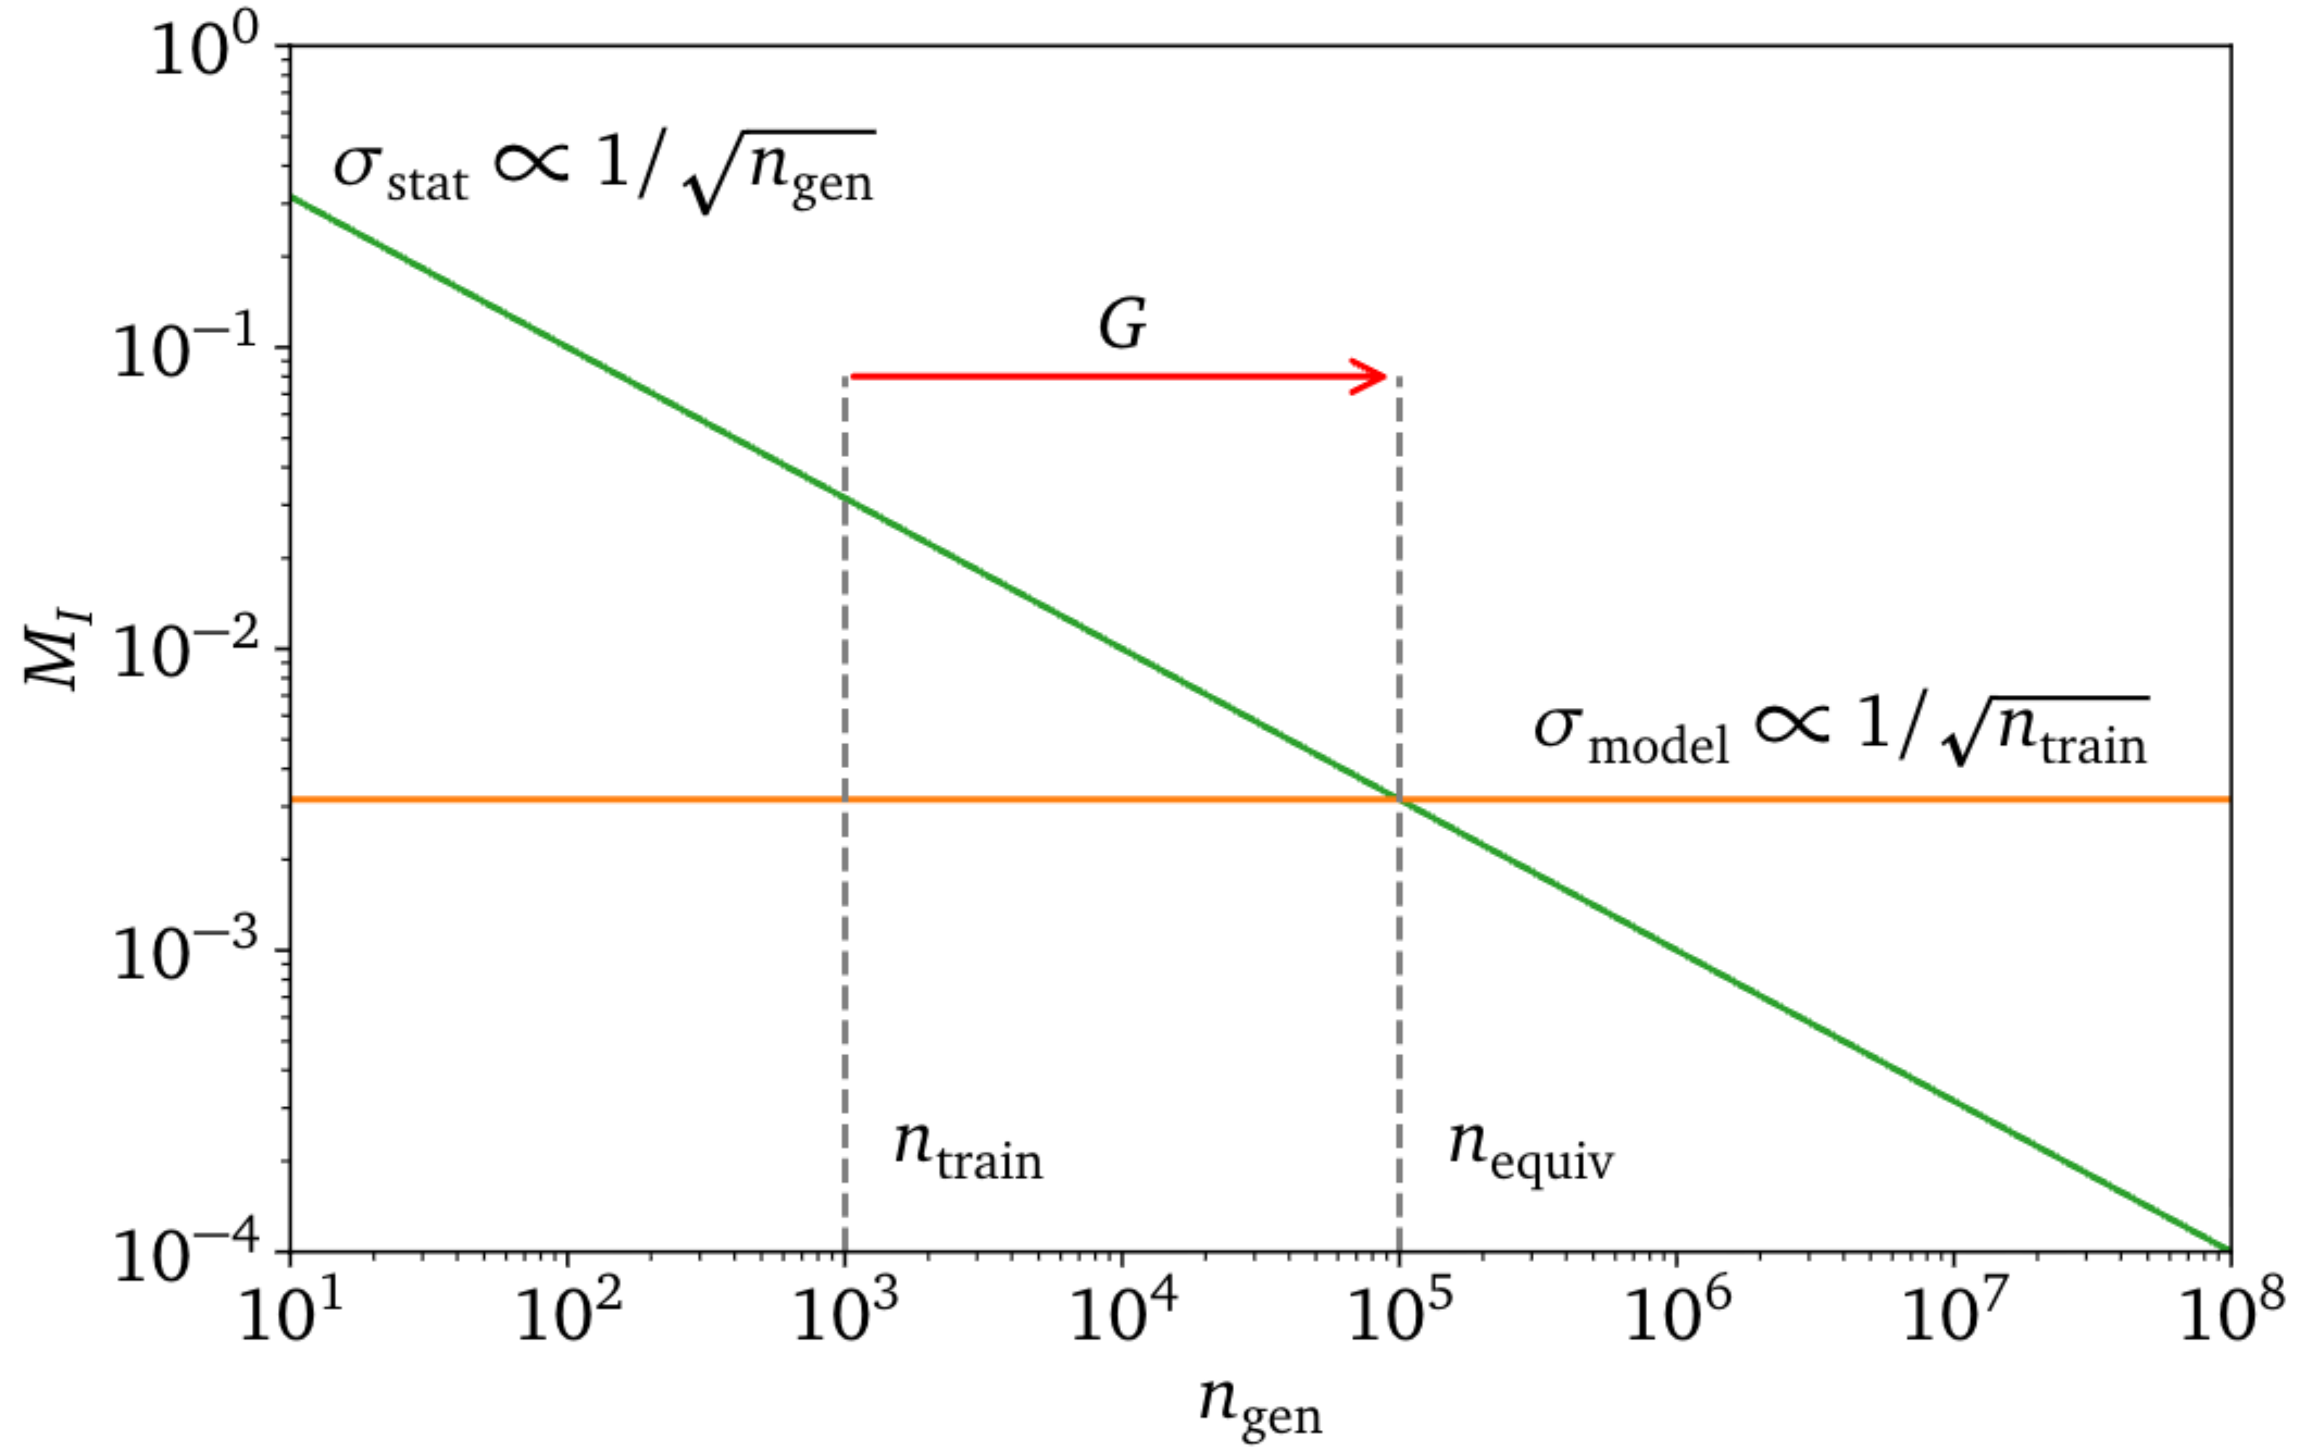

$n_{\text{equiv}} \leq n_{\text{train}}$: No Amplification

$n_{\text{equiv}} > n_{\text{train}}$: Amplification

$G = \dfrac{n_{\text{eqiv}}}{n_{\text{train}}}$: Amplification factor

Quantifying Amplification

Example: Gaussian Ring Radius

- Plot σ_{gen}
- Get $\sigma_{\text{stat}}(n)$ from falling section
- Get σ_{model} from asymptotic part

$$\sigma_{\text{stat}}(n)\Big|_{n=n_{\text{equiv}}} \stackrel{!}{=} \sigma_{\text{model}}$$

$$G = \frac{n_{\text{eqiv}}}{n_{\text{train}}}$$

- Estimated G agrees with truth

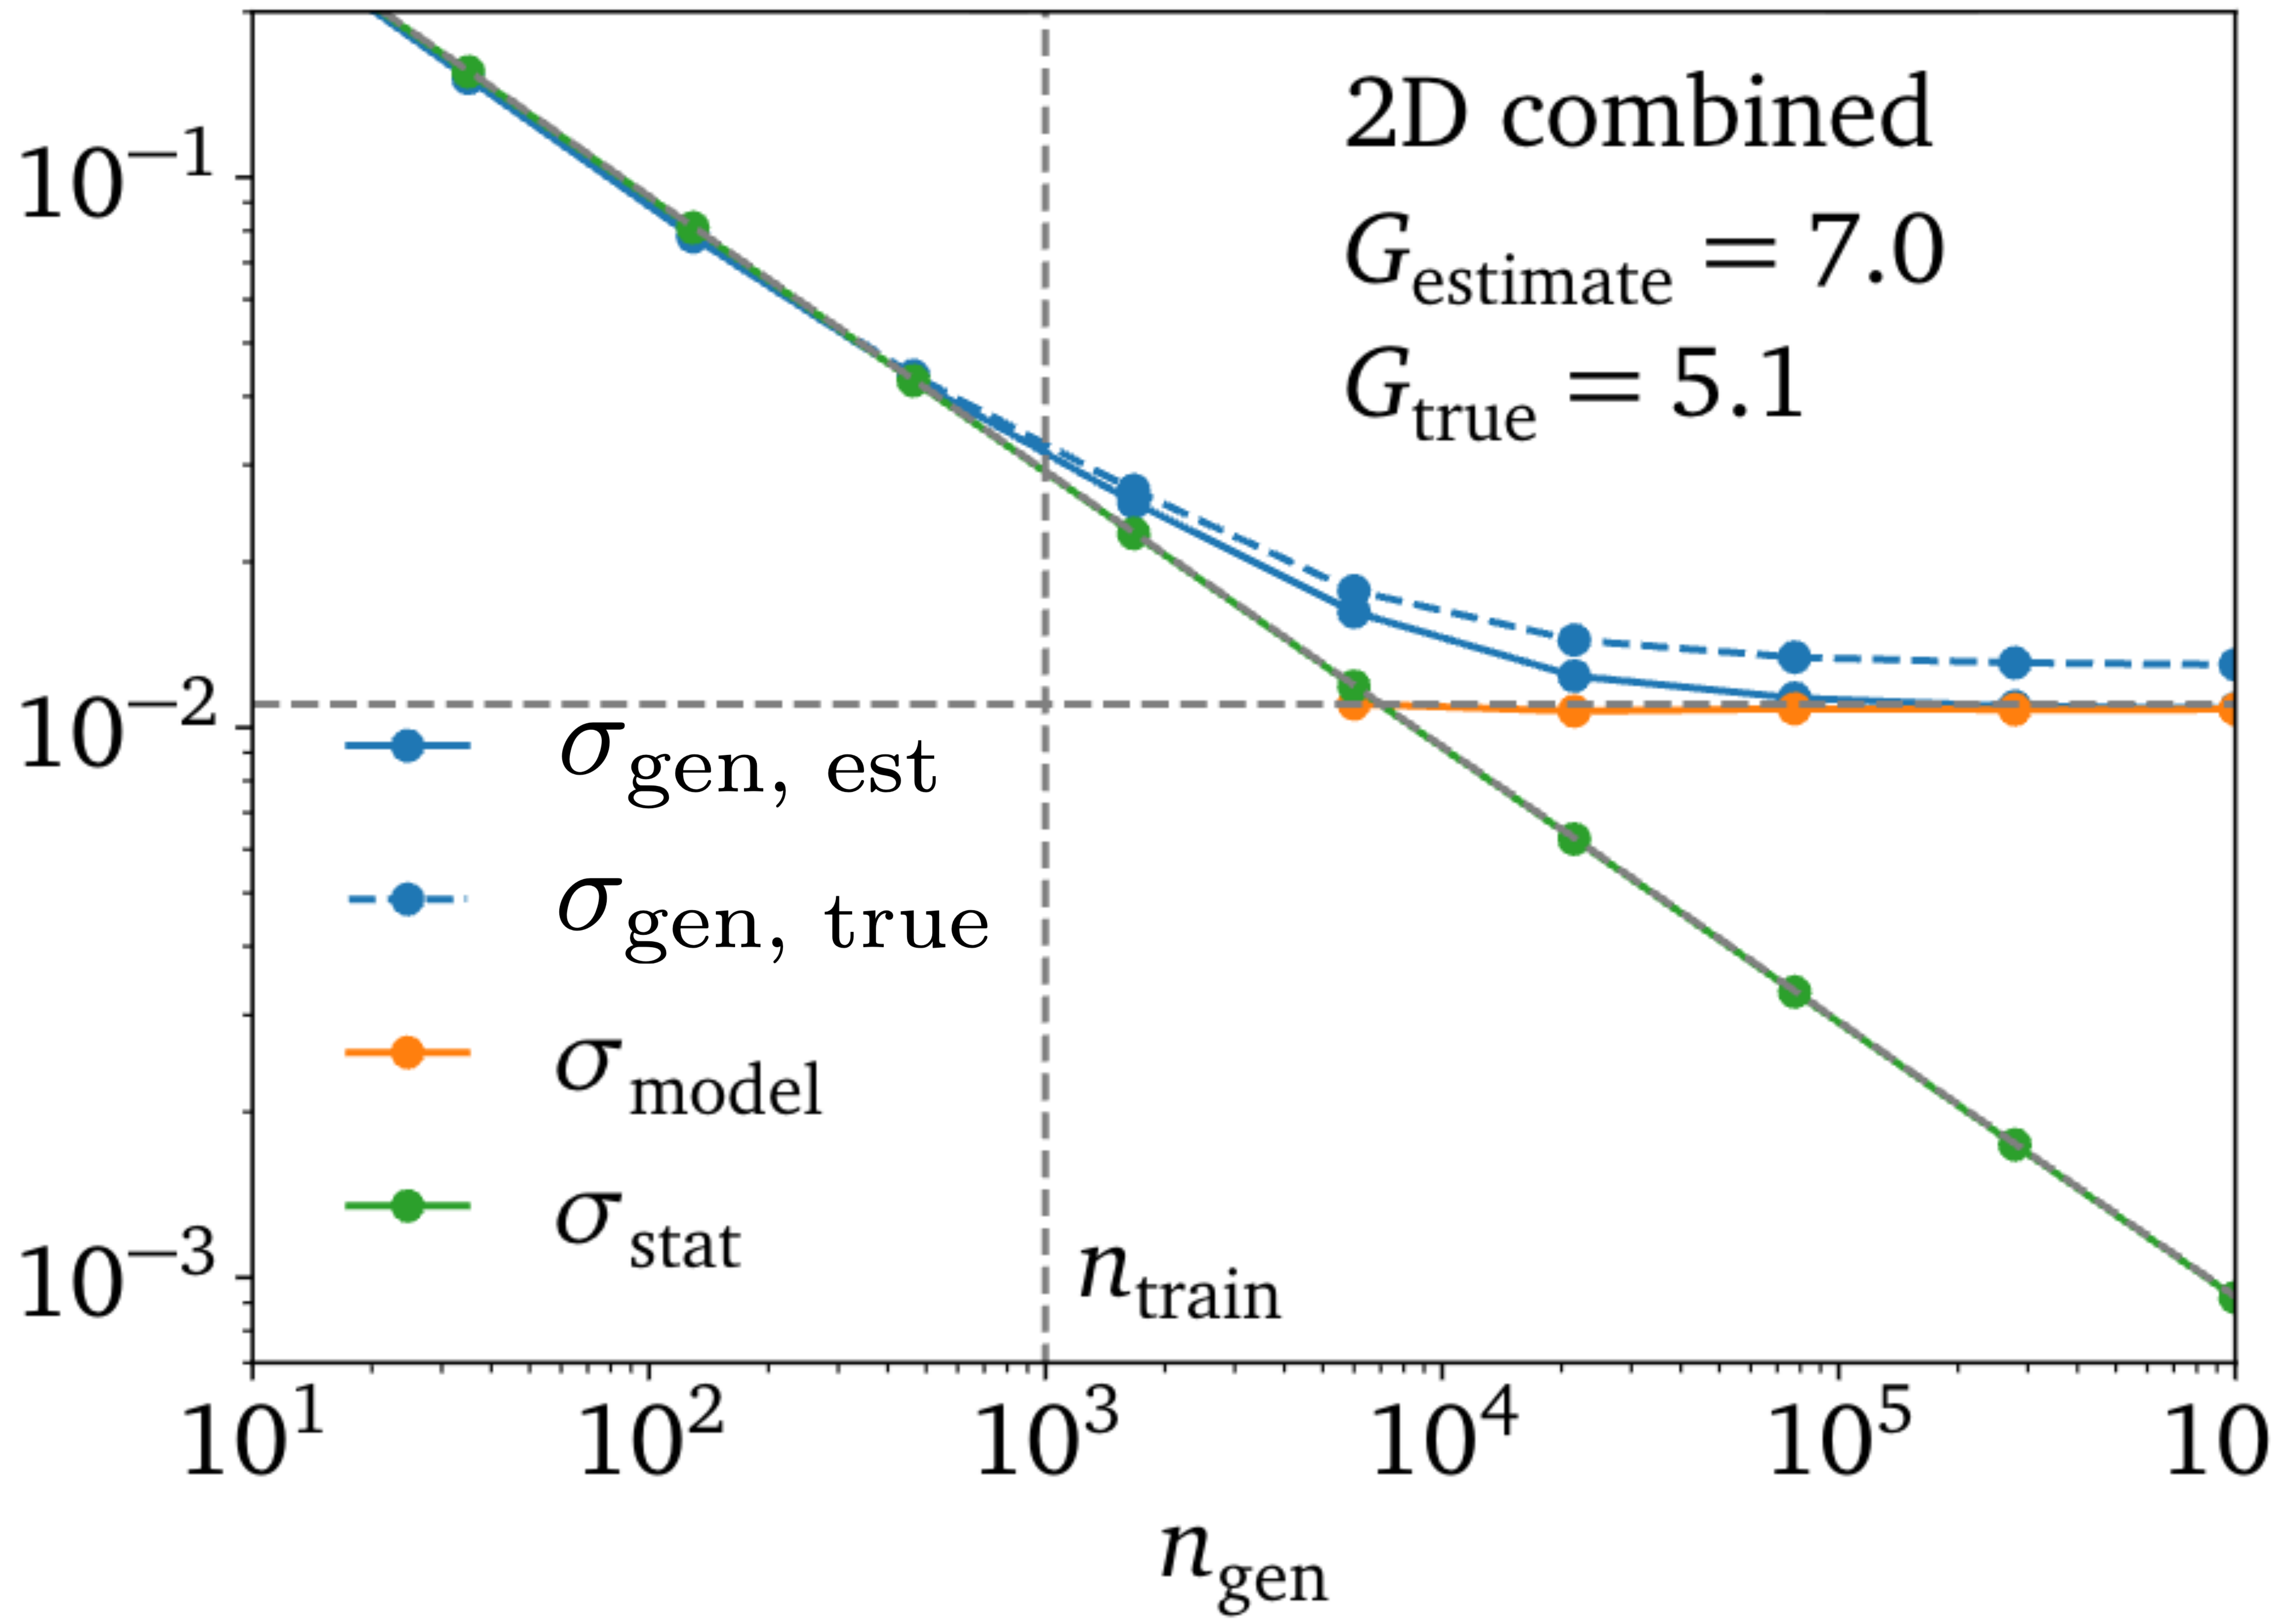

Bahl, **Diefenbacher** et al. Forecasting Generative Amplification 2509.08048

Quantifying Amplification

How to get σ_{gen} ?

- Uncertainty quantification?
  - Only σ_{model} not σ_{stat}

→ Need metric $M\left[D^n_{\text{gen}}, p_{\text{true}}\right]$ agreement (generated) set vs. true density

$$M\left[D^n_{\text{true}}, p_{\text{true}}\right] \approx \sigma_{\text{stat}} \qquad\qquad M\left[D^n_{\text{gen}}, p_{\text{true}}\right]\Big|_{n_{\text{gen}}\to\infty} \to \sigma_{\text{model}}$$

$$\boxed{\begin{aligned} \sigma_{\text{stat}}(n)\Big|_{n=n_{\text{equiv}}} &\overset{!}{=} \sigma_{\text{model}} \\ &\Updownarrow \\ M\left[D^n_{\text{true}}, p_{\text{true}}\right]\Big|_{n=n_{\text{equiv}}} &\overset{!}{=} M\left[D^n_{\text{gen}}, p_{\text{true}}\right]\Big|_{n_{\text{gen}}\to\infty} \end{aligned}}$$

Quantifying Amplification

First Metric: Averaging Amplification

- Fractions inside phase space region V

$$I = \int_V dx\, p_{\text{true}}(x)$$

$$\bar{I} = \frac{1}{n_{\text{gen}}} \sum_{x \in D_{\text{gen}}^{n_{\text{gen}}}} \mathbf{1}_{x \in V}$$

- Expected difference as metric

$$M_I\left[D_{\text{gen}}^{n_{\text{gen}}}, p_{\text{true}}\right] = \left\langle \left[I(p_{\text{true}}) - \bar{I}(D_{\text{gen}}^{n_{\text{gen}}}) \right]^2 \right\rangle$$

- Estimate this variance via BNN

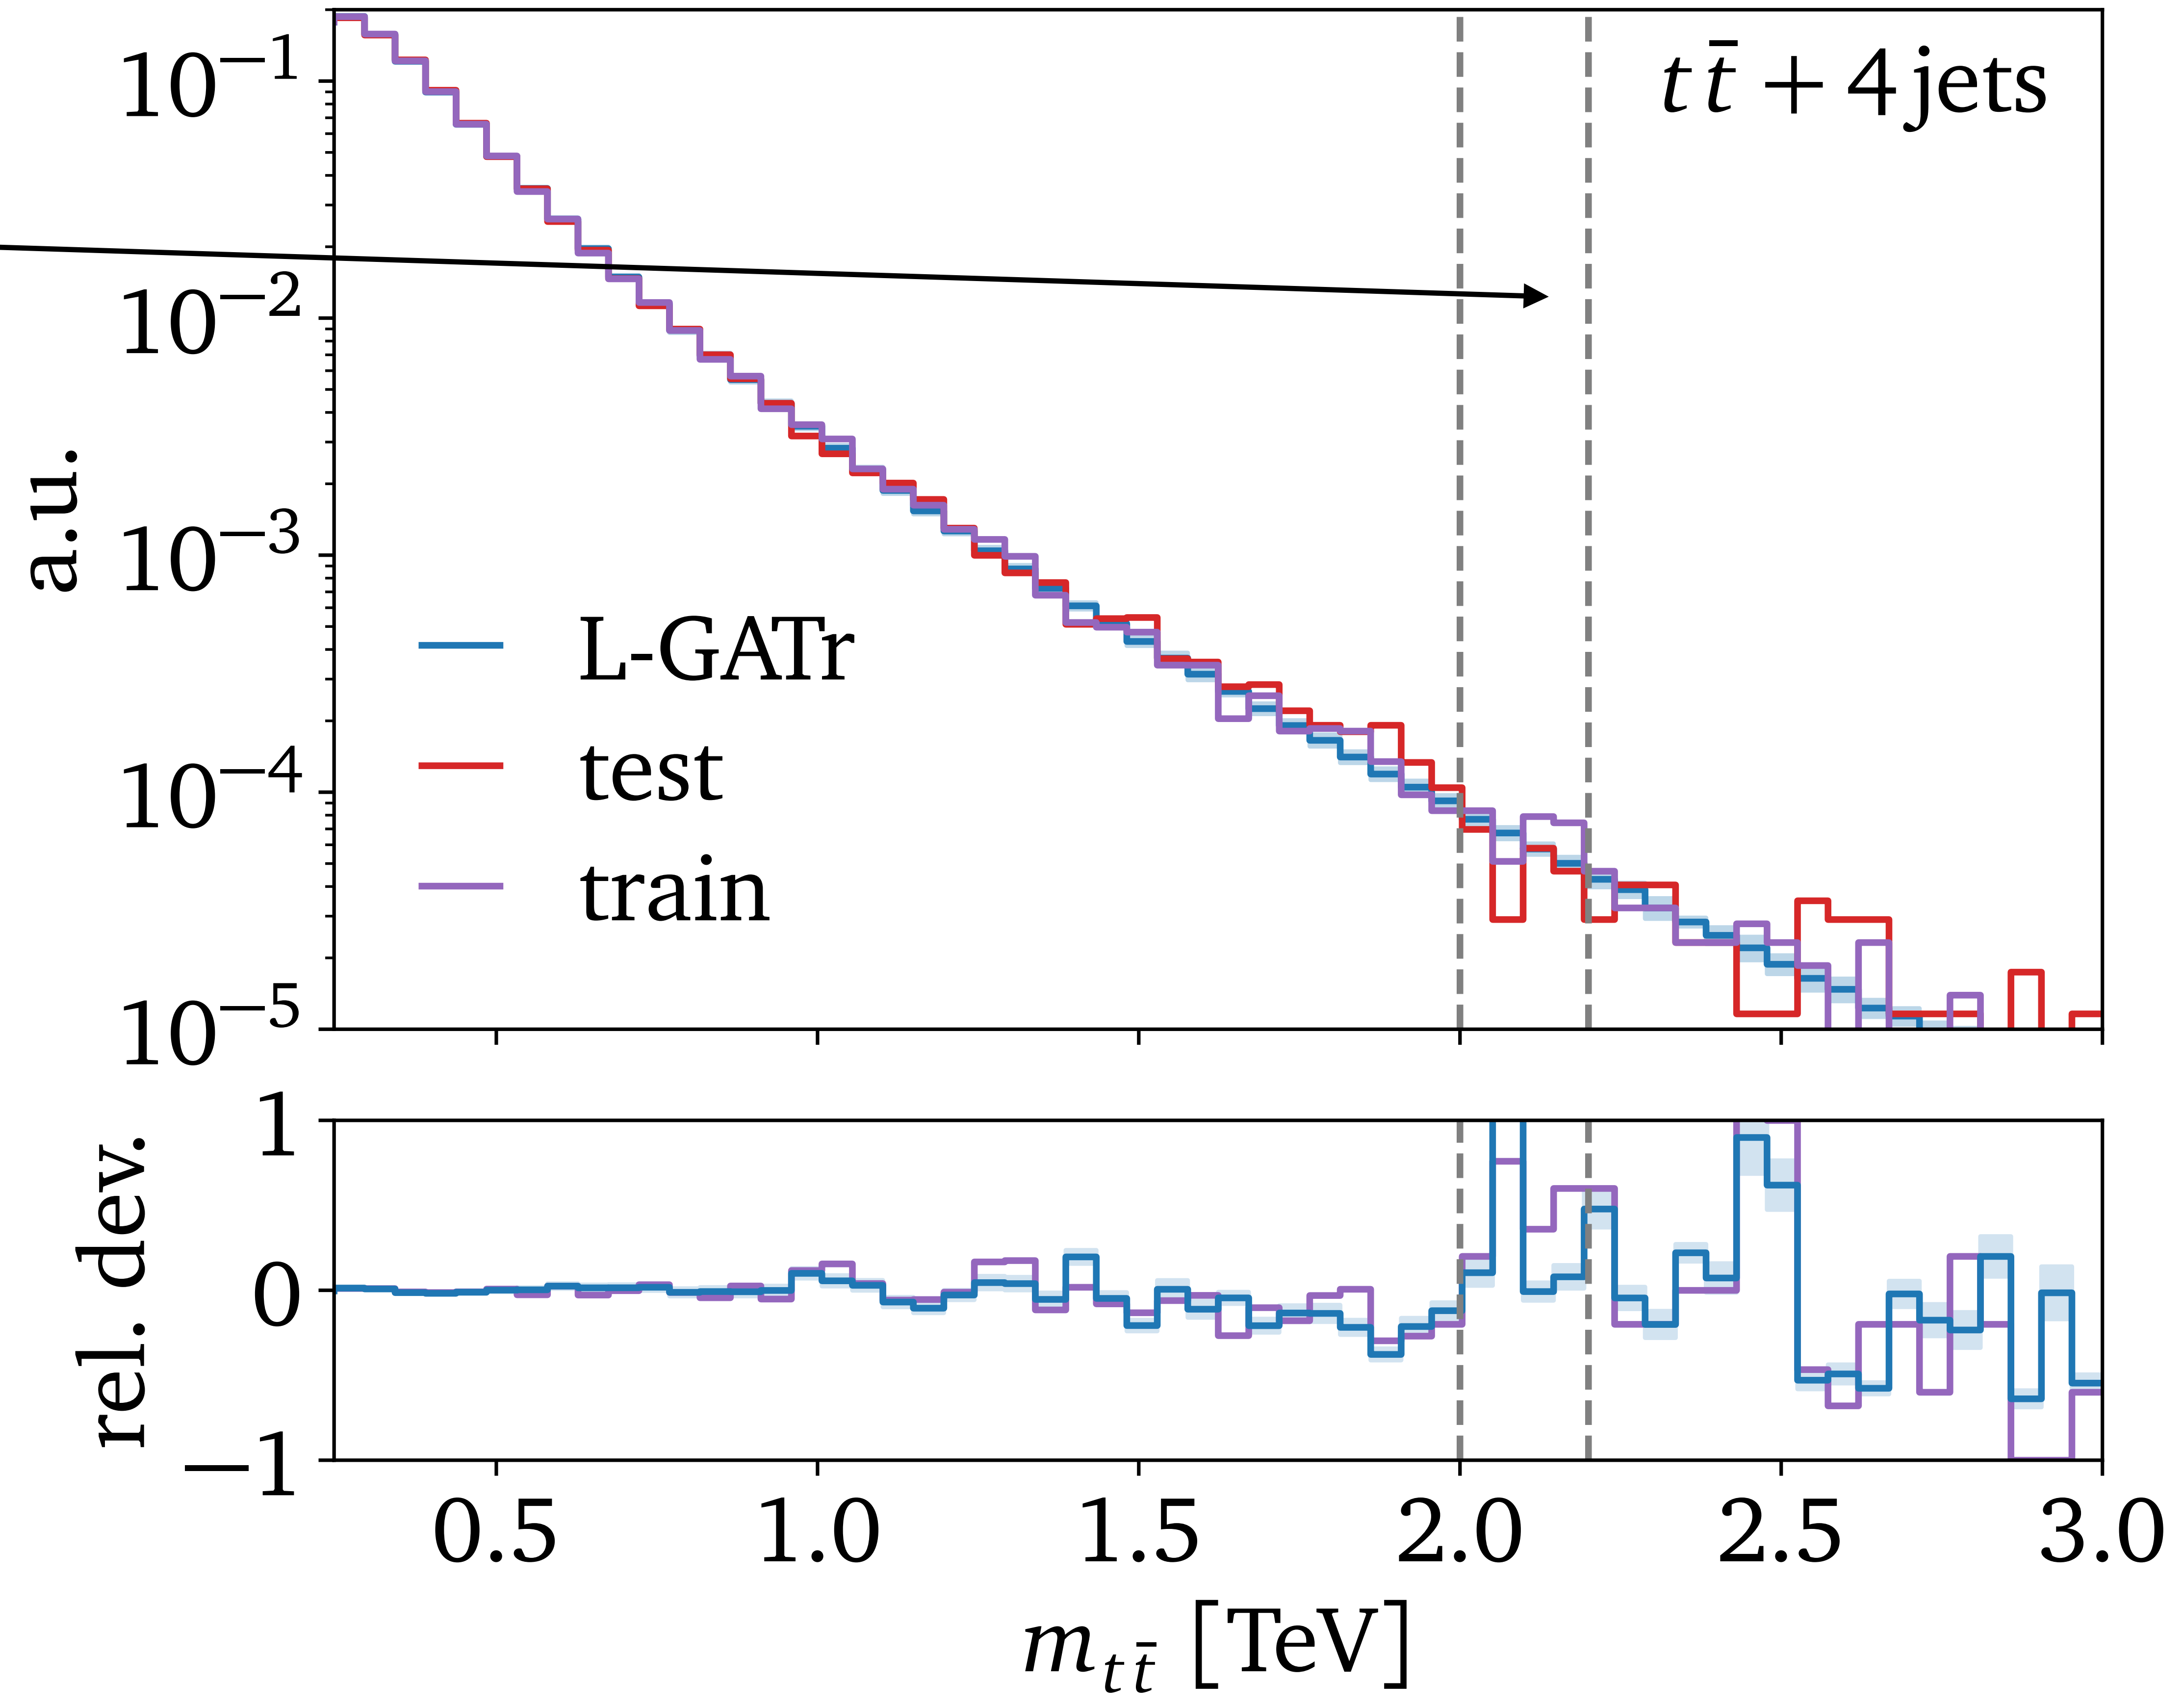

Bahl, **Diefenbacher** et. al. Forecasting Generative Amplification 2509.08048

Quantifying Amplification

Example: Di-Top + 4 jet events

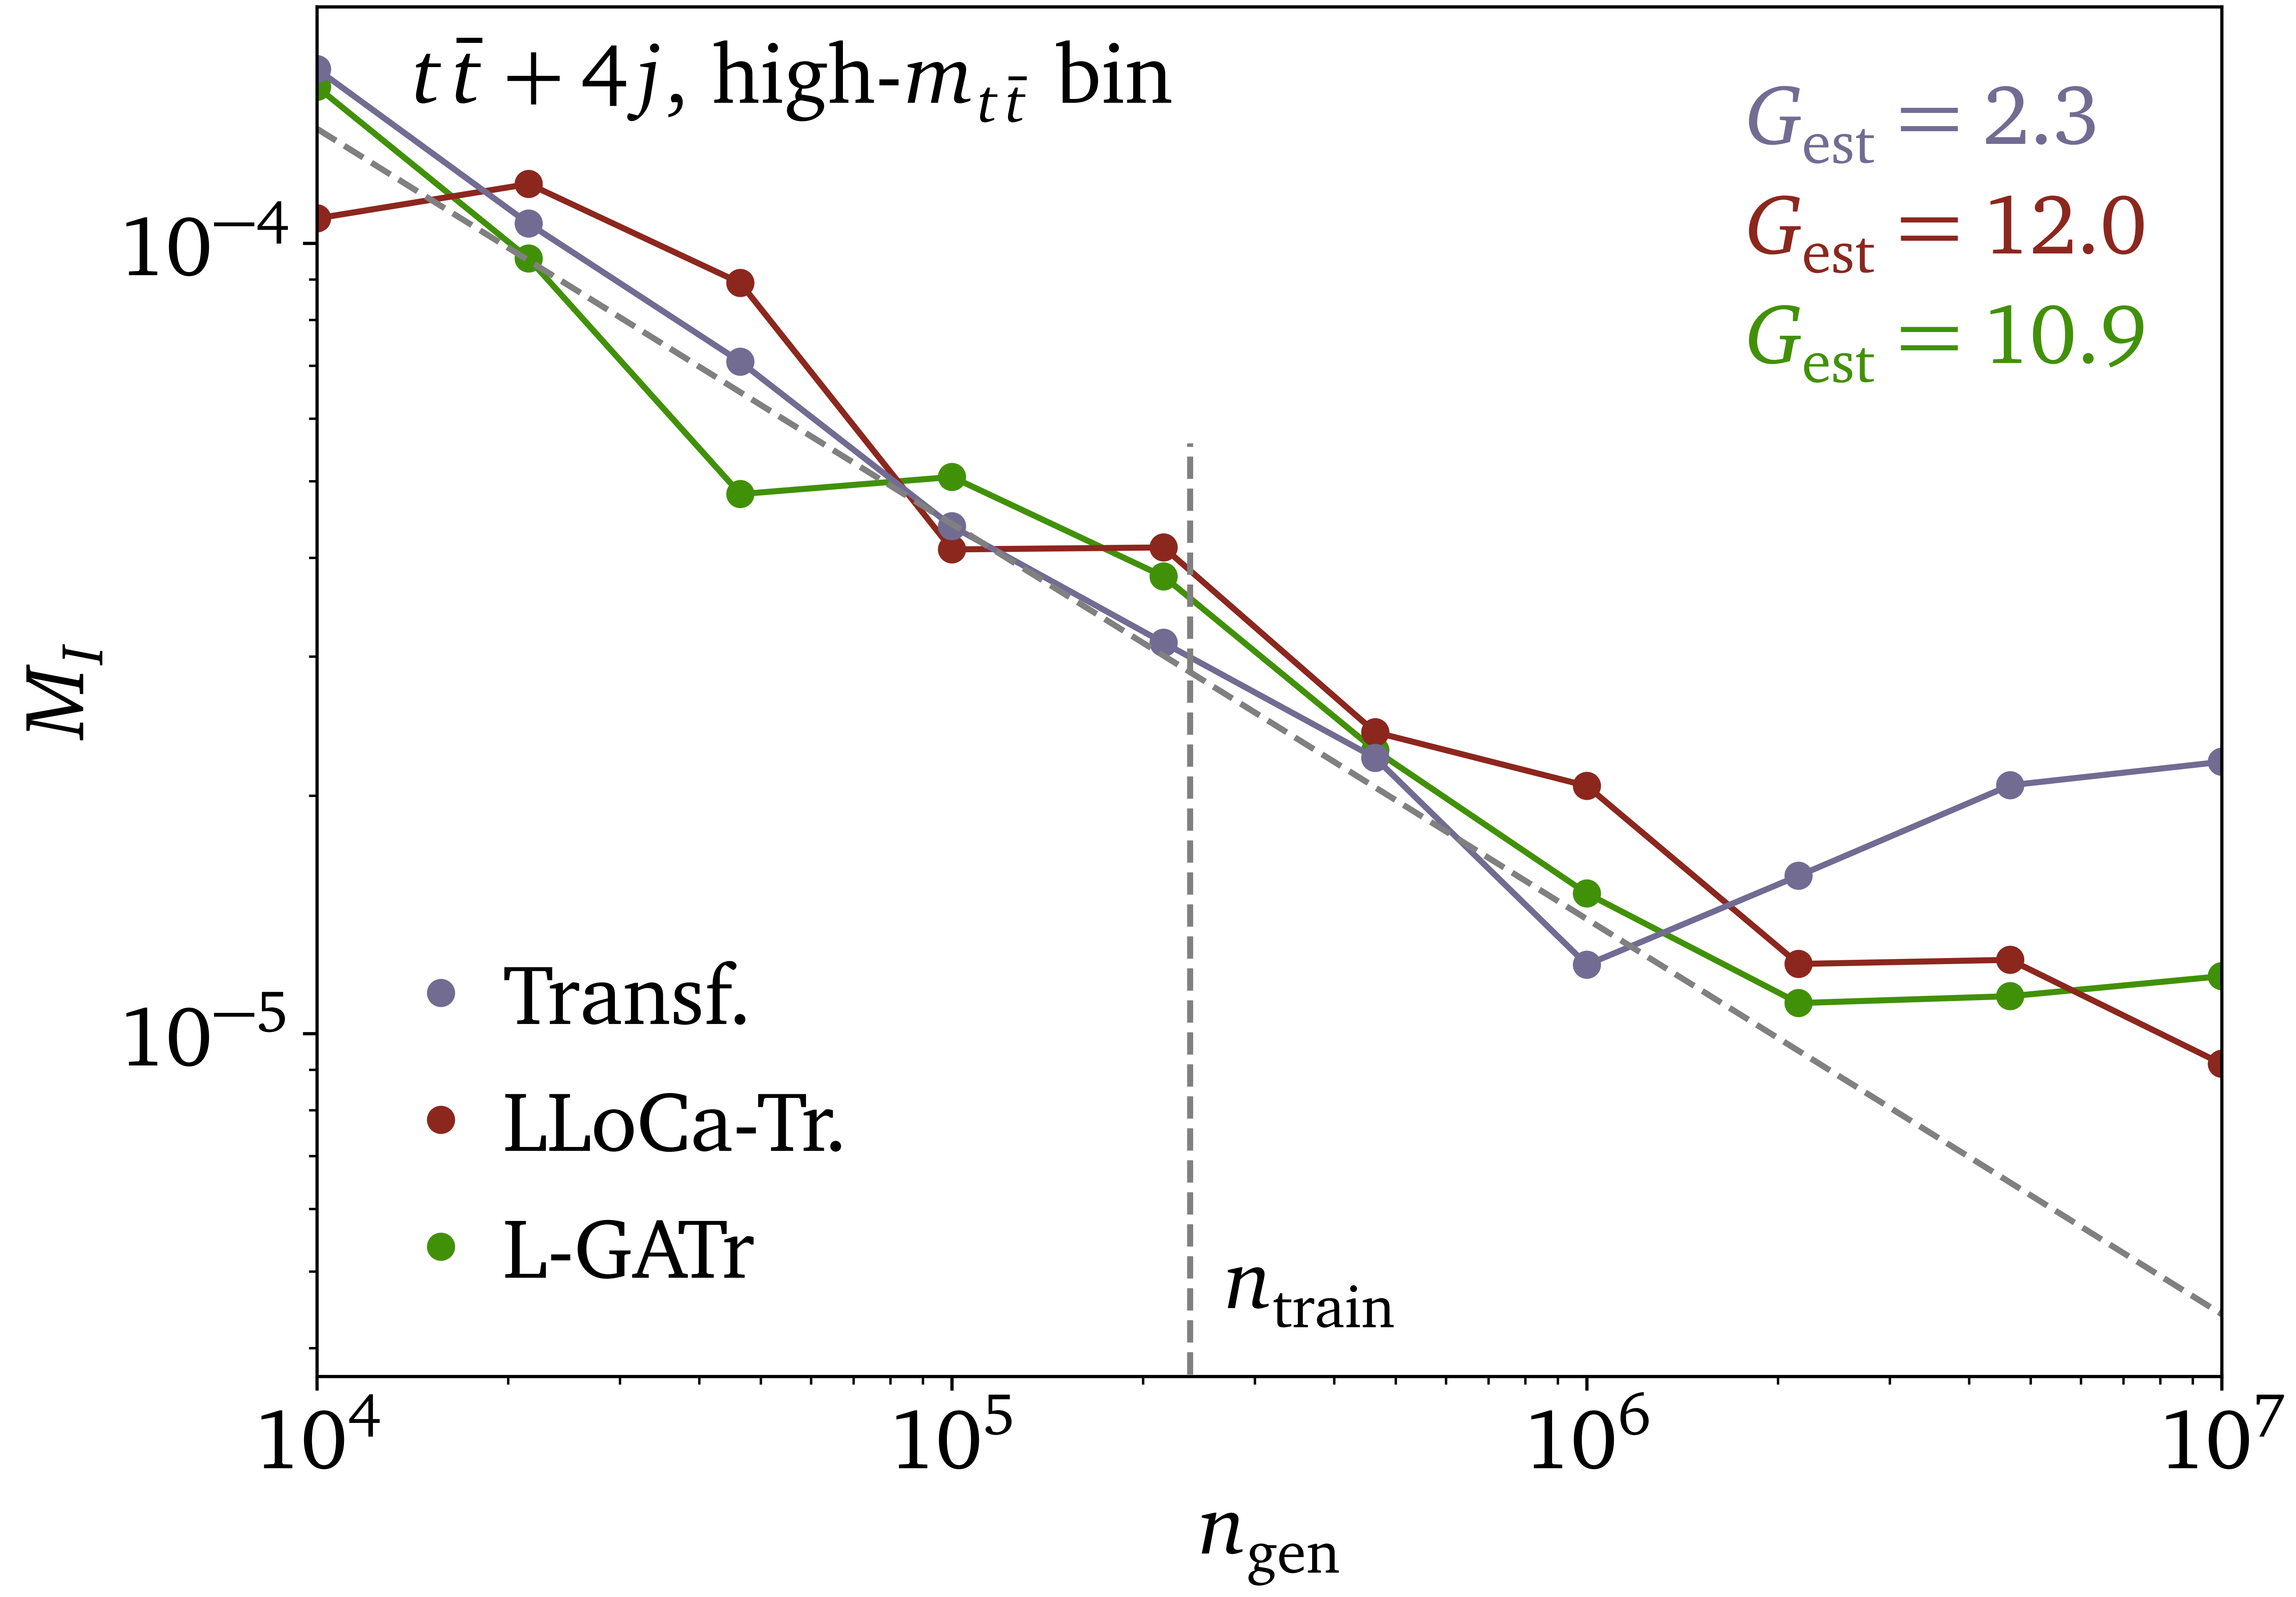

Bahl, **Diefenbacher** et al. Forecasting Generative Amplification 2509.08048

Quantifying Amplification

Second Metric: Differential Amplification

- Approximate unknown distribution with data sample

$$M\left[D^{n}_{\text{true}}, D^{n_{\text{train}}}_{\text{true}}\right] \approx M\left[D^{n}_{\text{true}}, p_{\text{true}}\right] \quad \text{with} \quad D^{n_{\text{train}}}_{\text{true}} \sim p_{\text{true}}$$

$$M\left[D^{n_{\text{gen}}}_{\text{gen}}, D^{n_{\text{train}}}_{\text{true}}\right] \approx M\left[D^{n_{\text{gen}}}_{\text{gen}}, p_{\text{true}}\right]$$

- Use two sample Kolmogorov Smirnov test
- Has analytical expectation value

$$\langle M_{\text{KS}}\left[D^{n}_{\text{true}}, D^{m}_{\text{true}}\right]\rangle = \sqrt{\frac{\pi}{2}}\log(2)\sqrt{\frac{n+m}{nm}}$$

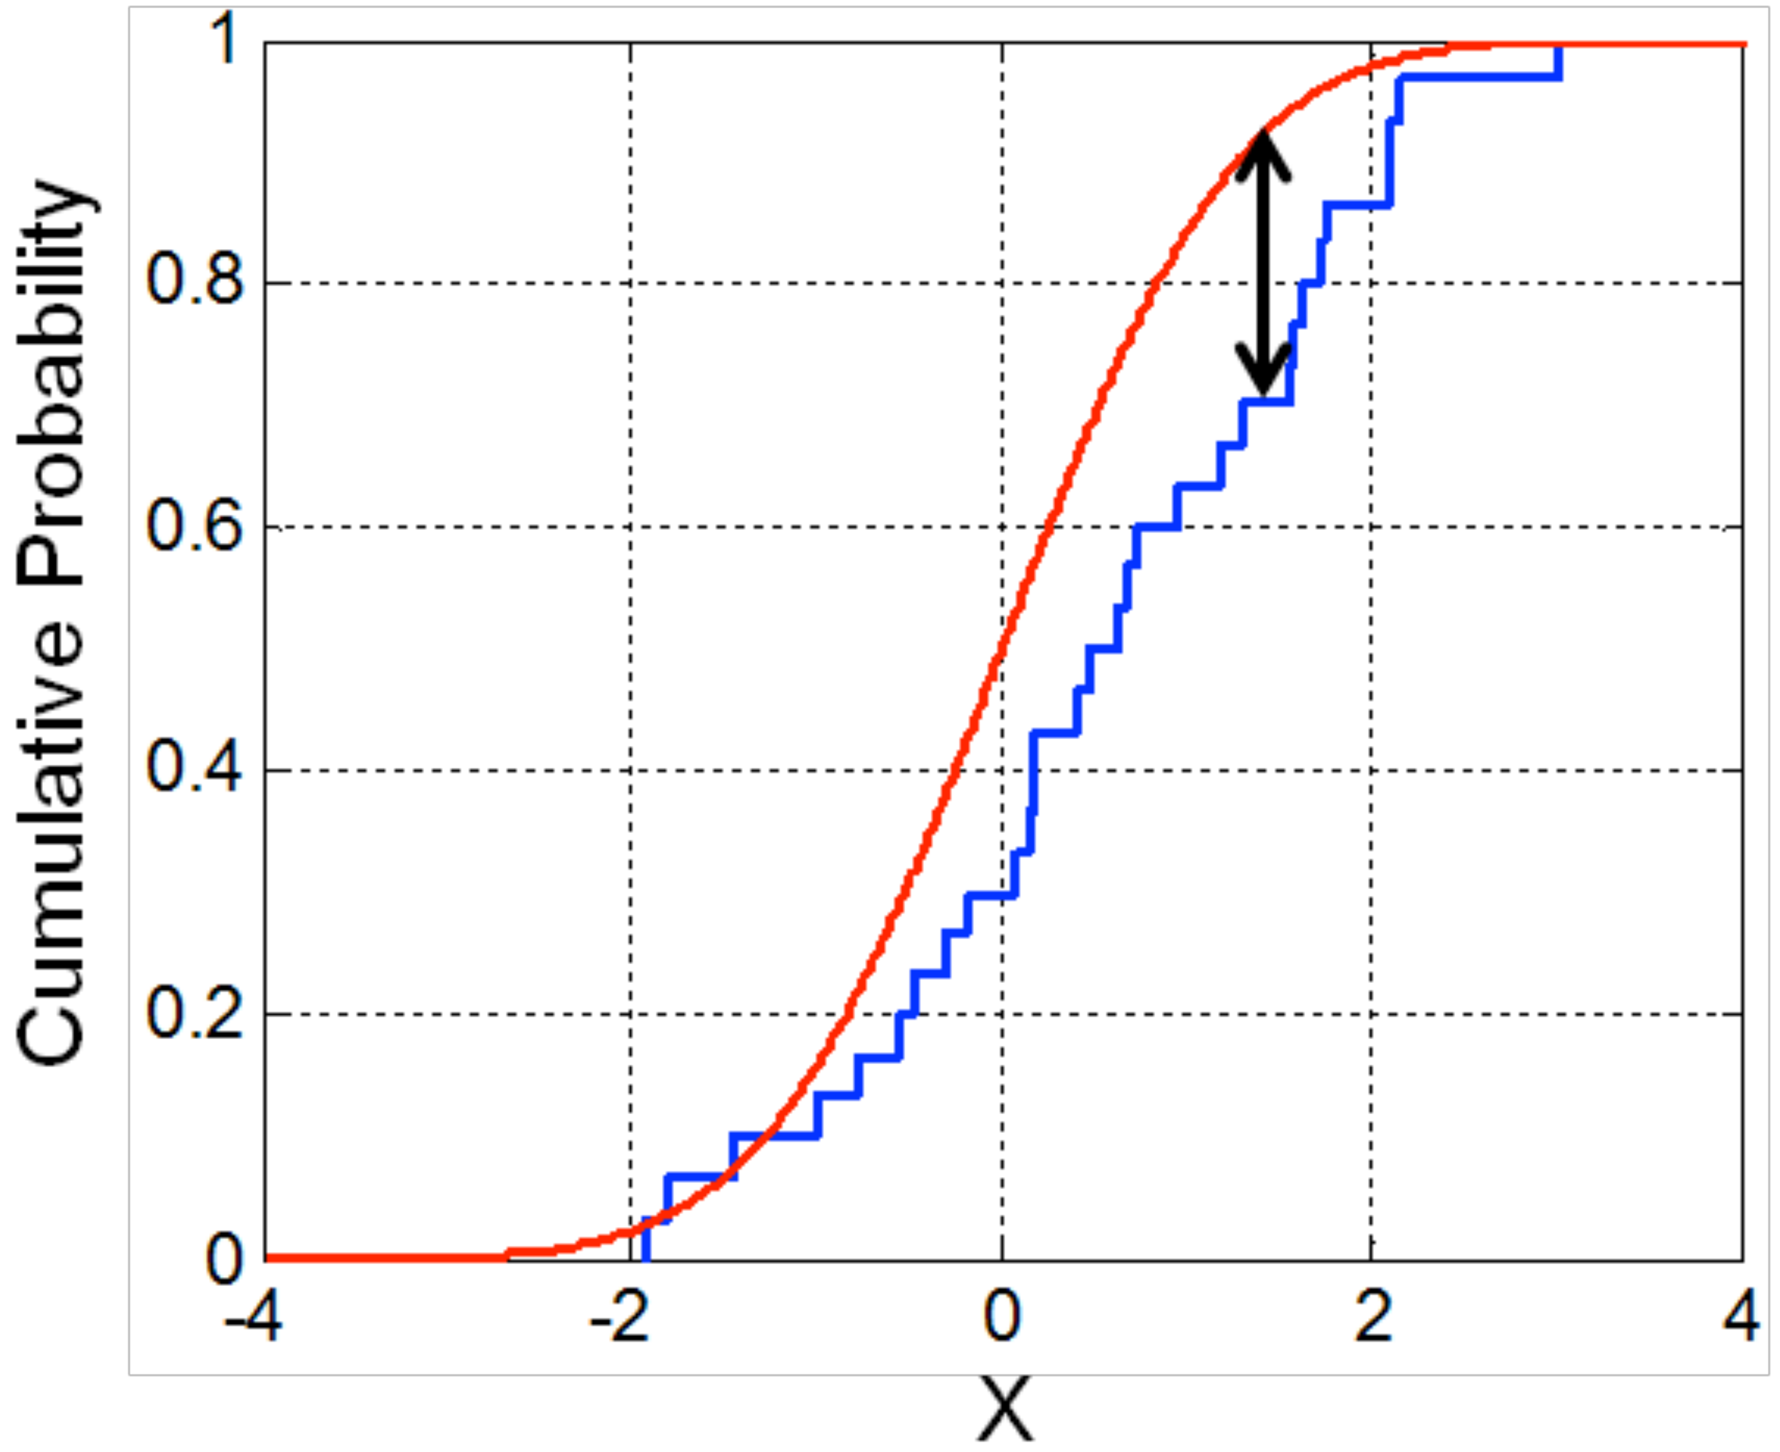

Bahl, **Diefenbacher** et al. Forecasting Generative Amplification 2509.08048

Quantifying Amplification

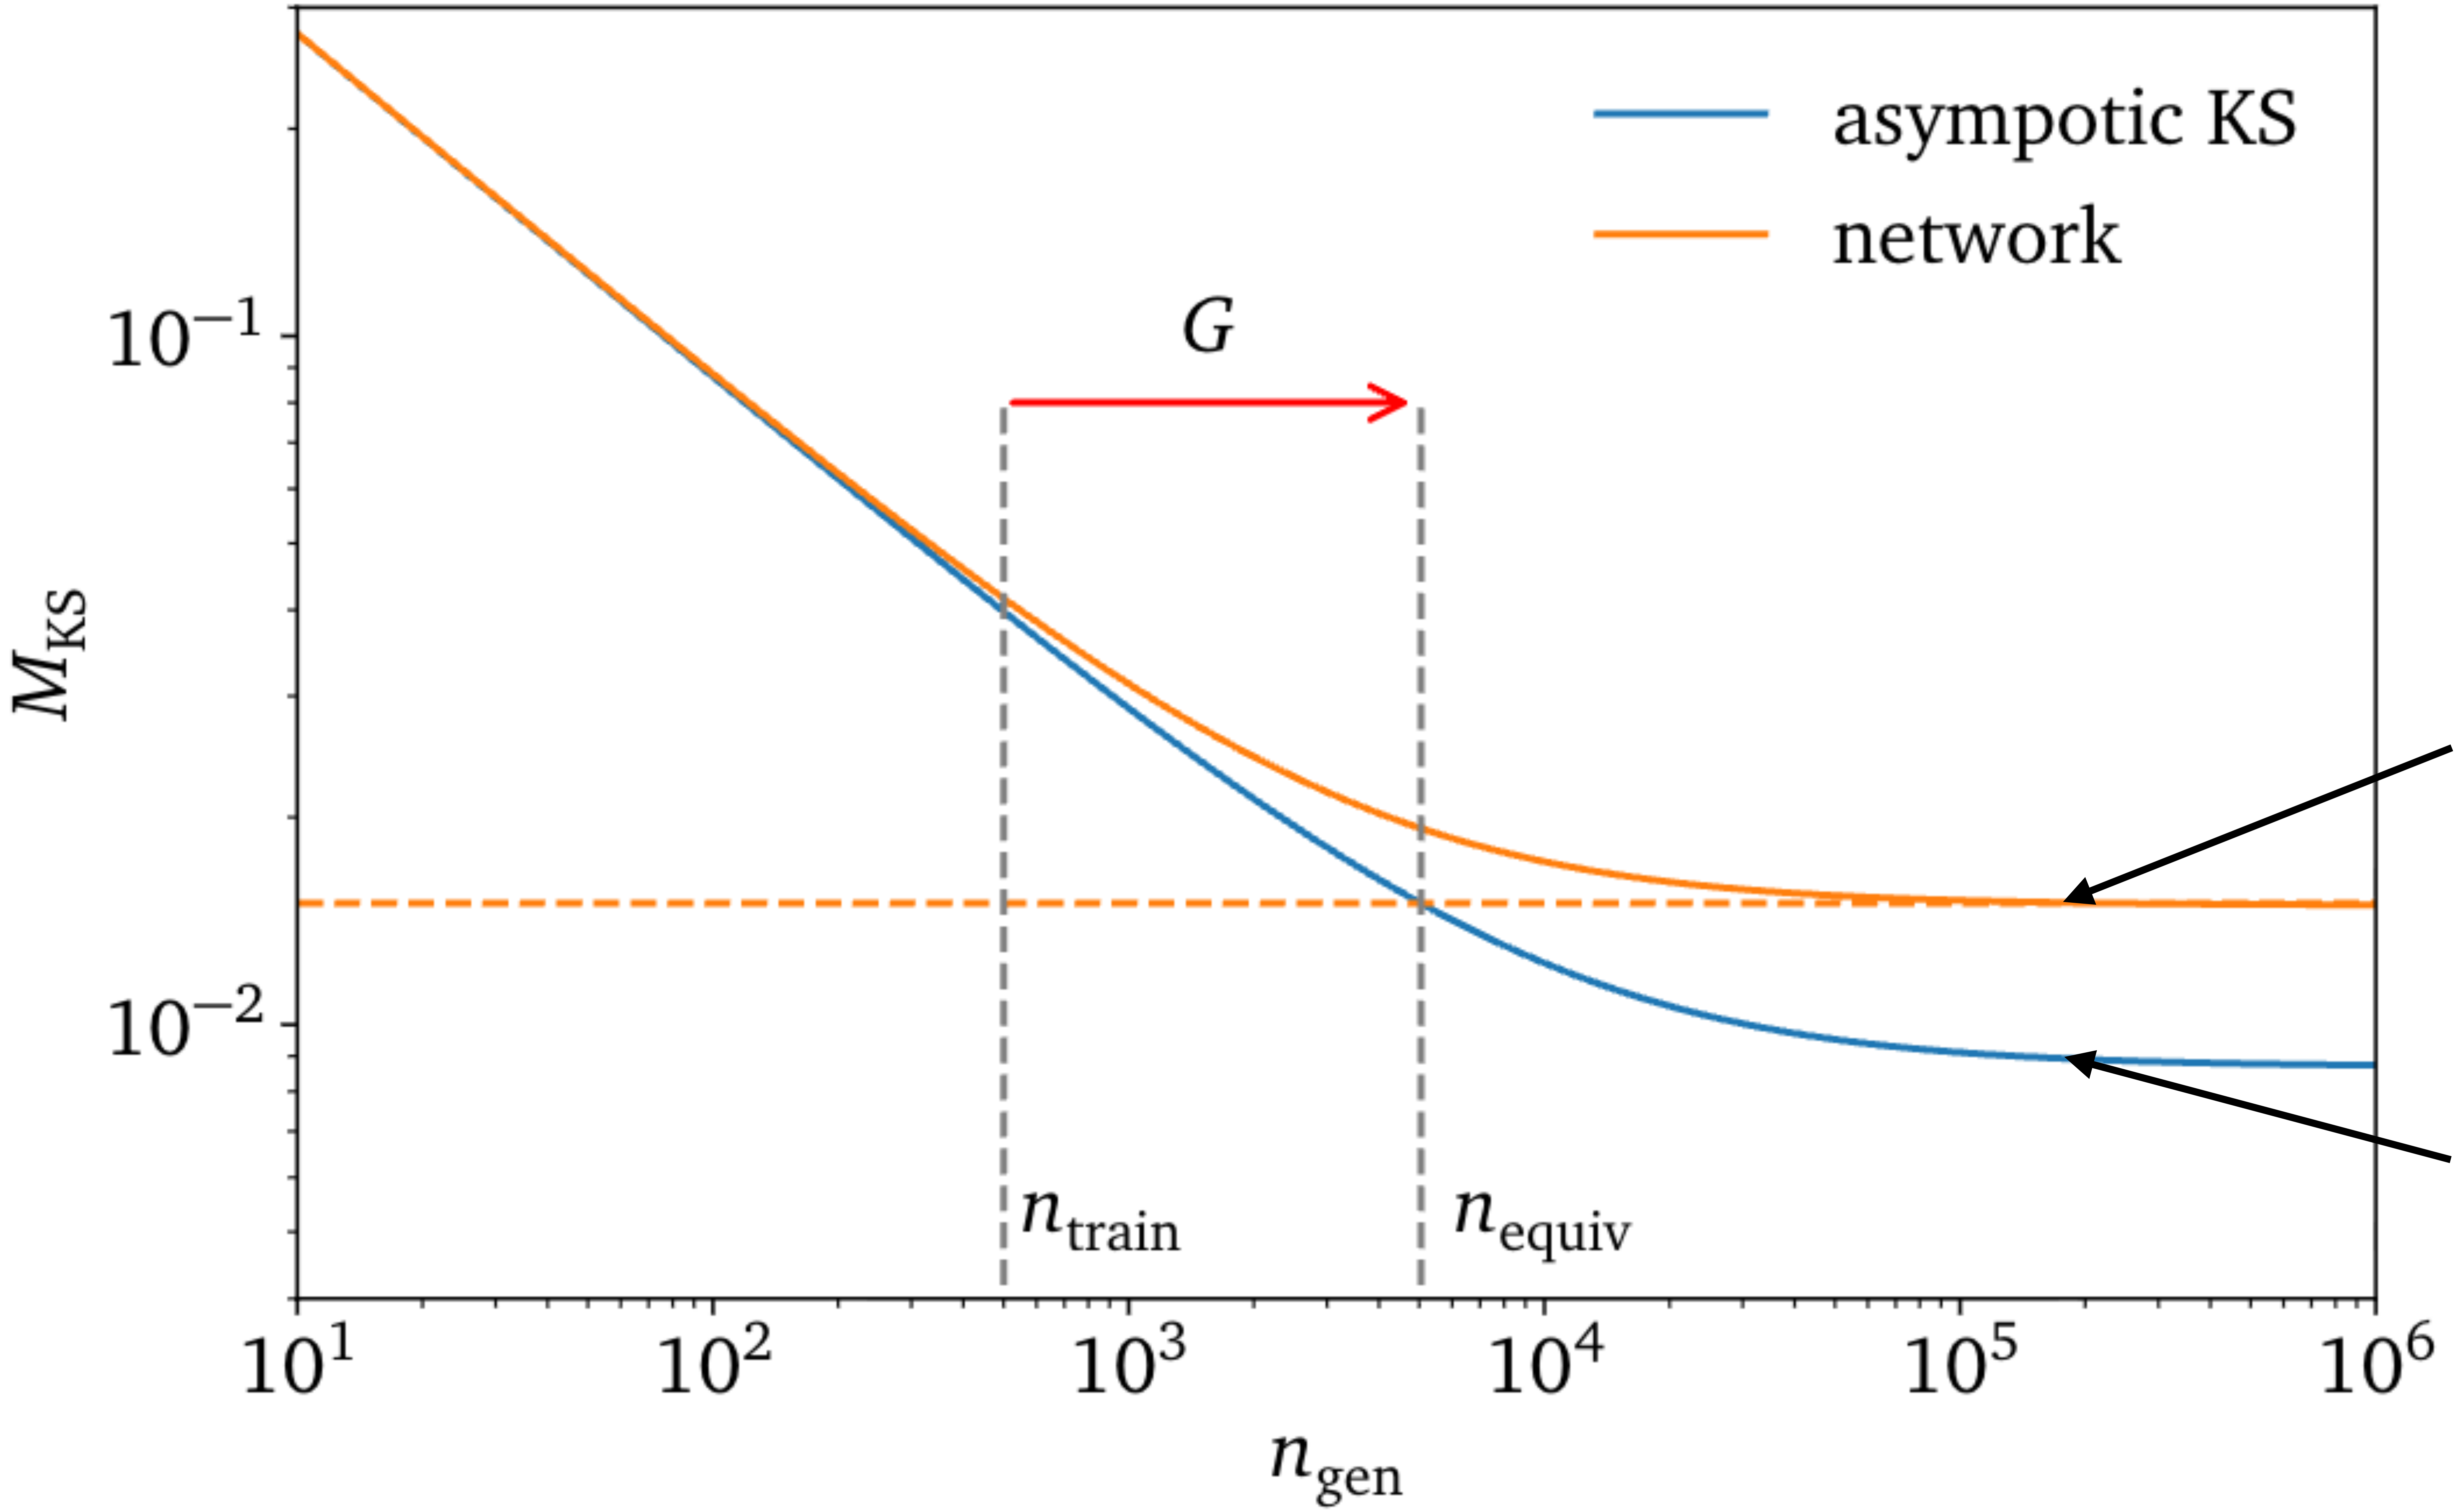

$$M_{\mathrm{KS}}\left[D_{\mathrm{gen}}^{n_{\mathrm{gen}}}, D_{\mathrm{true}}^{n_{\mathrm{train}}}\right]$$

$$\sqrt{\frac{\pi}{2}}\log(2)\sqrt{\frac{n_{\mathrm{gen}}+n_{\mathrm{train}}}{n_{\mathrm{gen}}n_{\mathrm{train}}}}$$

Bahl, **Diefenbacher** et al. Forecasting Generative Amplification 2509.08048

Example: Di-Top + 4 jet events with both methods

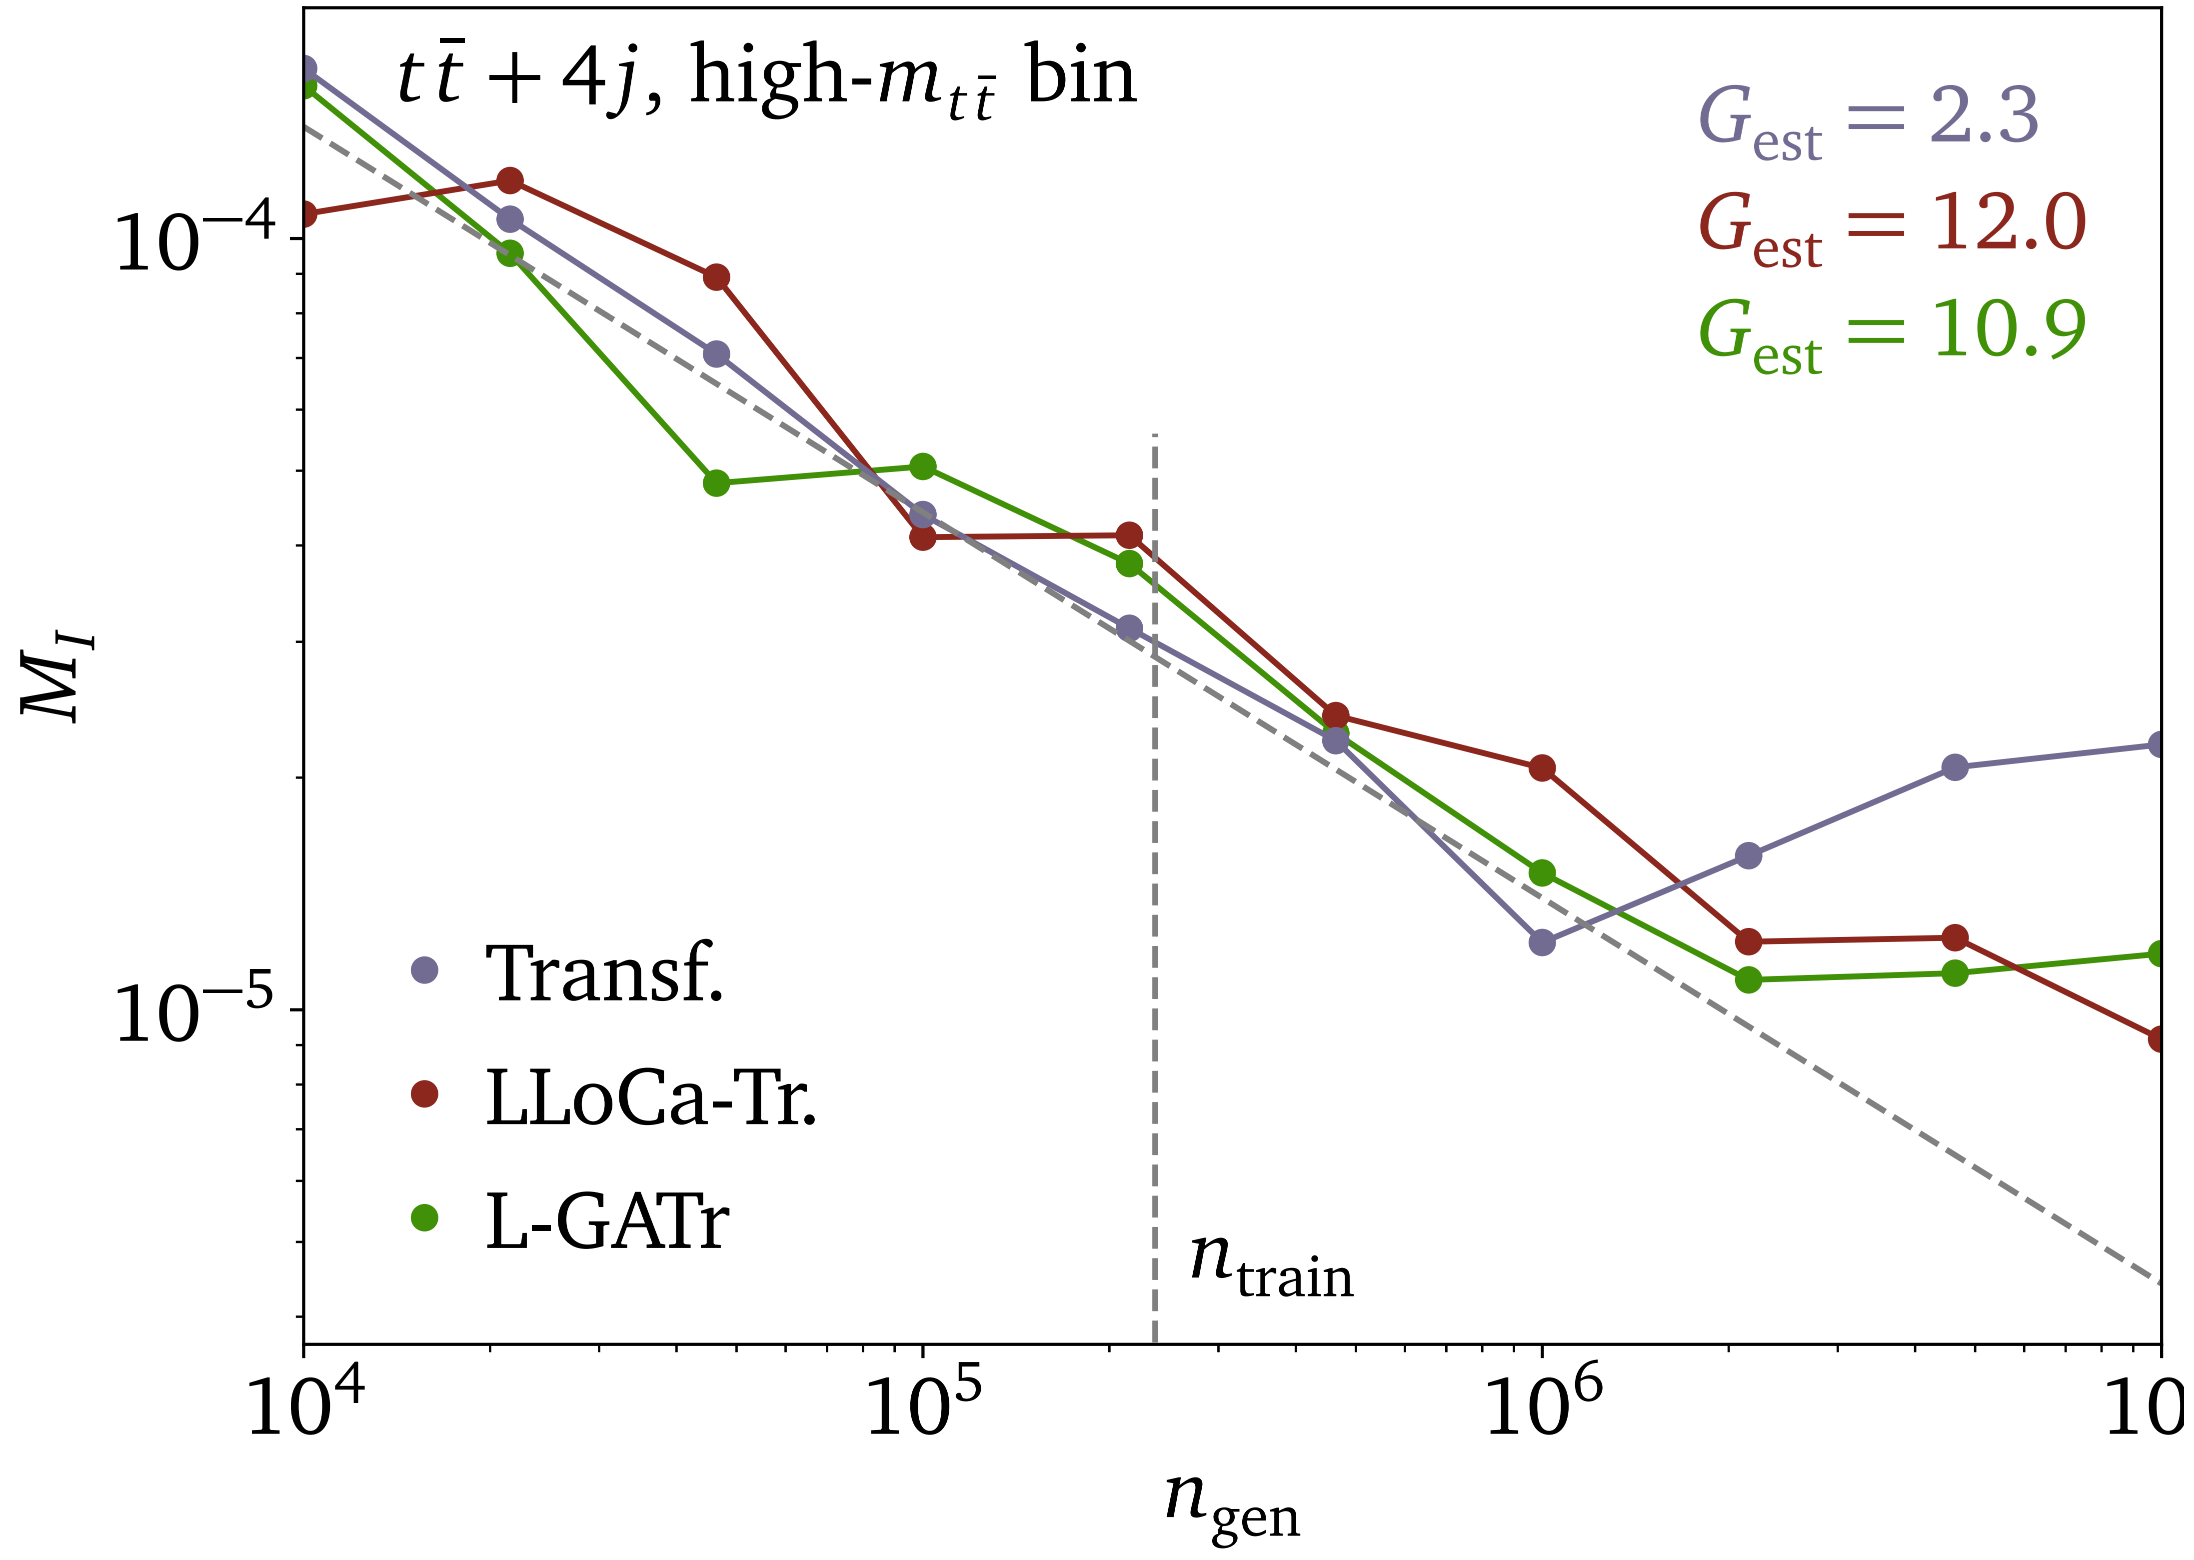

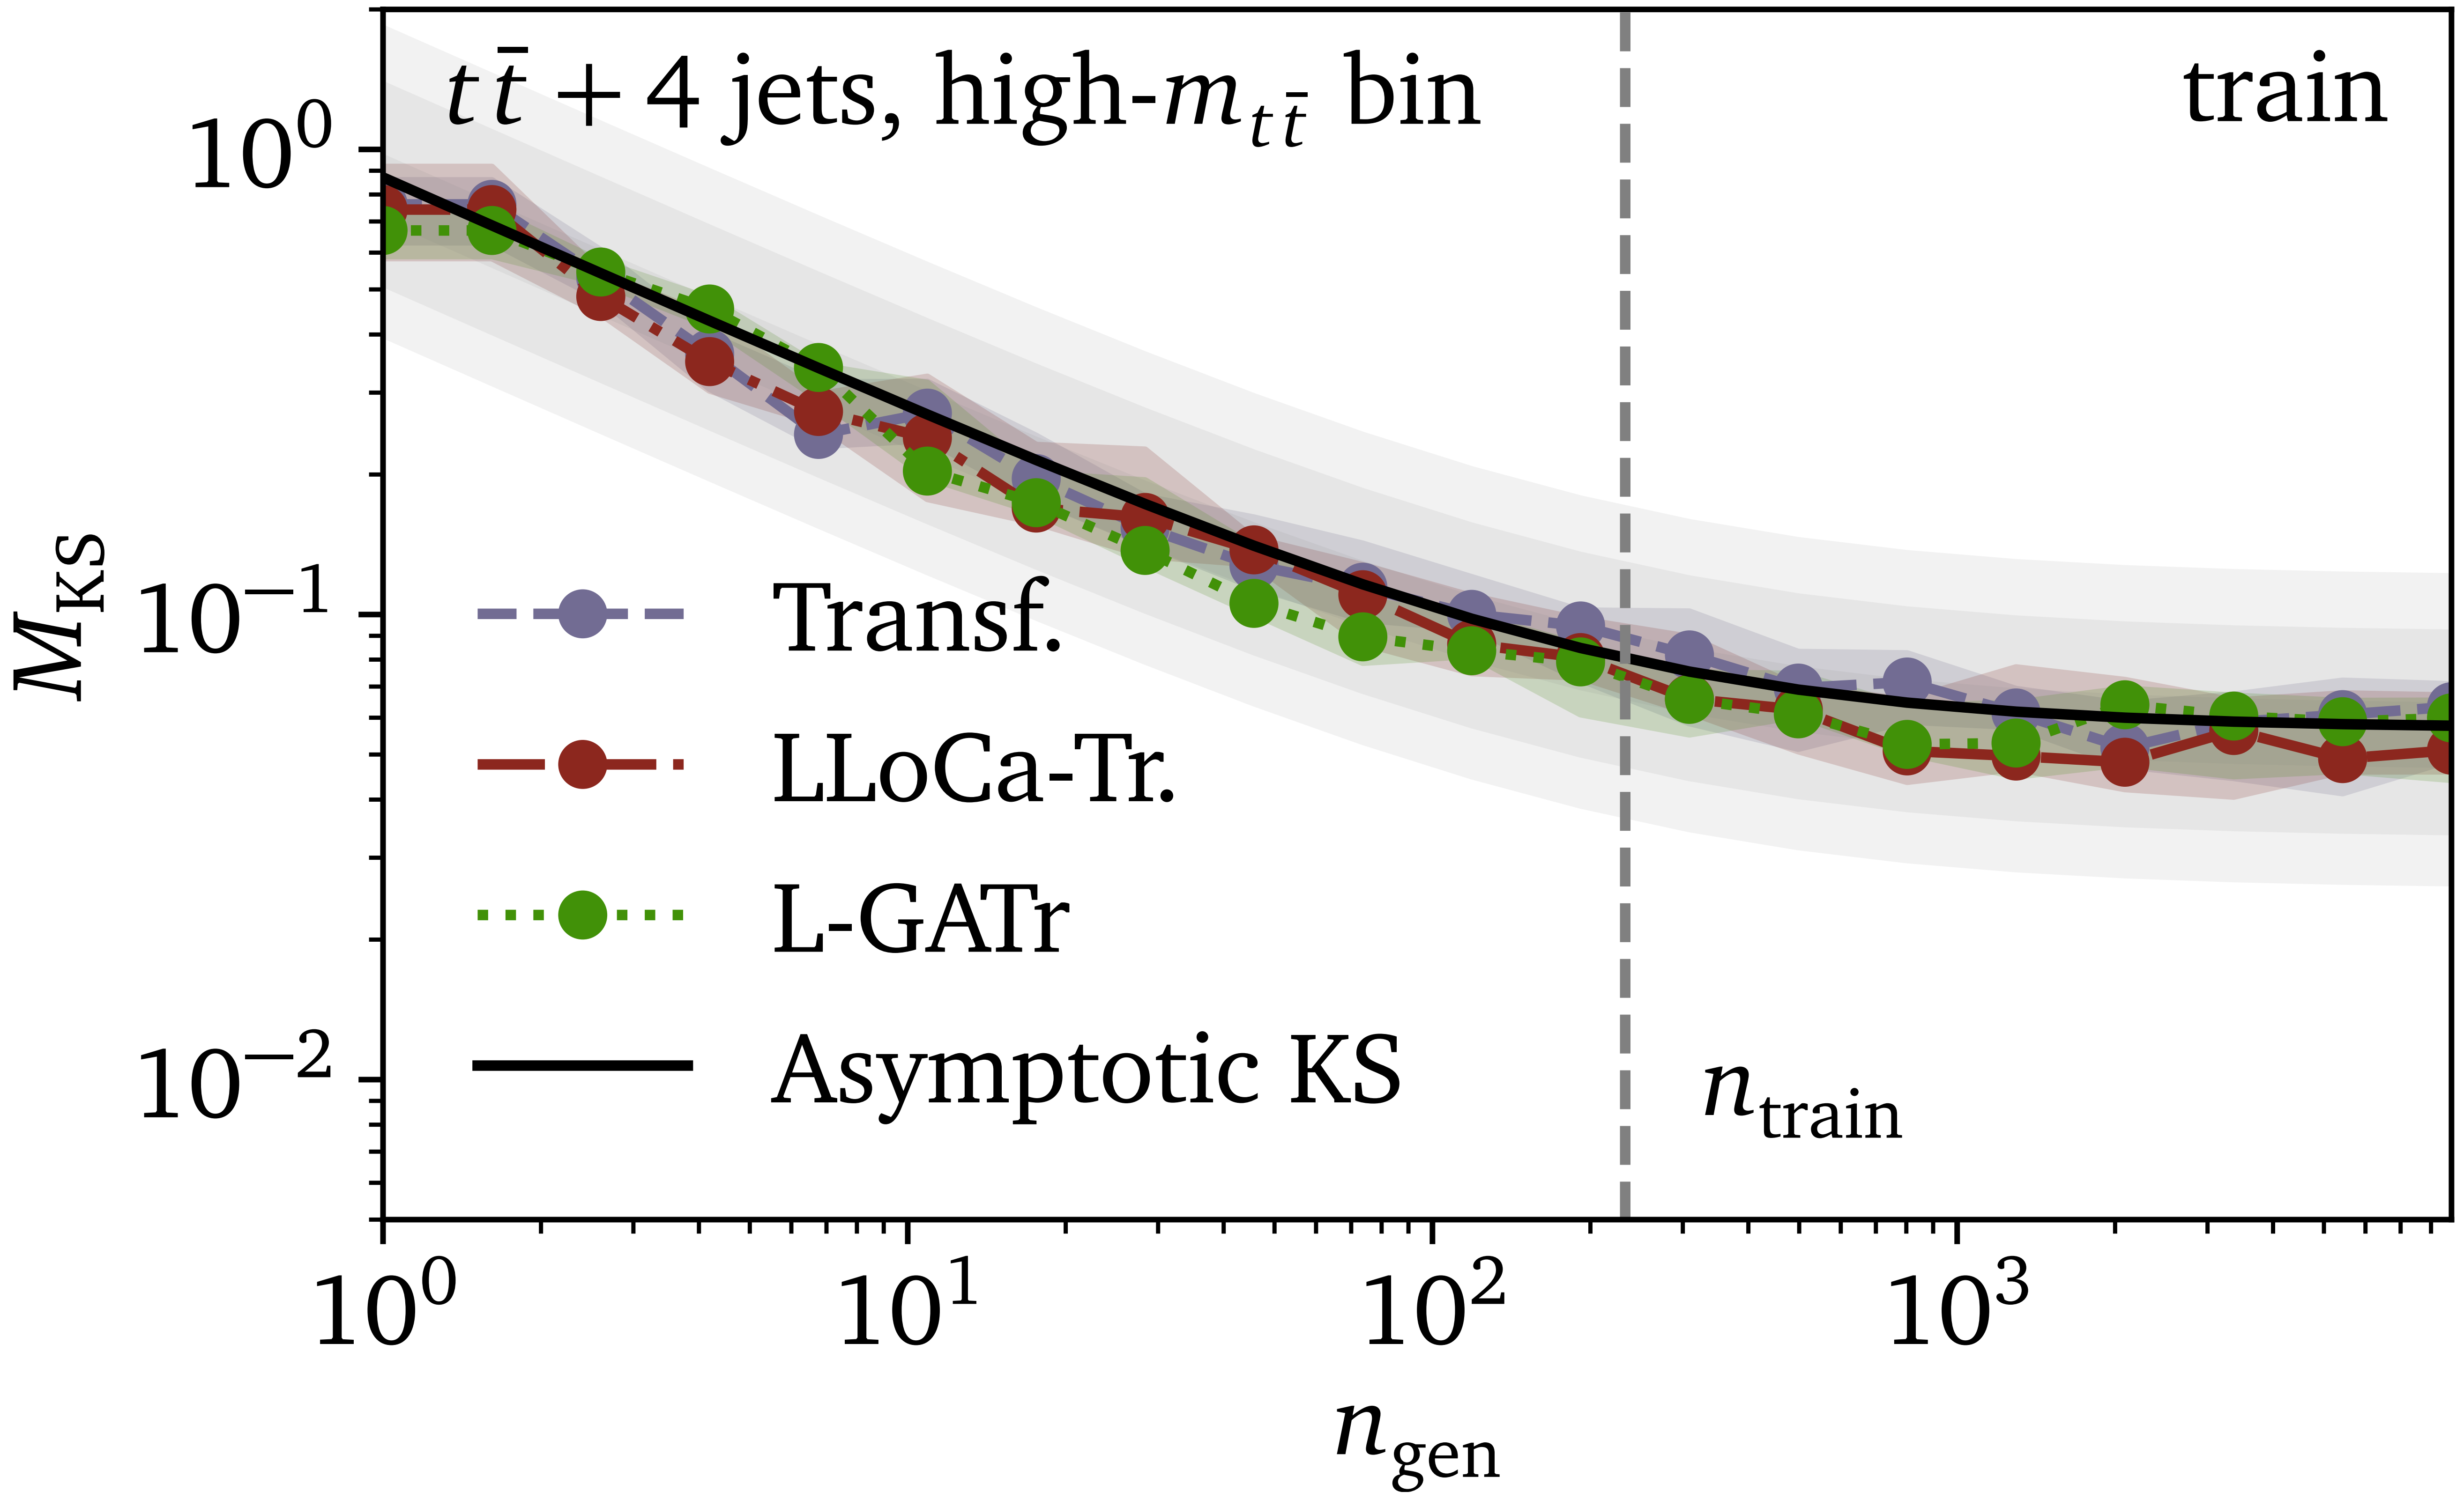

Bahl, **Diefenbacher** et al. Forecasting Generative Amplification 2509.08048

Conclusion

- **Amplification Investigation**
  - Generative networks can sample beyond training data
  - Quantify by how much
  - Core to generative uncertainties
- **Enables generative approaches in HEP**
  - Amplification key for generative simulations
  - Simulation-based inference, generative unfolding need uncertainties

Conclusion

- **ML methods bring unprecedented precision**
  - LHC environment highly complex
  - $t\bar{t}H$, CP violation measurements impossible without
  - Directly enabled by ML

- **Fundamental physics understanding**
  - $t\bar{t}H$: early universe Higgs potential
  - CP violation: matter-antimatter asymmetry, baryogenesis

Thank You

Backup

Quantifying Amplification

- **Example Metric: Averaging Amplification**
  - Fractions inside phase space region V:

$$I = \int_V dx\, p_{\text{true}}(x)$$

$$\bar{I} = \frac{1}{n_{\text{gen}}} \sum_{x \in D_{\text{gen}}^{n_{\text{gen}}}} \mathbf{1}_{x \in V}\,,$$

  - Expected difference as metric

$$M\left[D_{\text{gen}}^{n_{\text{gen}}}, p_{\text{true}}\right] = \left\langle \left[I(p_{\text{true}}) - \bar{I}(D_{\text{gen}}^{n_{\text{gen}}})\right]^2 \right\rangle$$

  - Estimate via Bayesian Network variance in V

Radius of Gaussian Ring

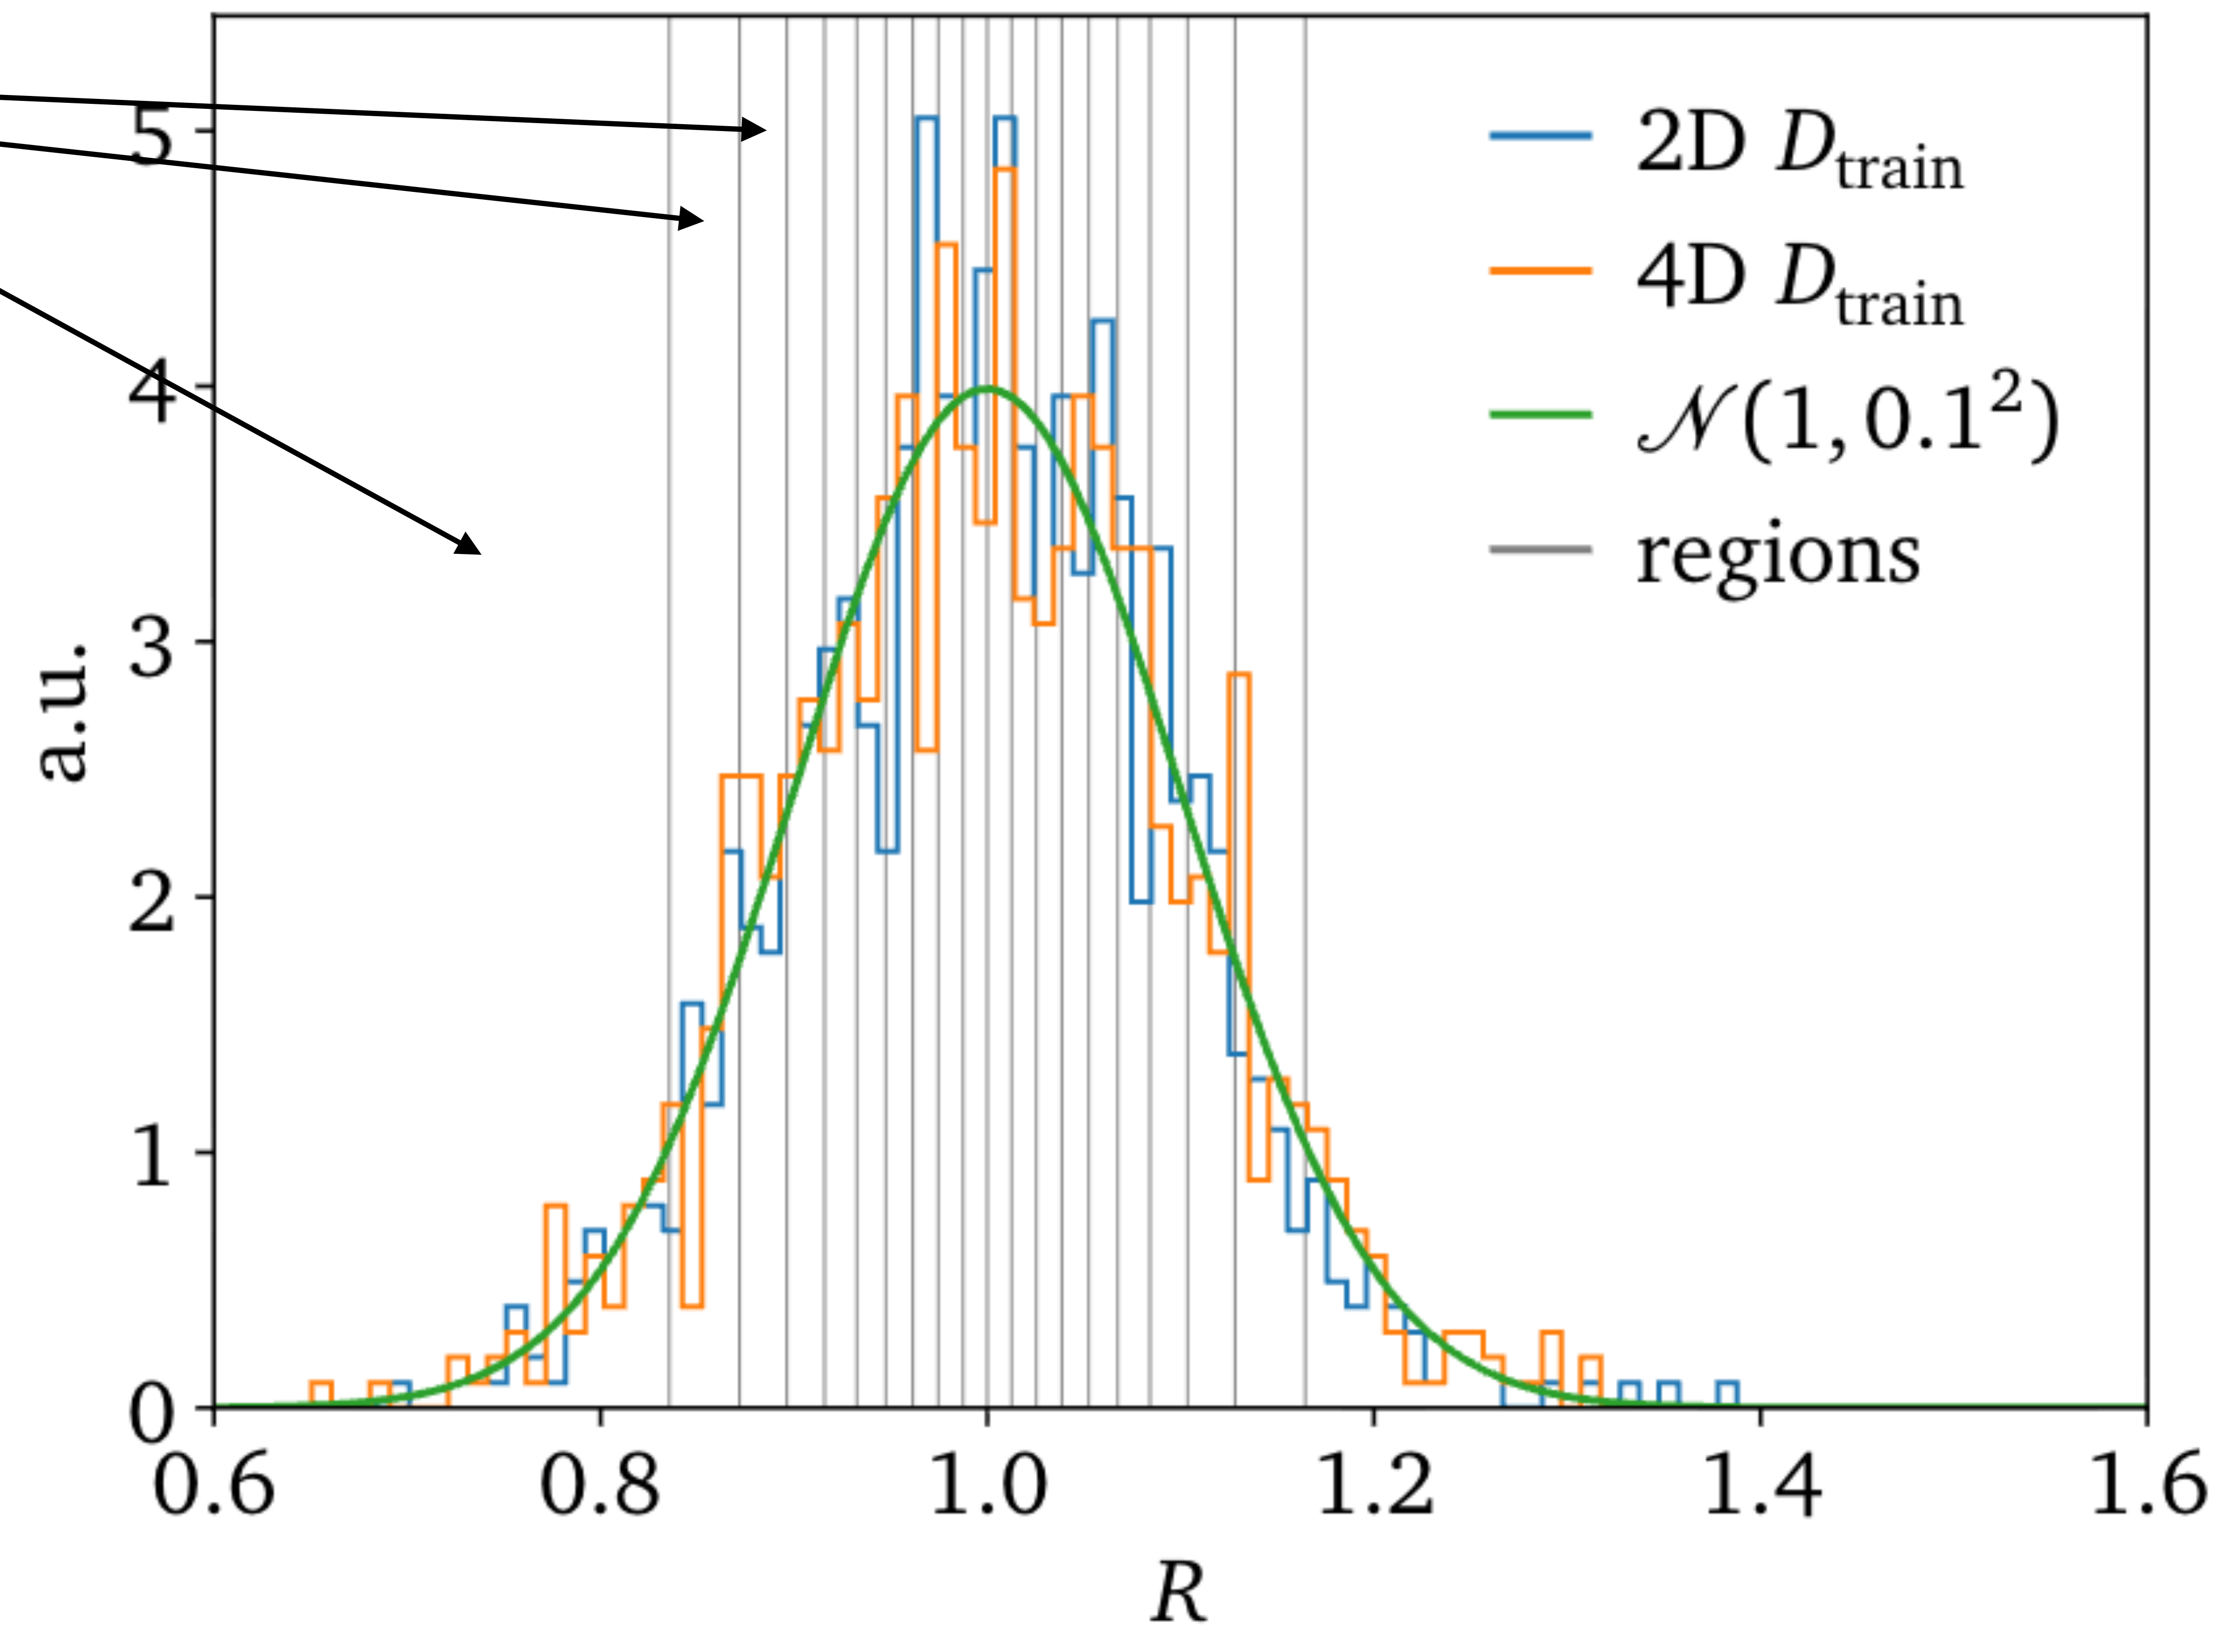

Bahl, **Diefenbacher** et. al. Forecasting Generative Amplification 2509.08048

Quantifying Amplification

- **Example: Gaussian Ring Radius**
  - Plot $M_I\left[D_{\text{gen}}^{n\text{gen}}, p_{\text{data}}\right](n_{\text{gen}})$
  - Get $\sigma_{\text{stat}}(n)$ from falling section
  - Get σ_{model} from asymptotic part
  - Combine regions V_i

$$\sigma_{\text{stat}}^2 = \sum_i^{n_v} \sigma_{\text{stat},V_i}^2 \quad \sigma_{\text{model}}^2 = \sum_i^{n_v} \sigma_{\text{model},V_i}^2$$

$$\sigma_{\text{stat}}(n_{\text{equiv}}) \stackrel{!}{=} \sigma_{\text{model}}$$

  - Estimate G, agrees with truth

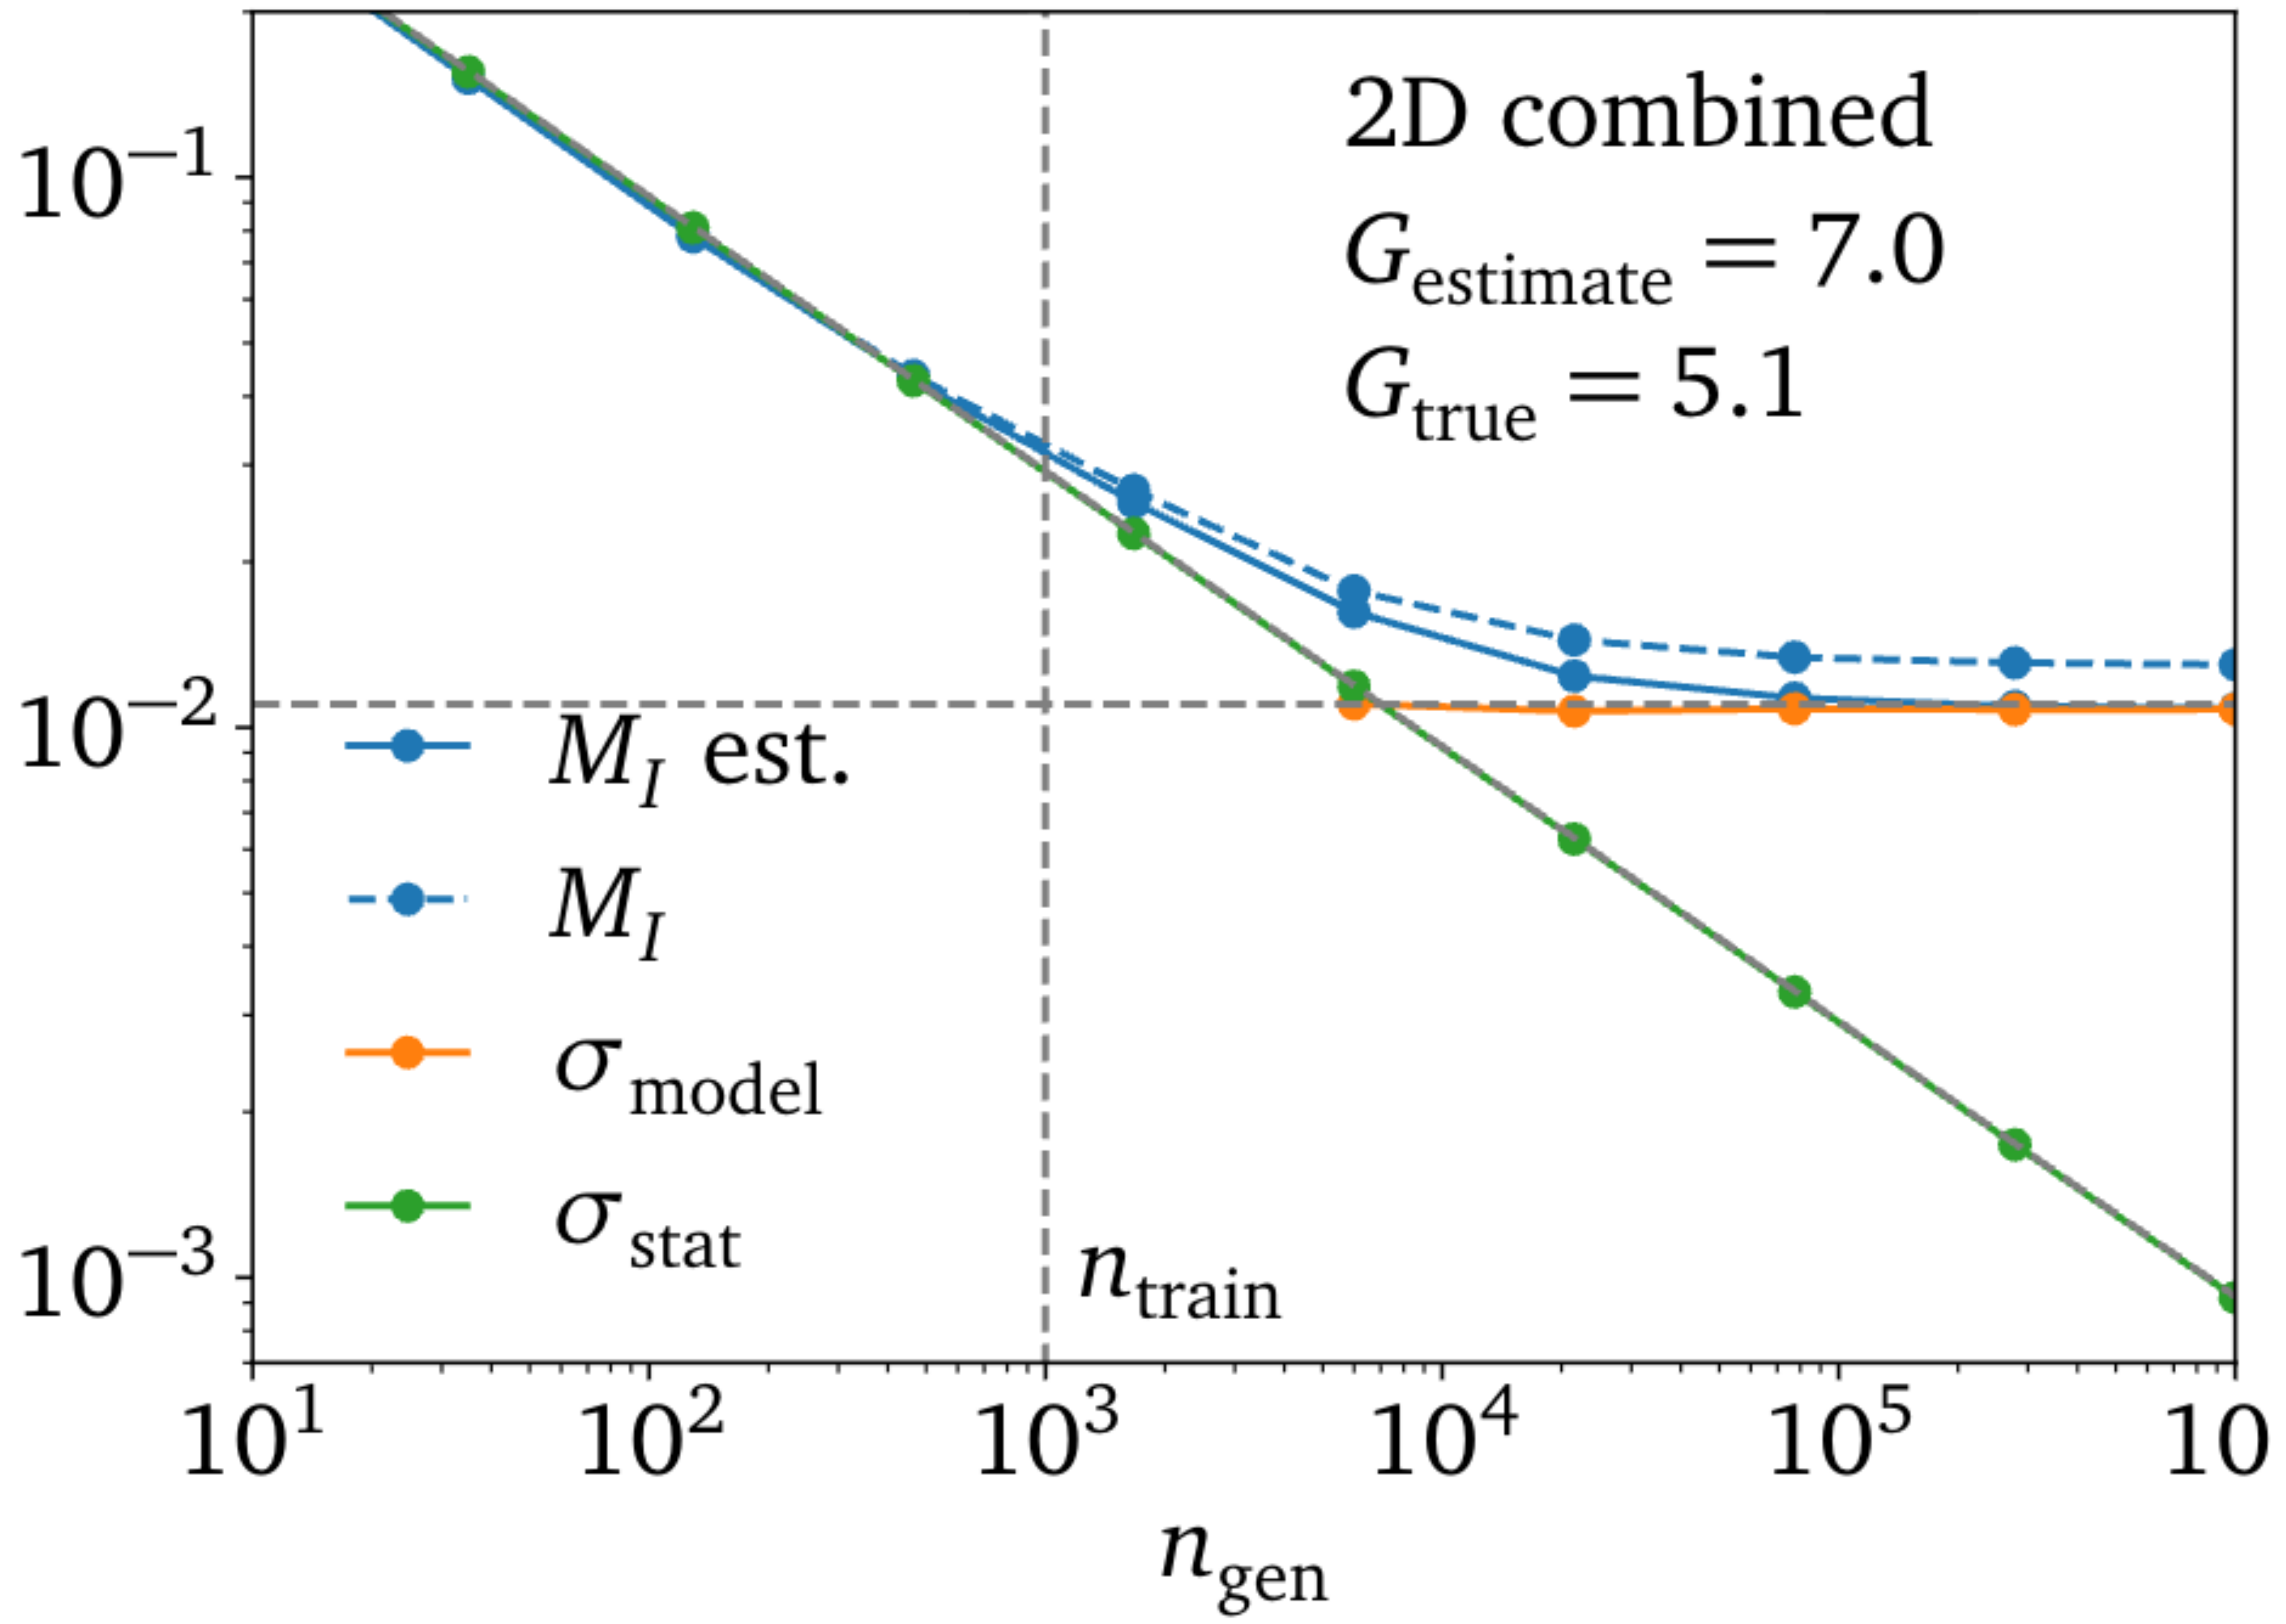

Bahl, **Diefenbacher** et. al. Forecasting Generative Amplification 2509.08048

Quantifying Amplification

- Example: Di-Top + 4 jet events

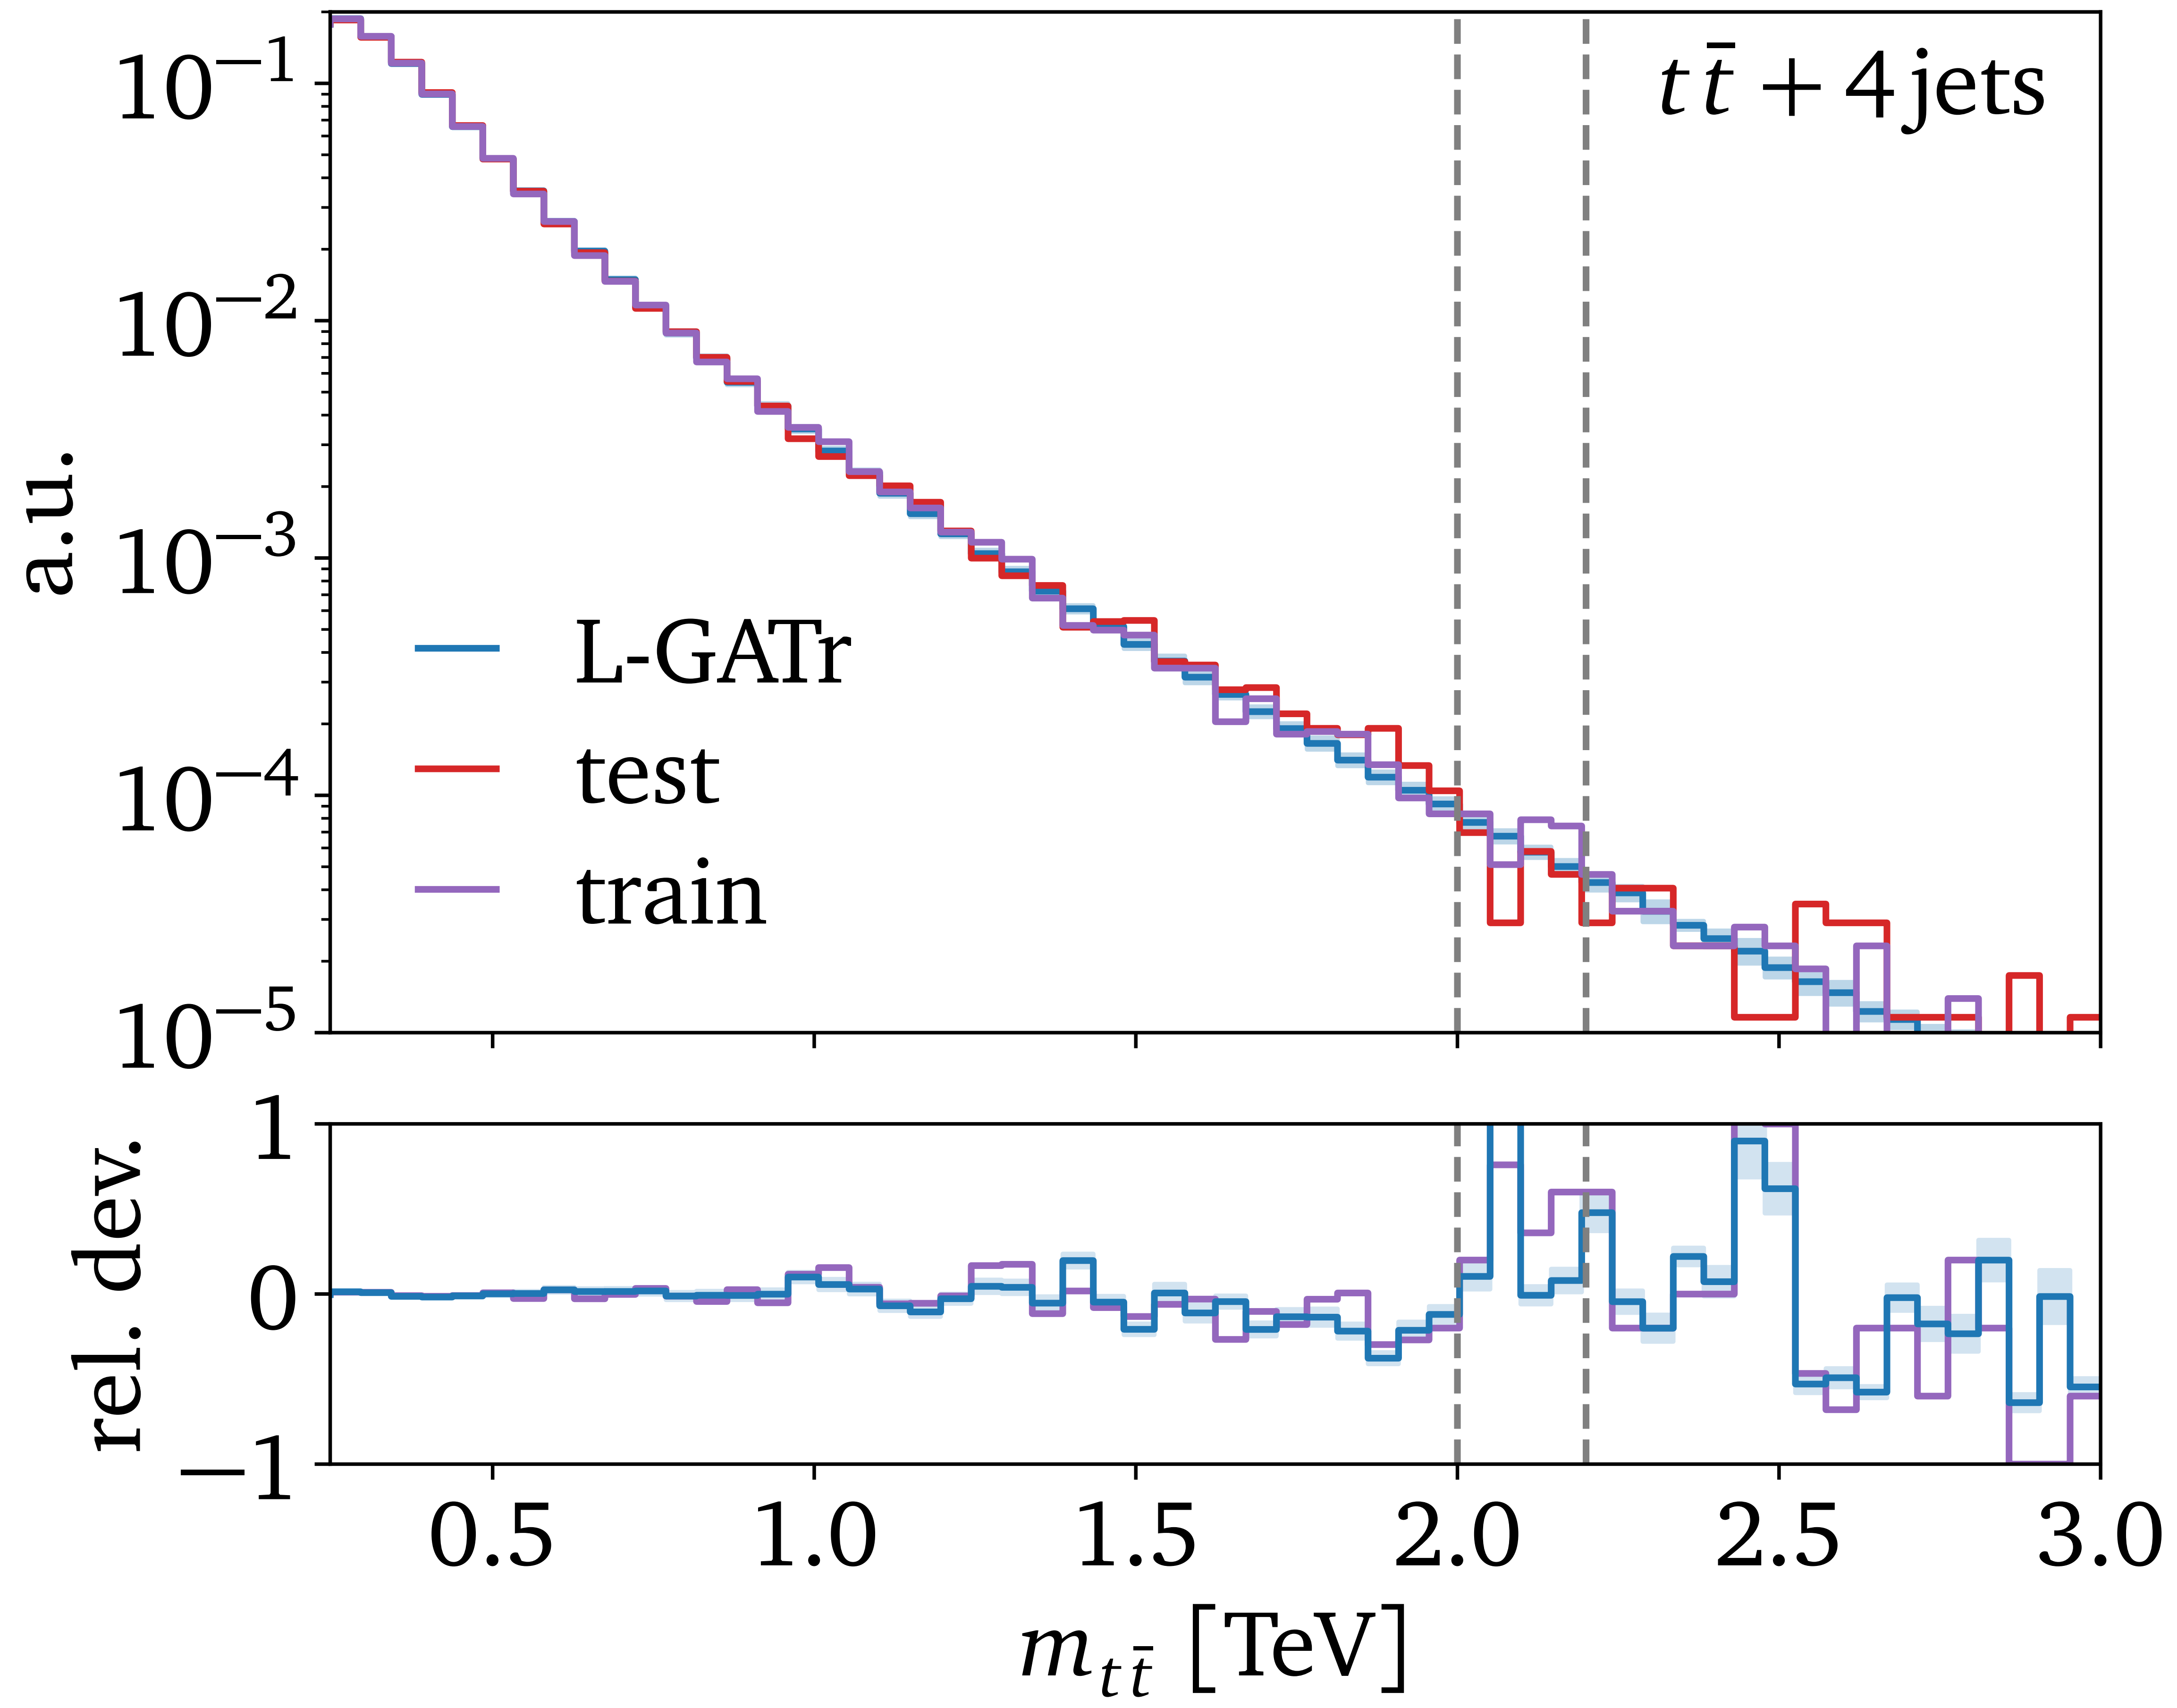

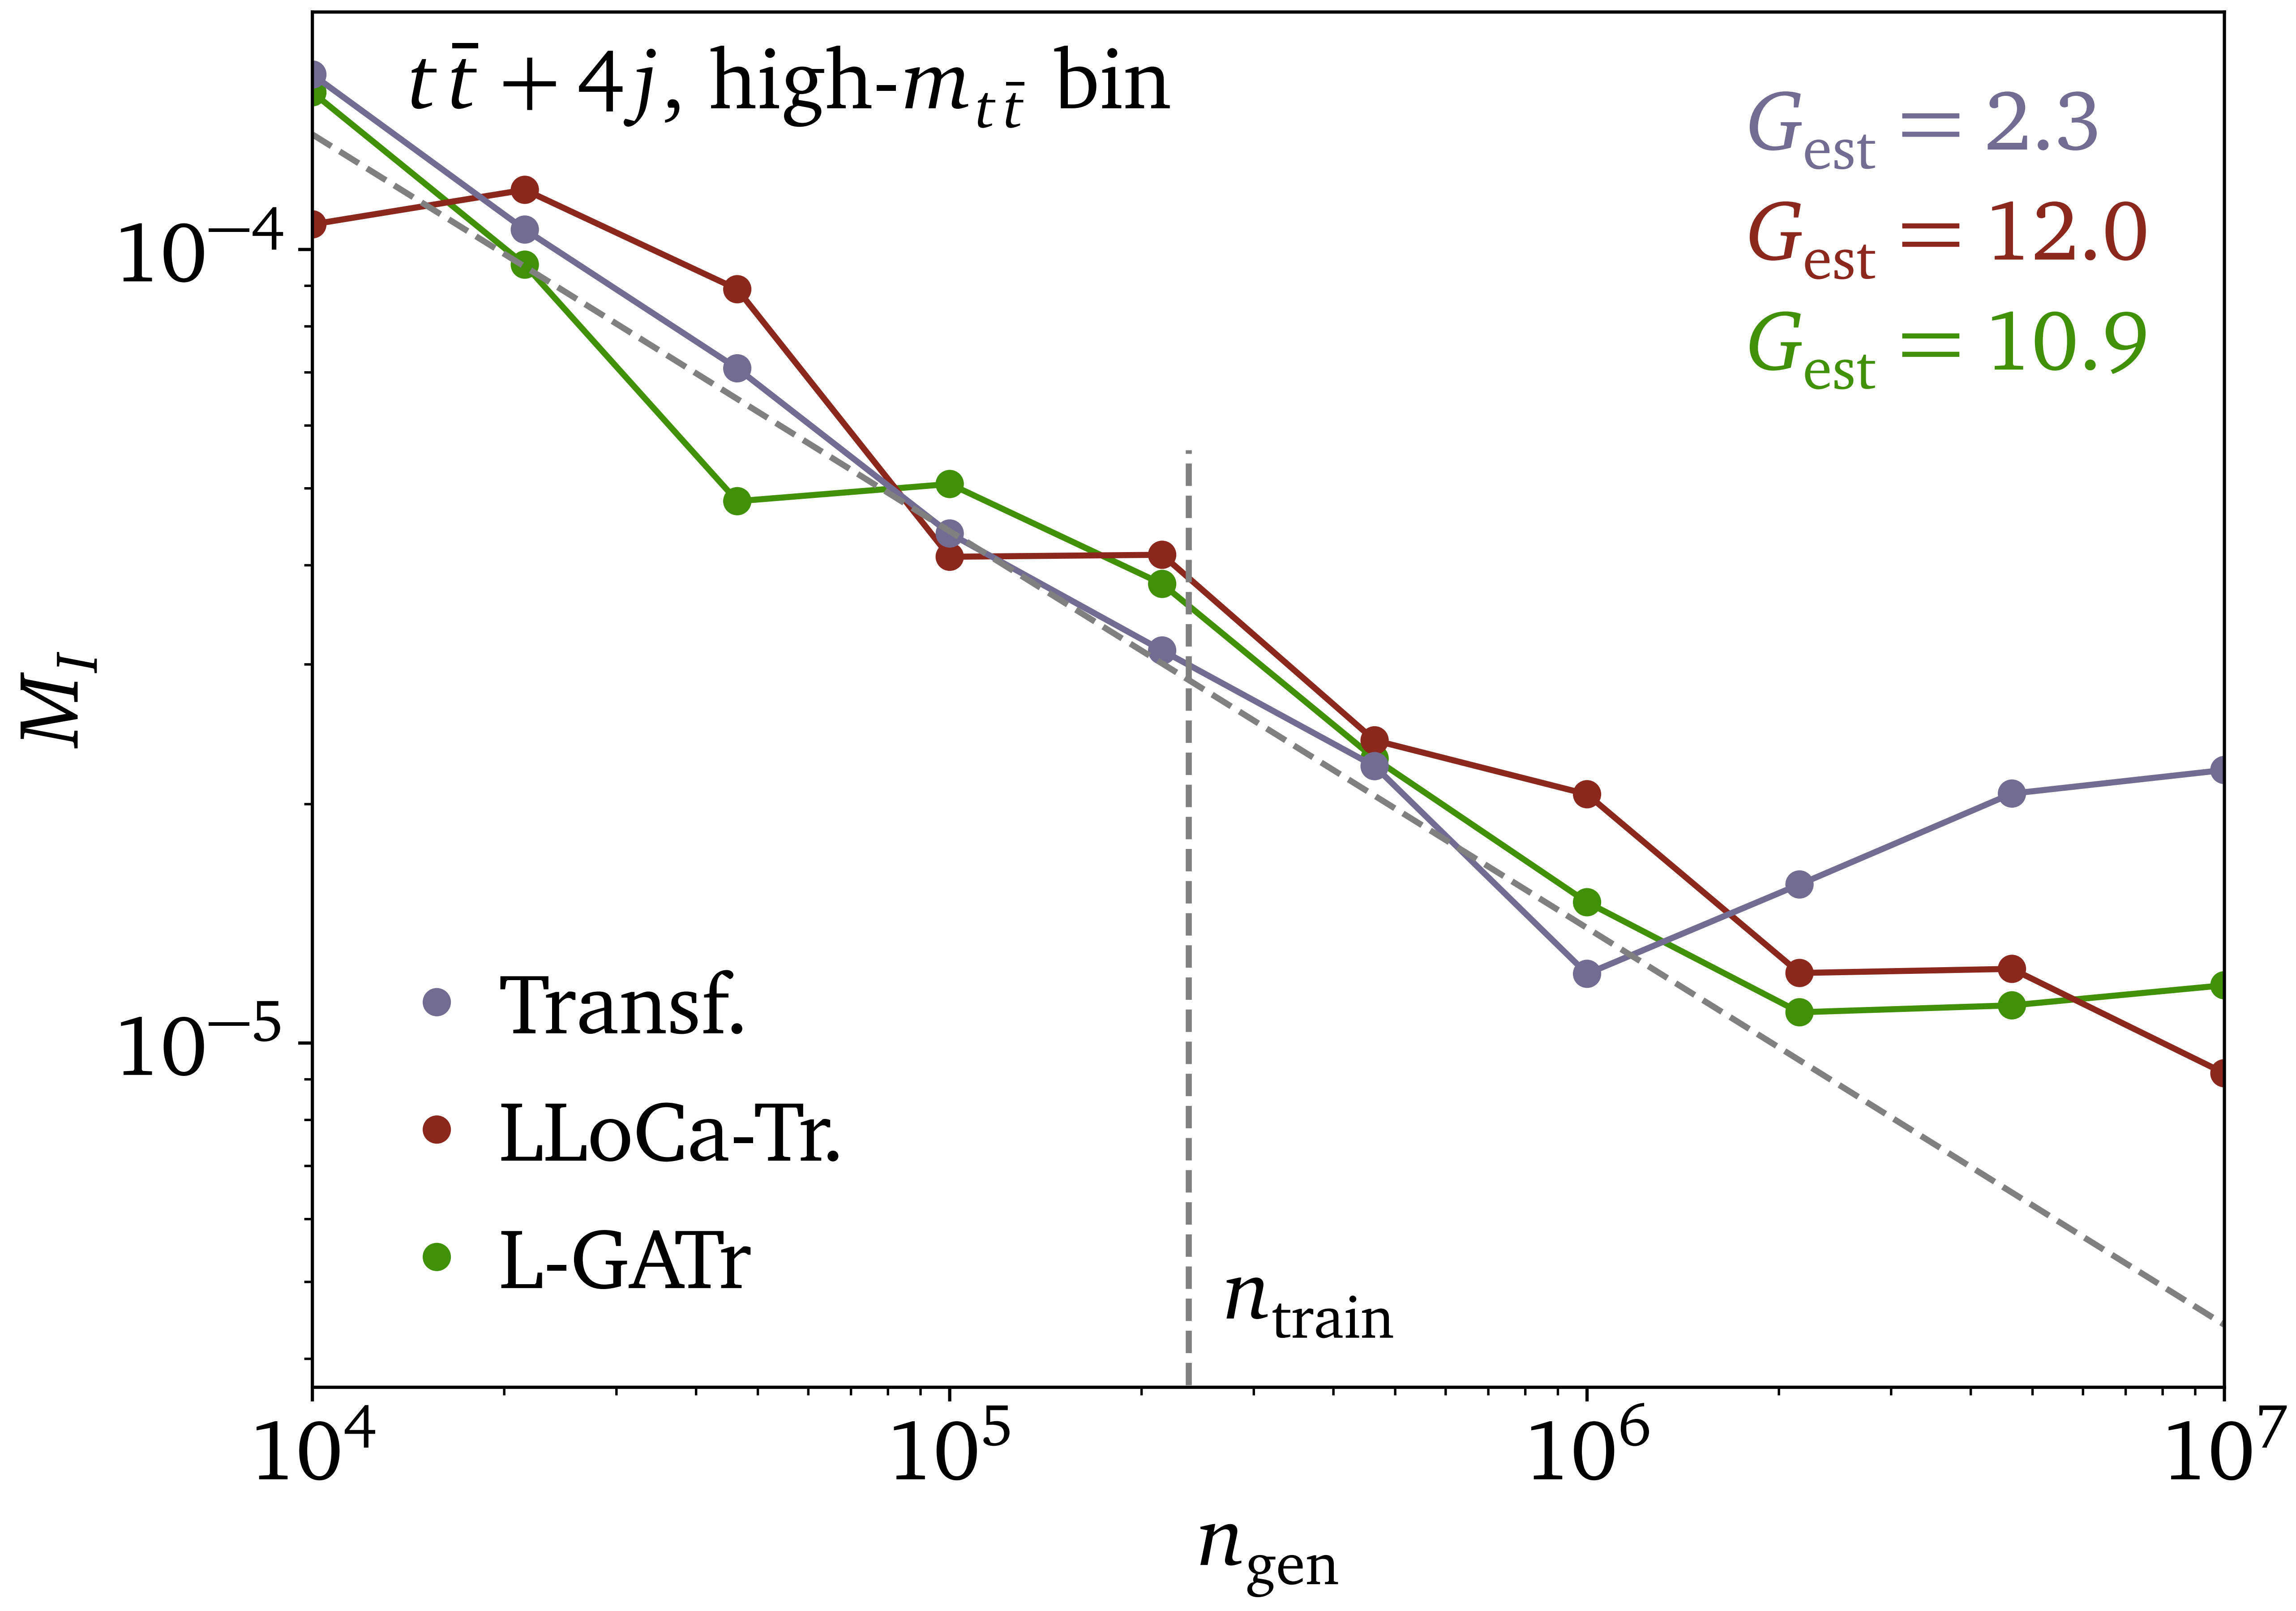

Bahl, **Diefenbacher** et. al. Forecasting Generative Amplification 2509.08048

AI Uncertainties

- **Generalization Uncertainty: Generative**
  - Cannot learn $p_{\text{gen}}(p_{\text{true}})$
  - Instead $p_{\text{gen}}(D_{\text{true}}^{n_{\text{train}}})$ with $D_{\text{true}}^{n_{\text{train}}} \sim p_{\text{true}}$
  - Extrapolation in Fourier space
  - What happens at higher resolutions than in training?

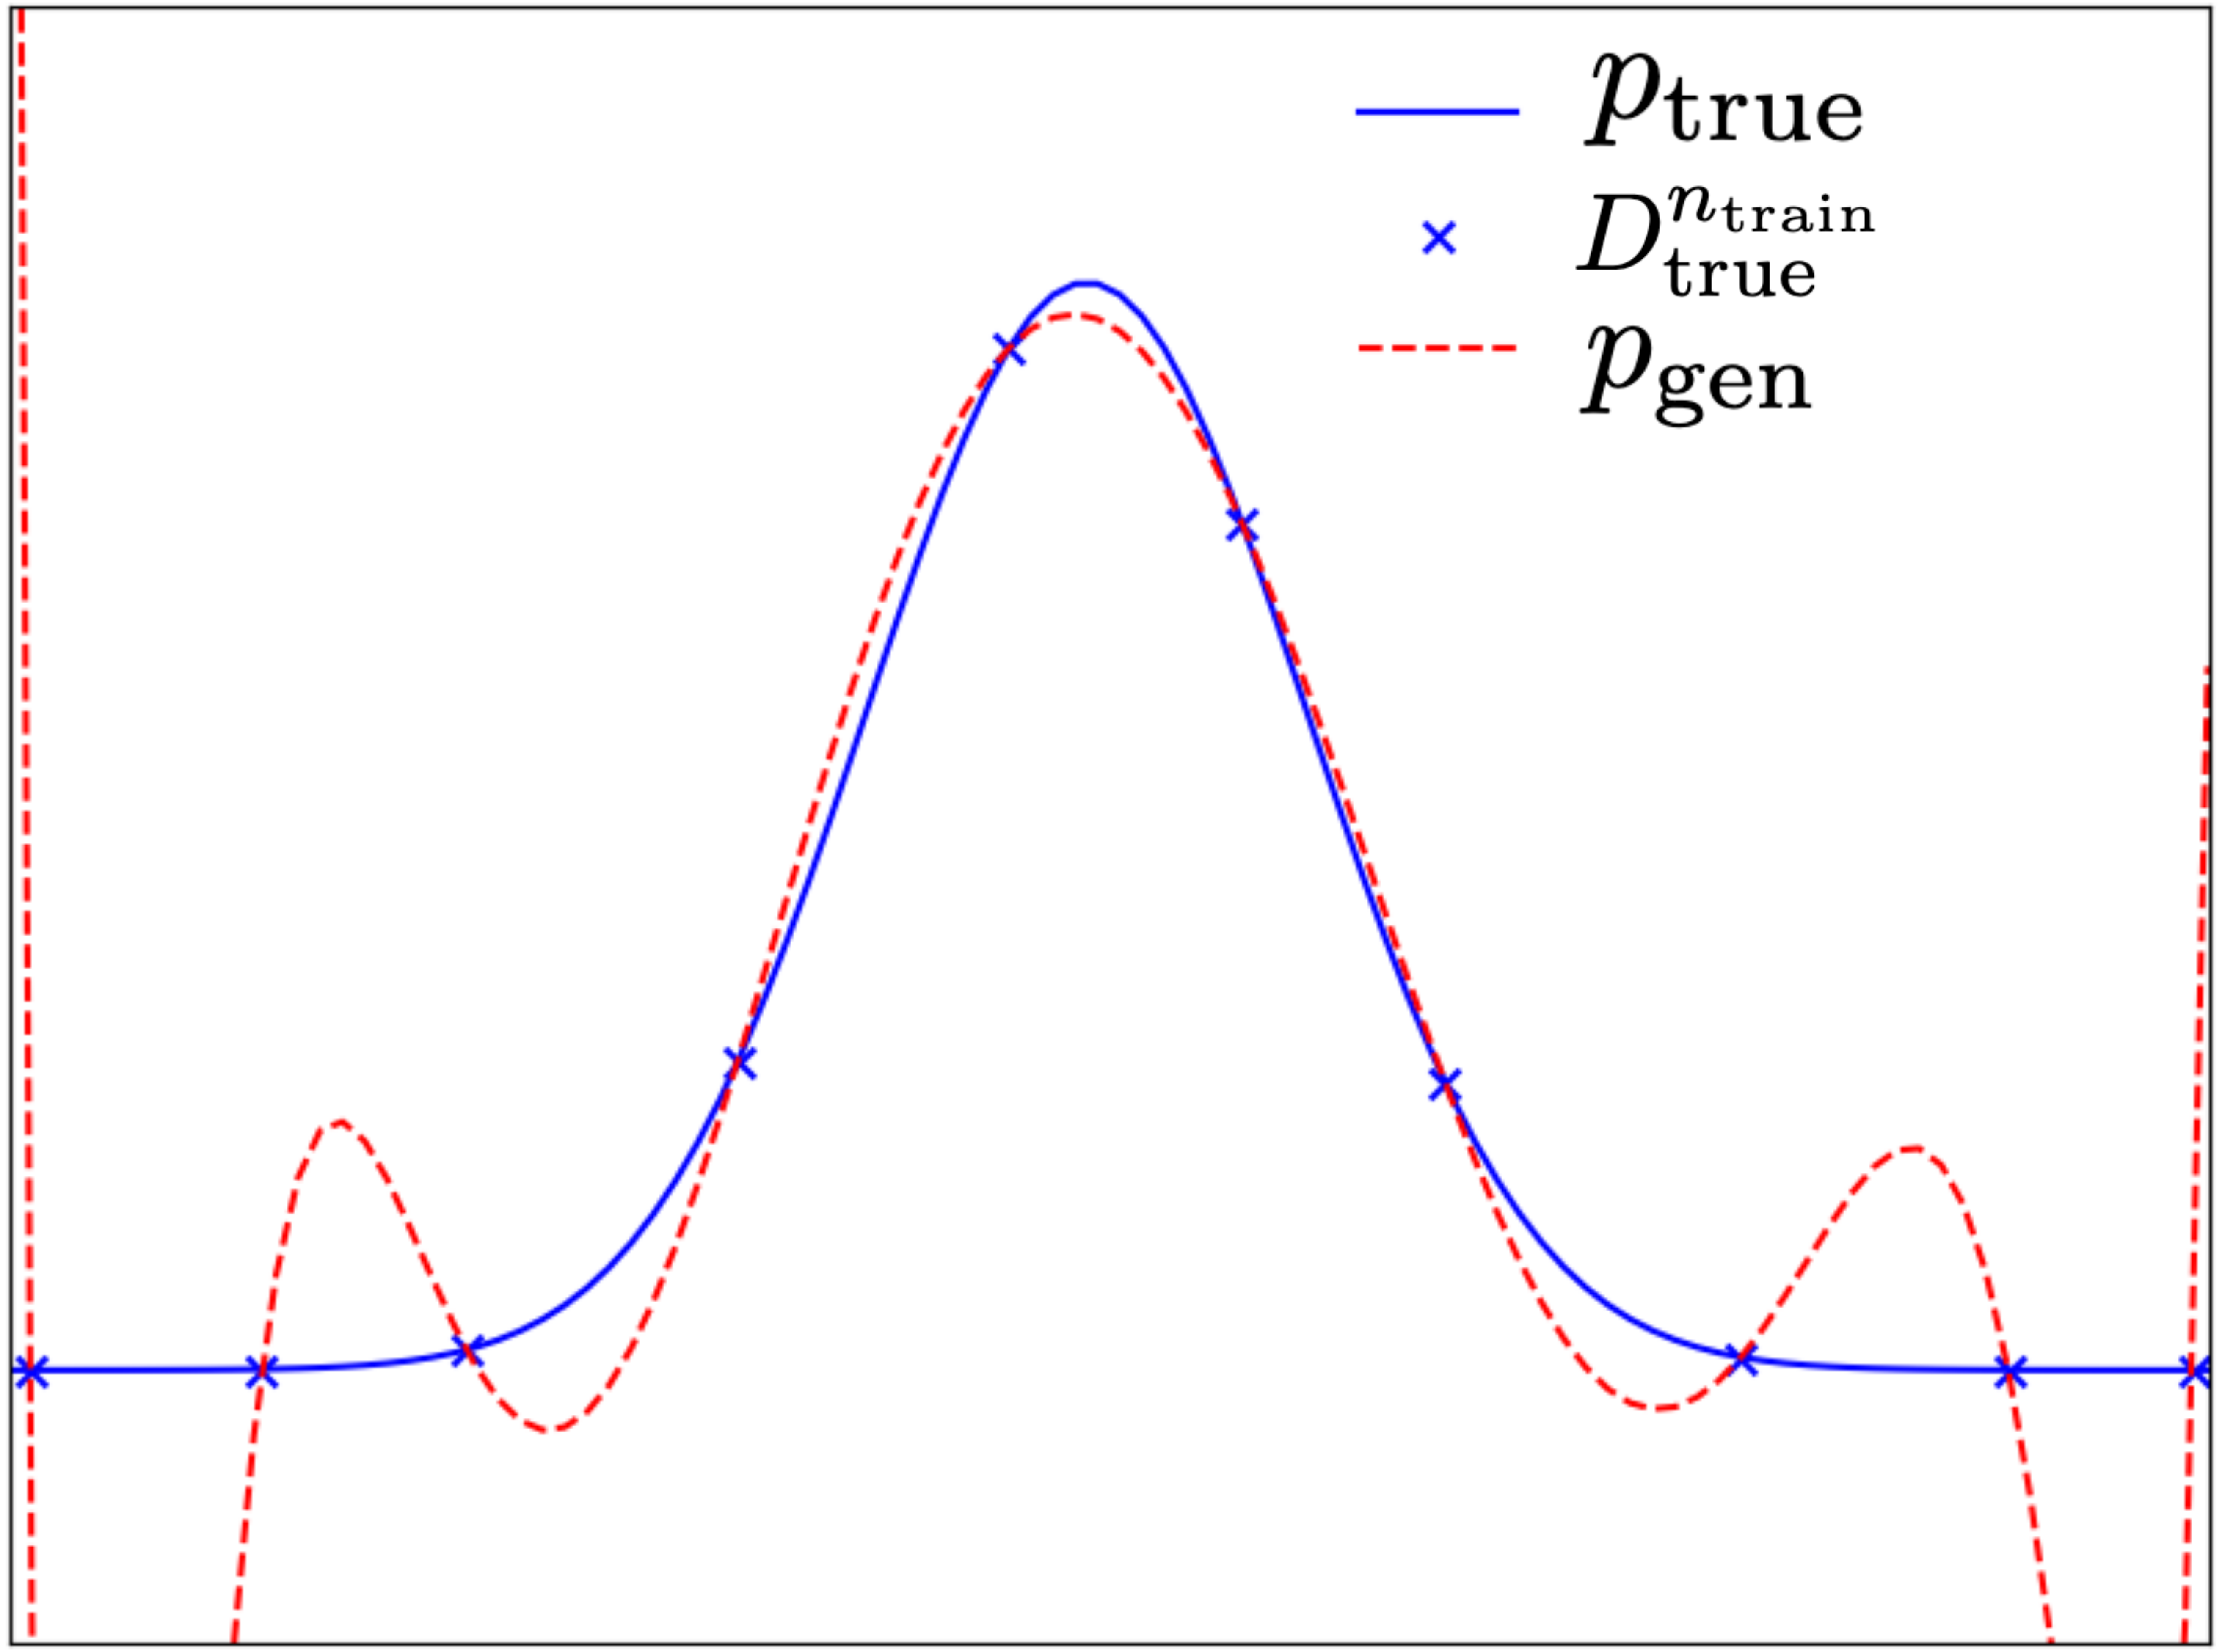

Amplification

- **Measure for generalization uncertainty** σ_{model}

$$\sigma_{\text{model}} \approx M\left[p_{\text{gen}}, p_{\text{true}}\right]$$

  - With metric M
  - p_{gen} intractable → approximate

$$D_{\text{gen}}^{n_{\text{gen}}} \sim p_{\text{gen}}$$

  - Sample to arbitrary precision

$$\sigma_{\text{model}} \approx M\left[D_{\text{gen}}^{n_{\text{gen}}}, p_{\text{true}}\right]\Big|_{n_{\text{gen}} \to \infty}$$

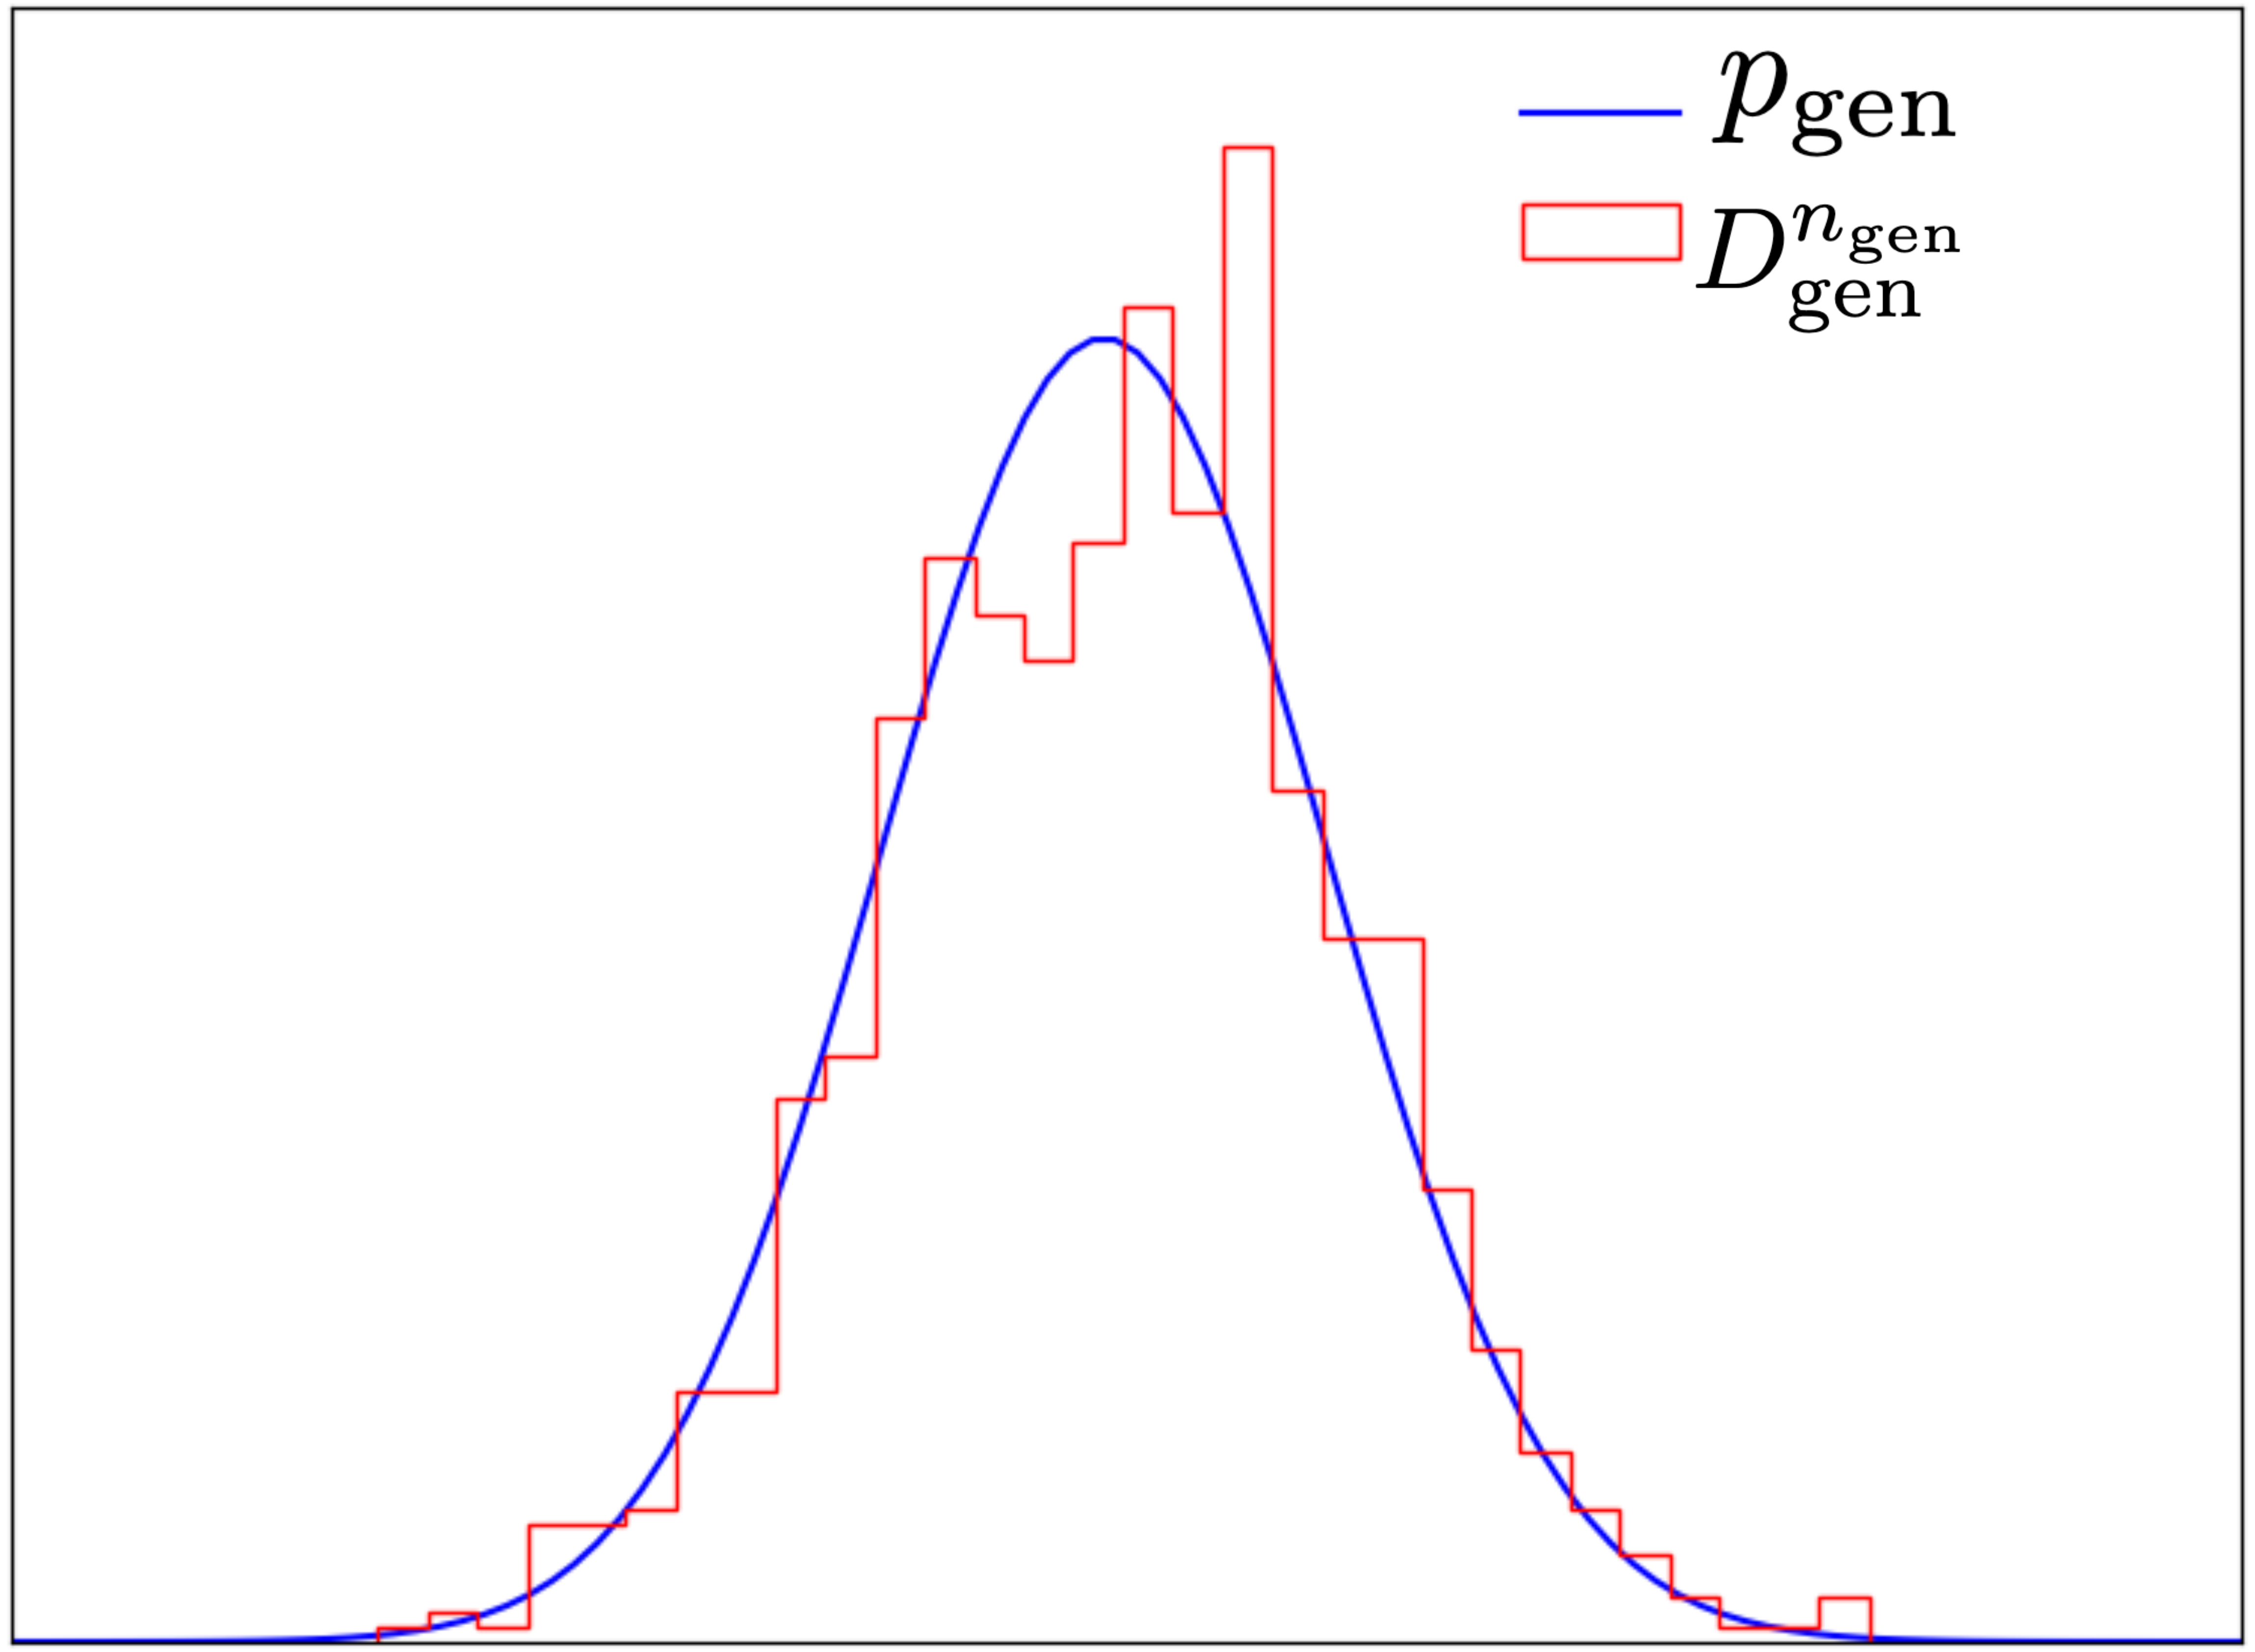

Generative Simulation

- **Generative networks**
  - Trained on samples $x \sim p_{\text{true}}(x)$

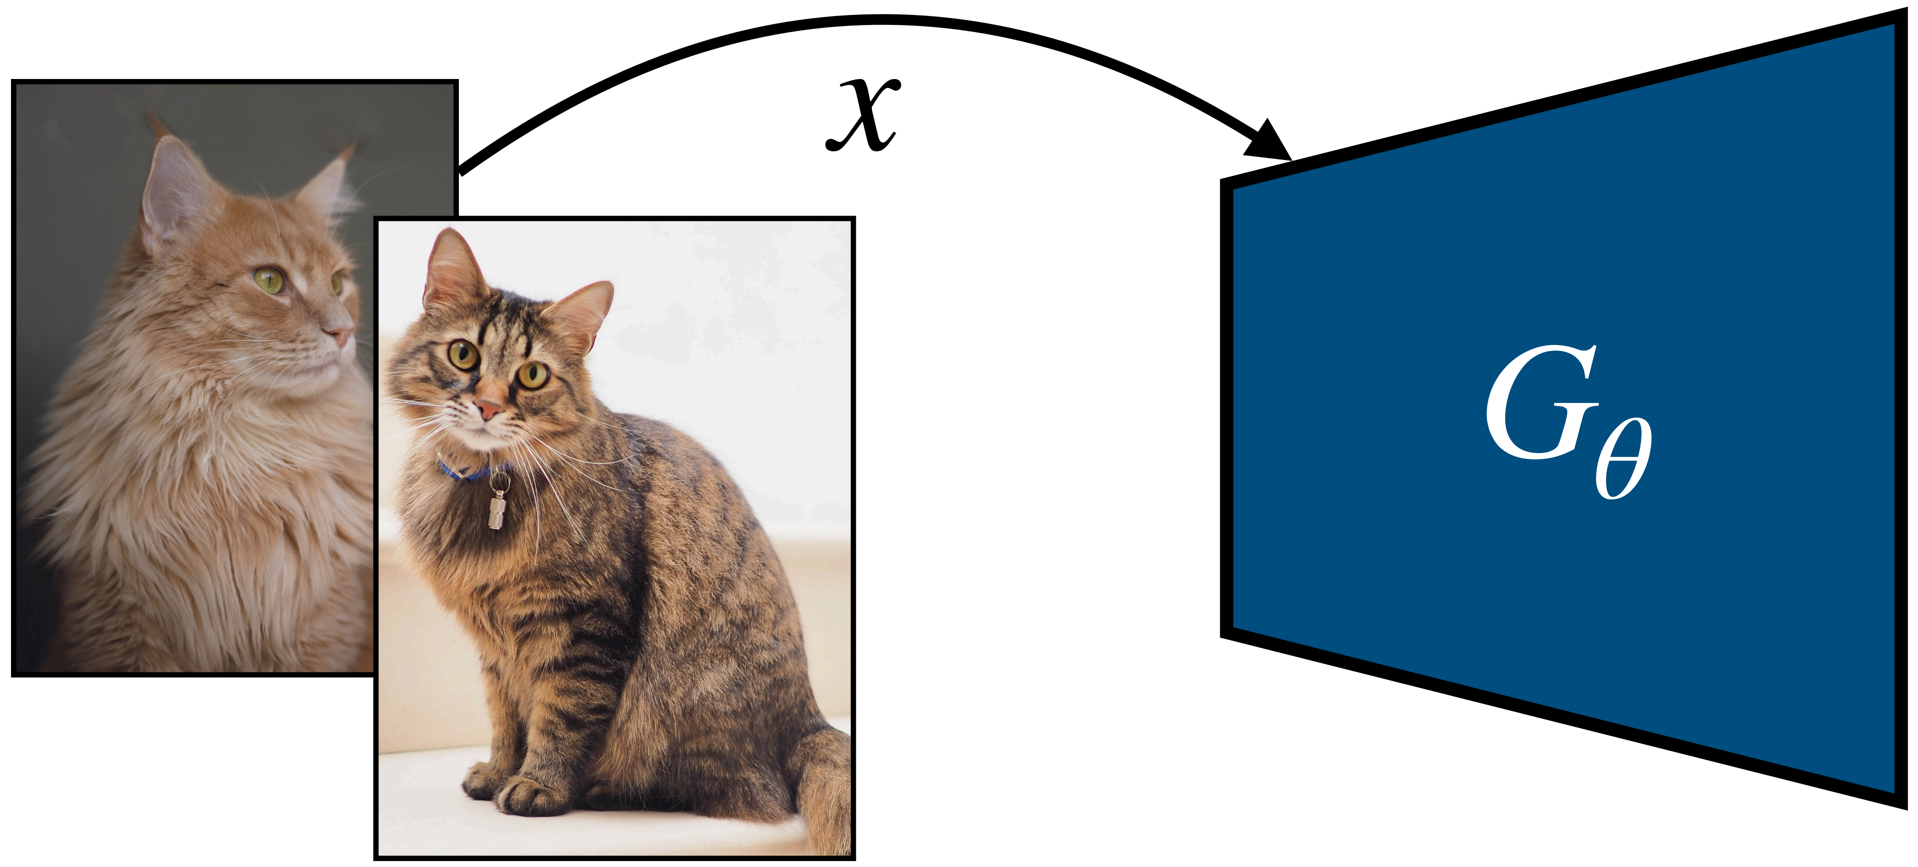

Generative Simulation

- **Generative networks**

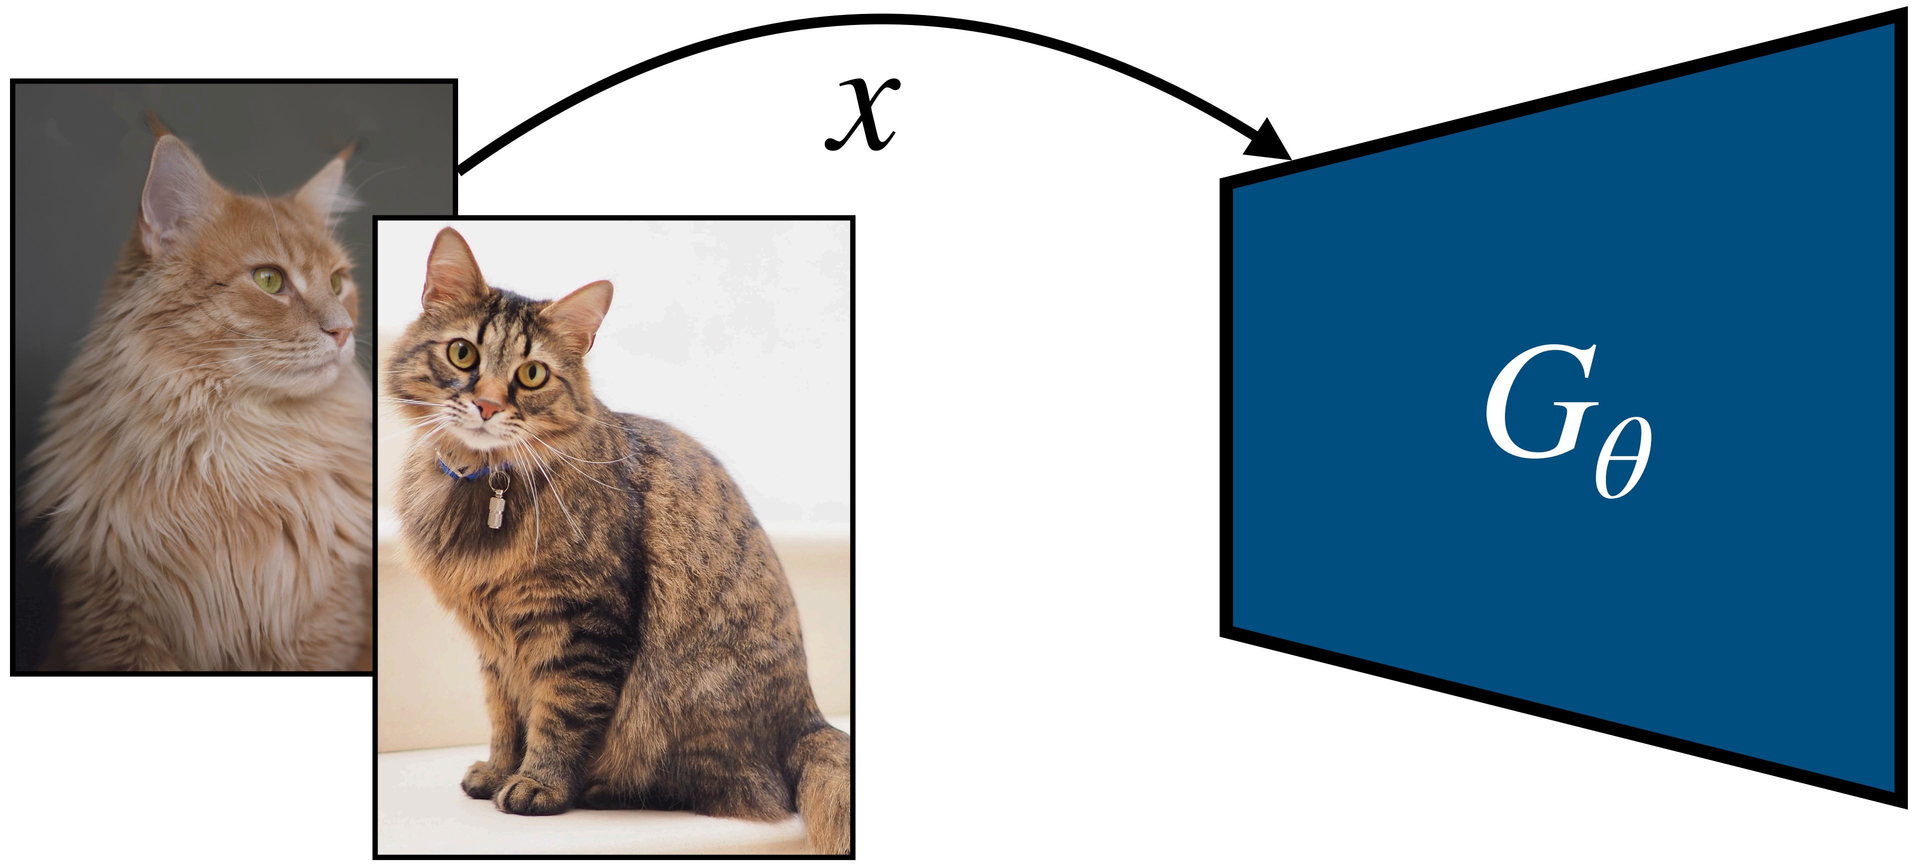

  - Trained on samples $x \sim p_{\text{true}}(x)$
  - Learns to map noise to $p_{\text{gen}}(x)$

$$z \sim \mathcal{N}(0,1) \xrightarrow{G_\theta(z|c)} x \sim p_{\text{gen}}(x)$$

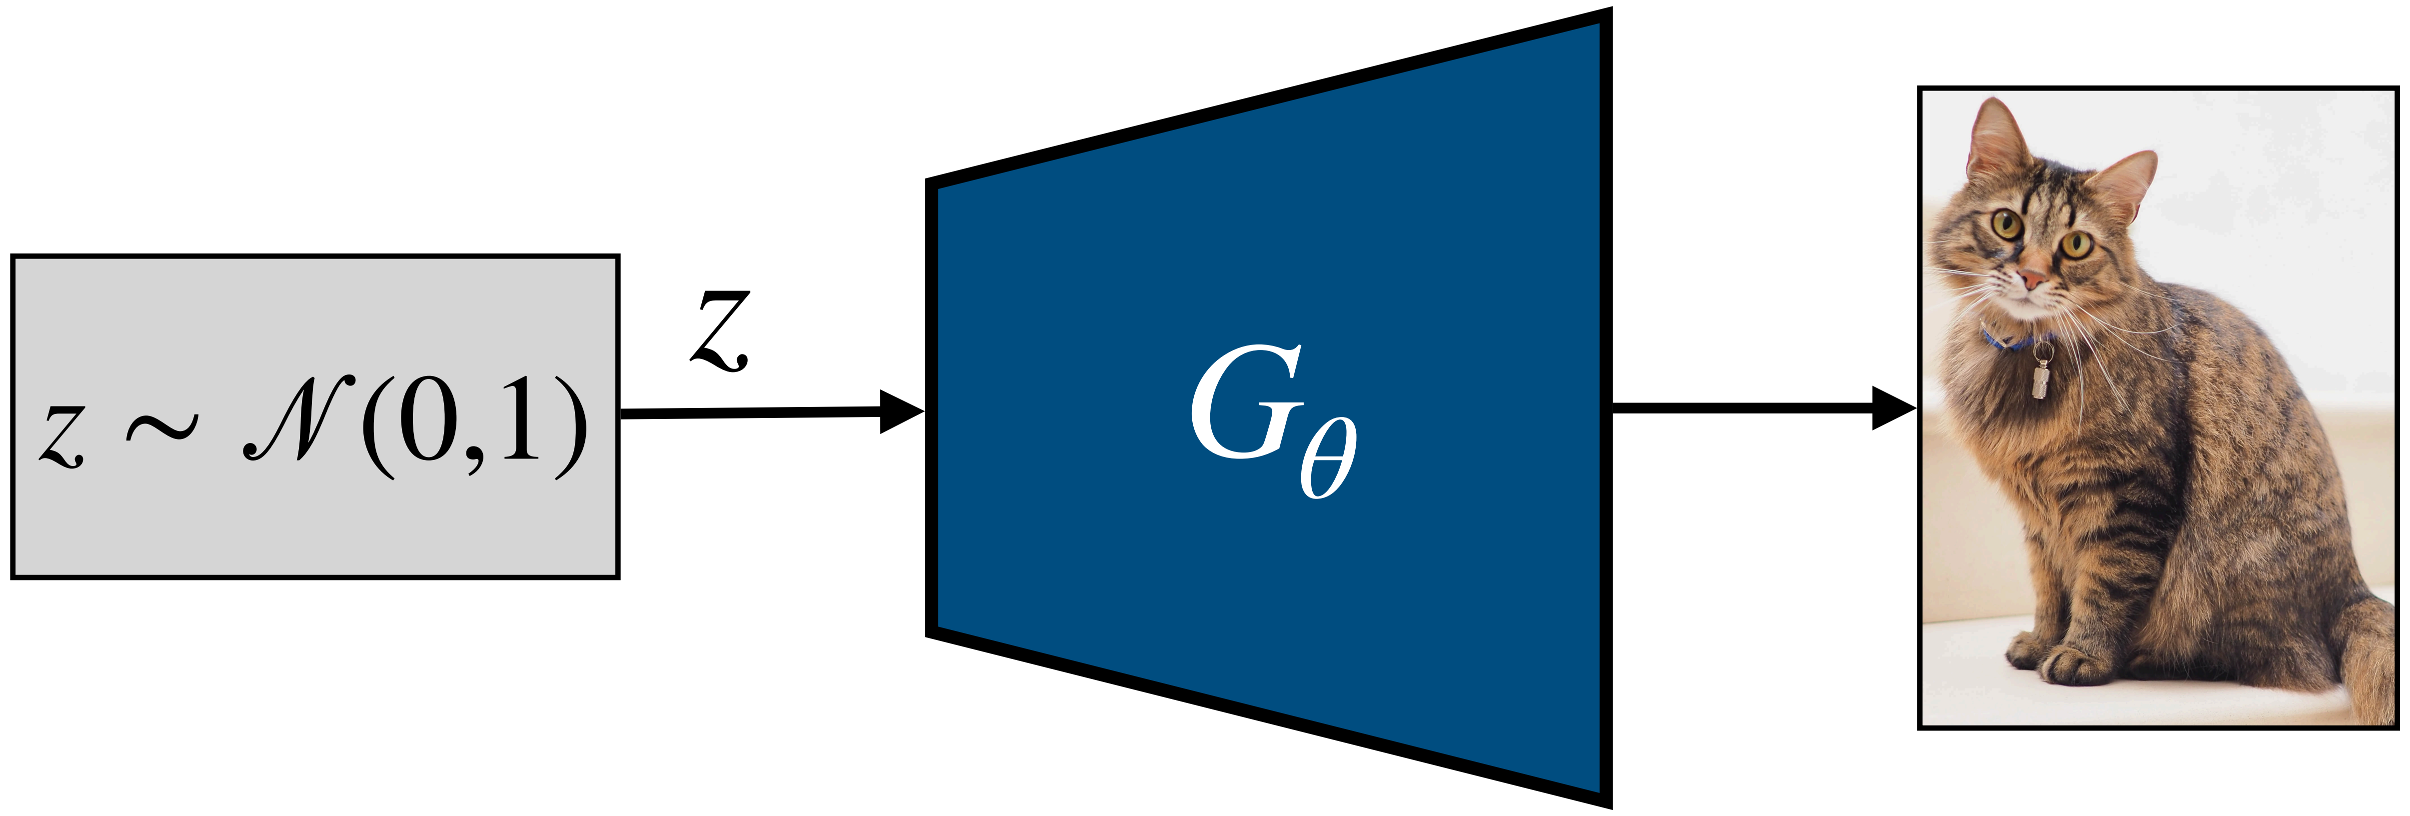

  - Goal: $p_{\text{gen}}(x) \approx p_{\text{true}}(x)$

AI Uncertainties

- **Statistical Uncertainty**
- **Systematic Uncertainty**
- **Generalization Uncertainty**
  - Inference data ≠ training data
  - Extrapolation outside training

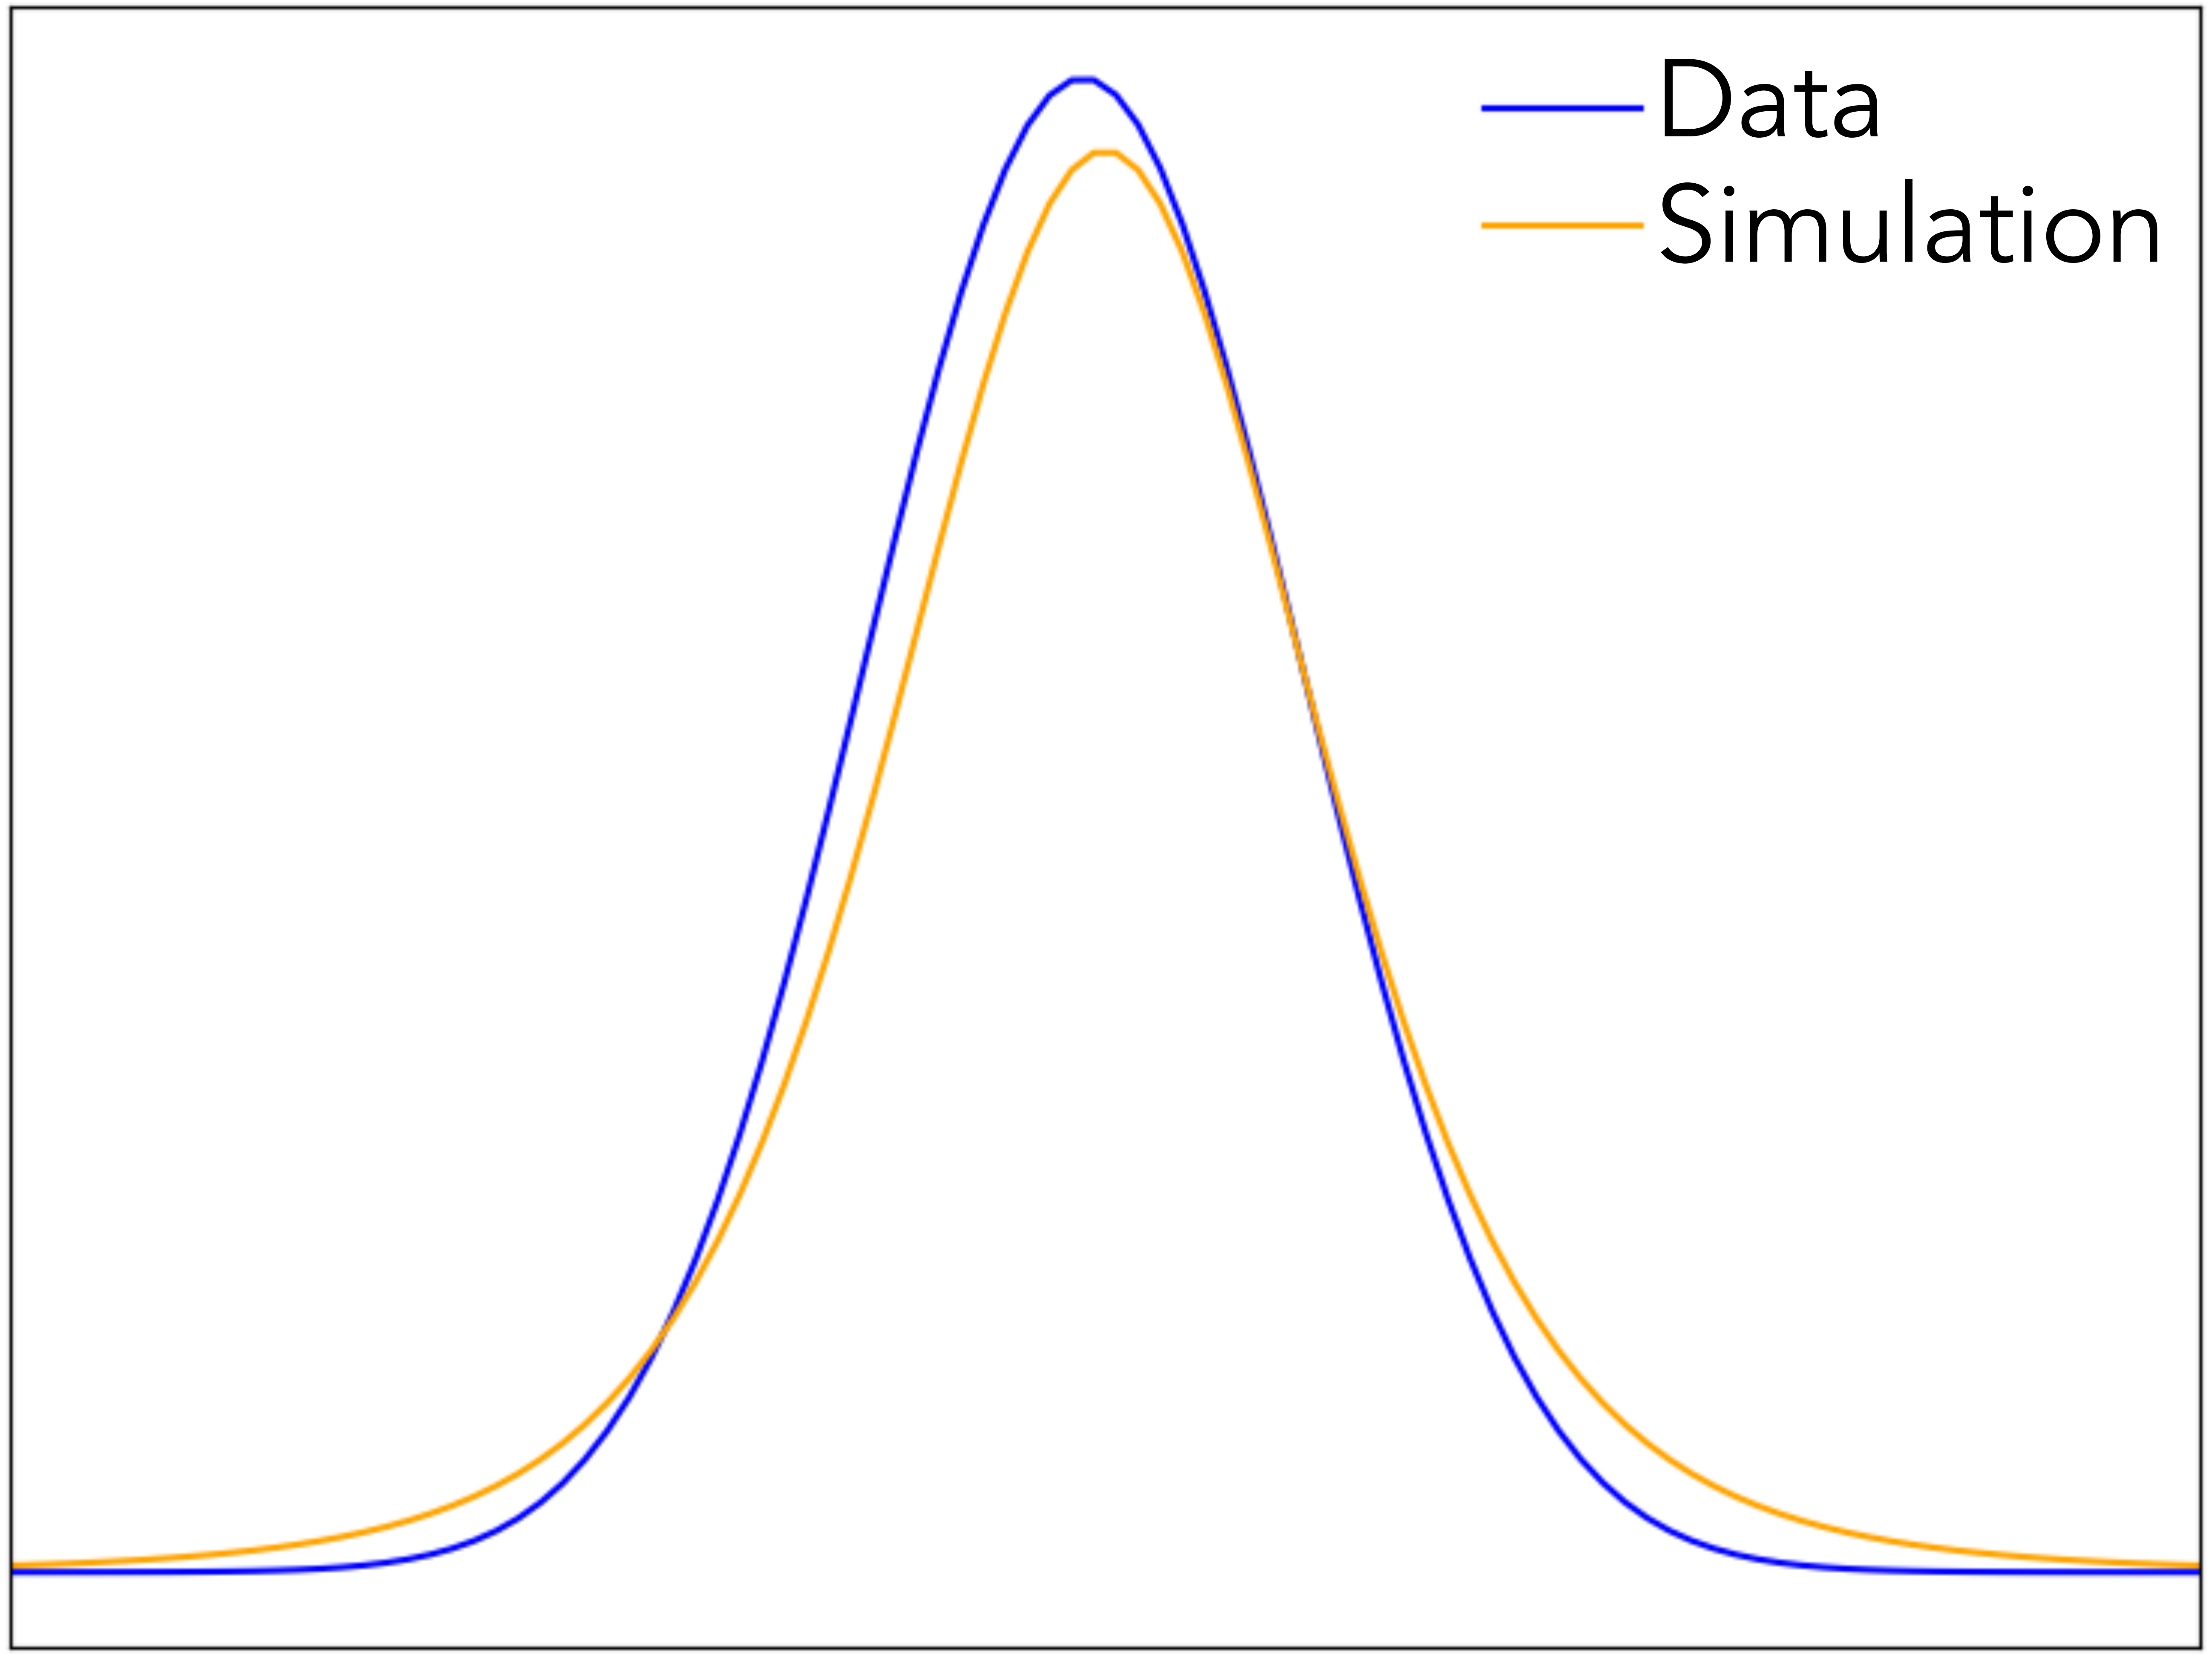

AI Uncertainties

- **Statistical Uncertainty σ_{stat}**
  - Or epistemic, reducible
  - Scales with training data size
- **Systematic Uncertainty σ_{syst}**
  - Or aleatoric, irreducible
  - Independent of training data size

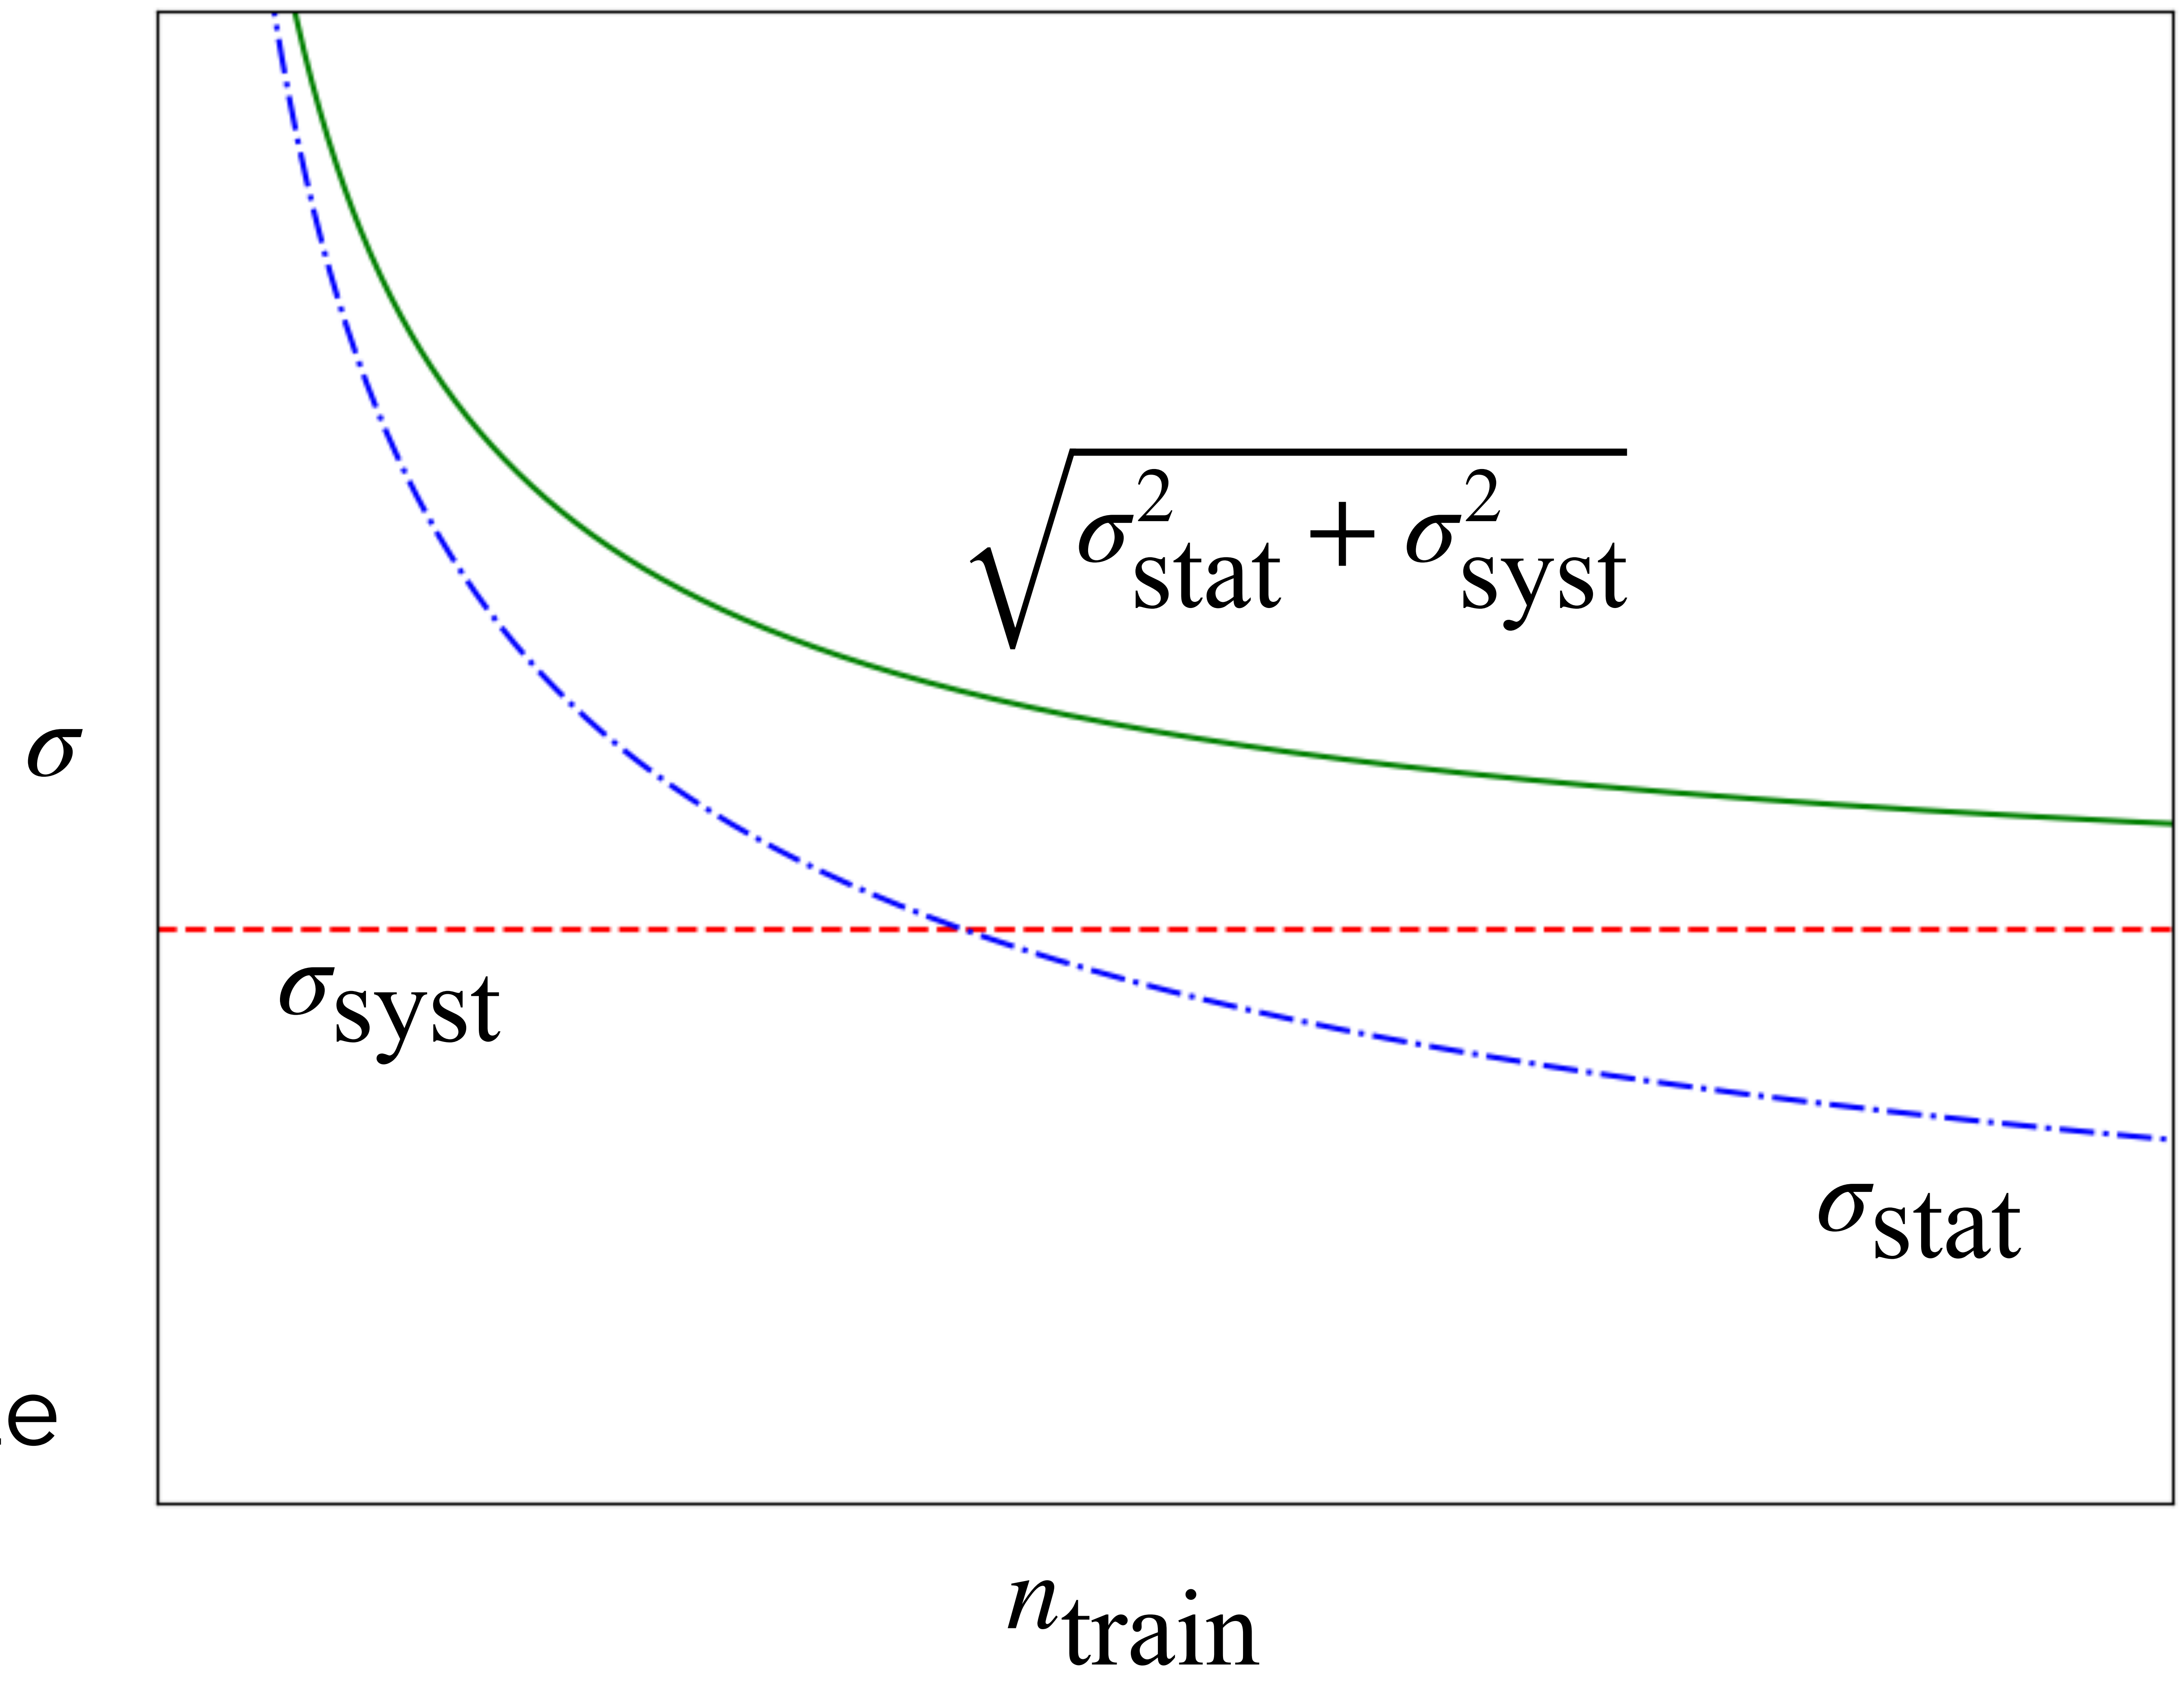

AI Uncertainties

Physics needs uncertainties

Classifier model:

- **Statistical Uncertainty** σ_{stat}
  - Or epistemic, reducible
  - Scales with training data size
- **Systematic Uncertainty** σ_{syst}
  - Or aleatoric, irreducible
  - Independent of training data size

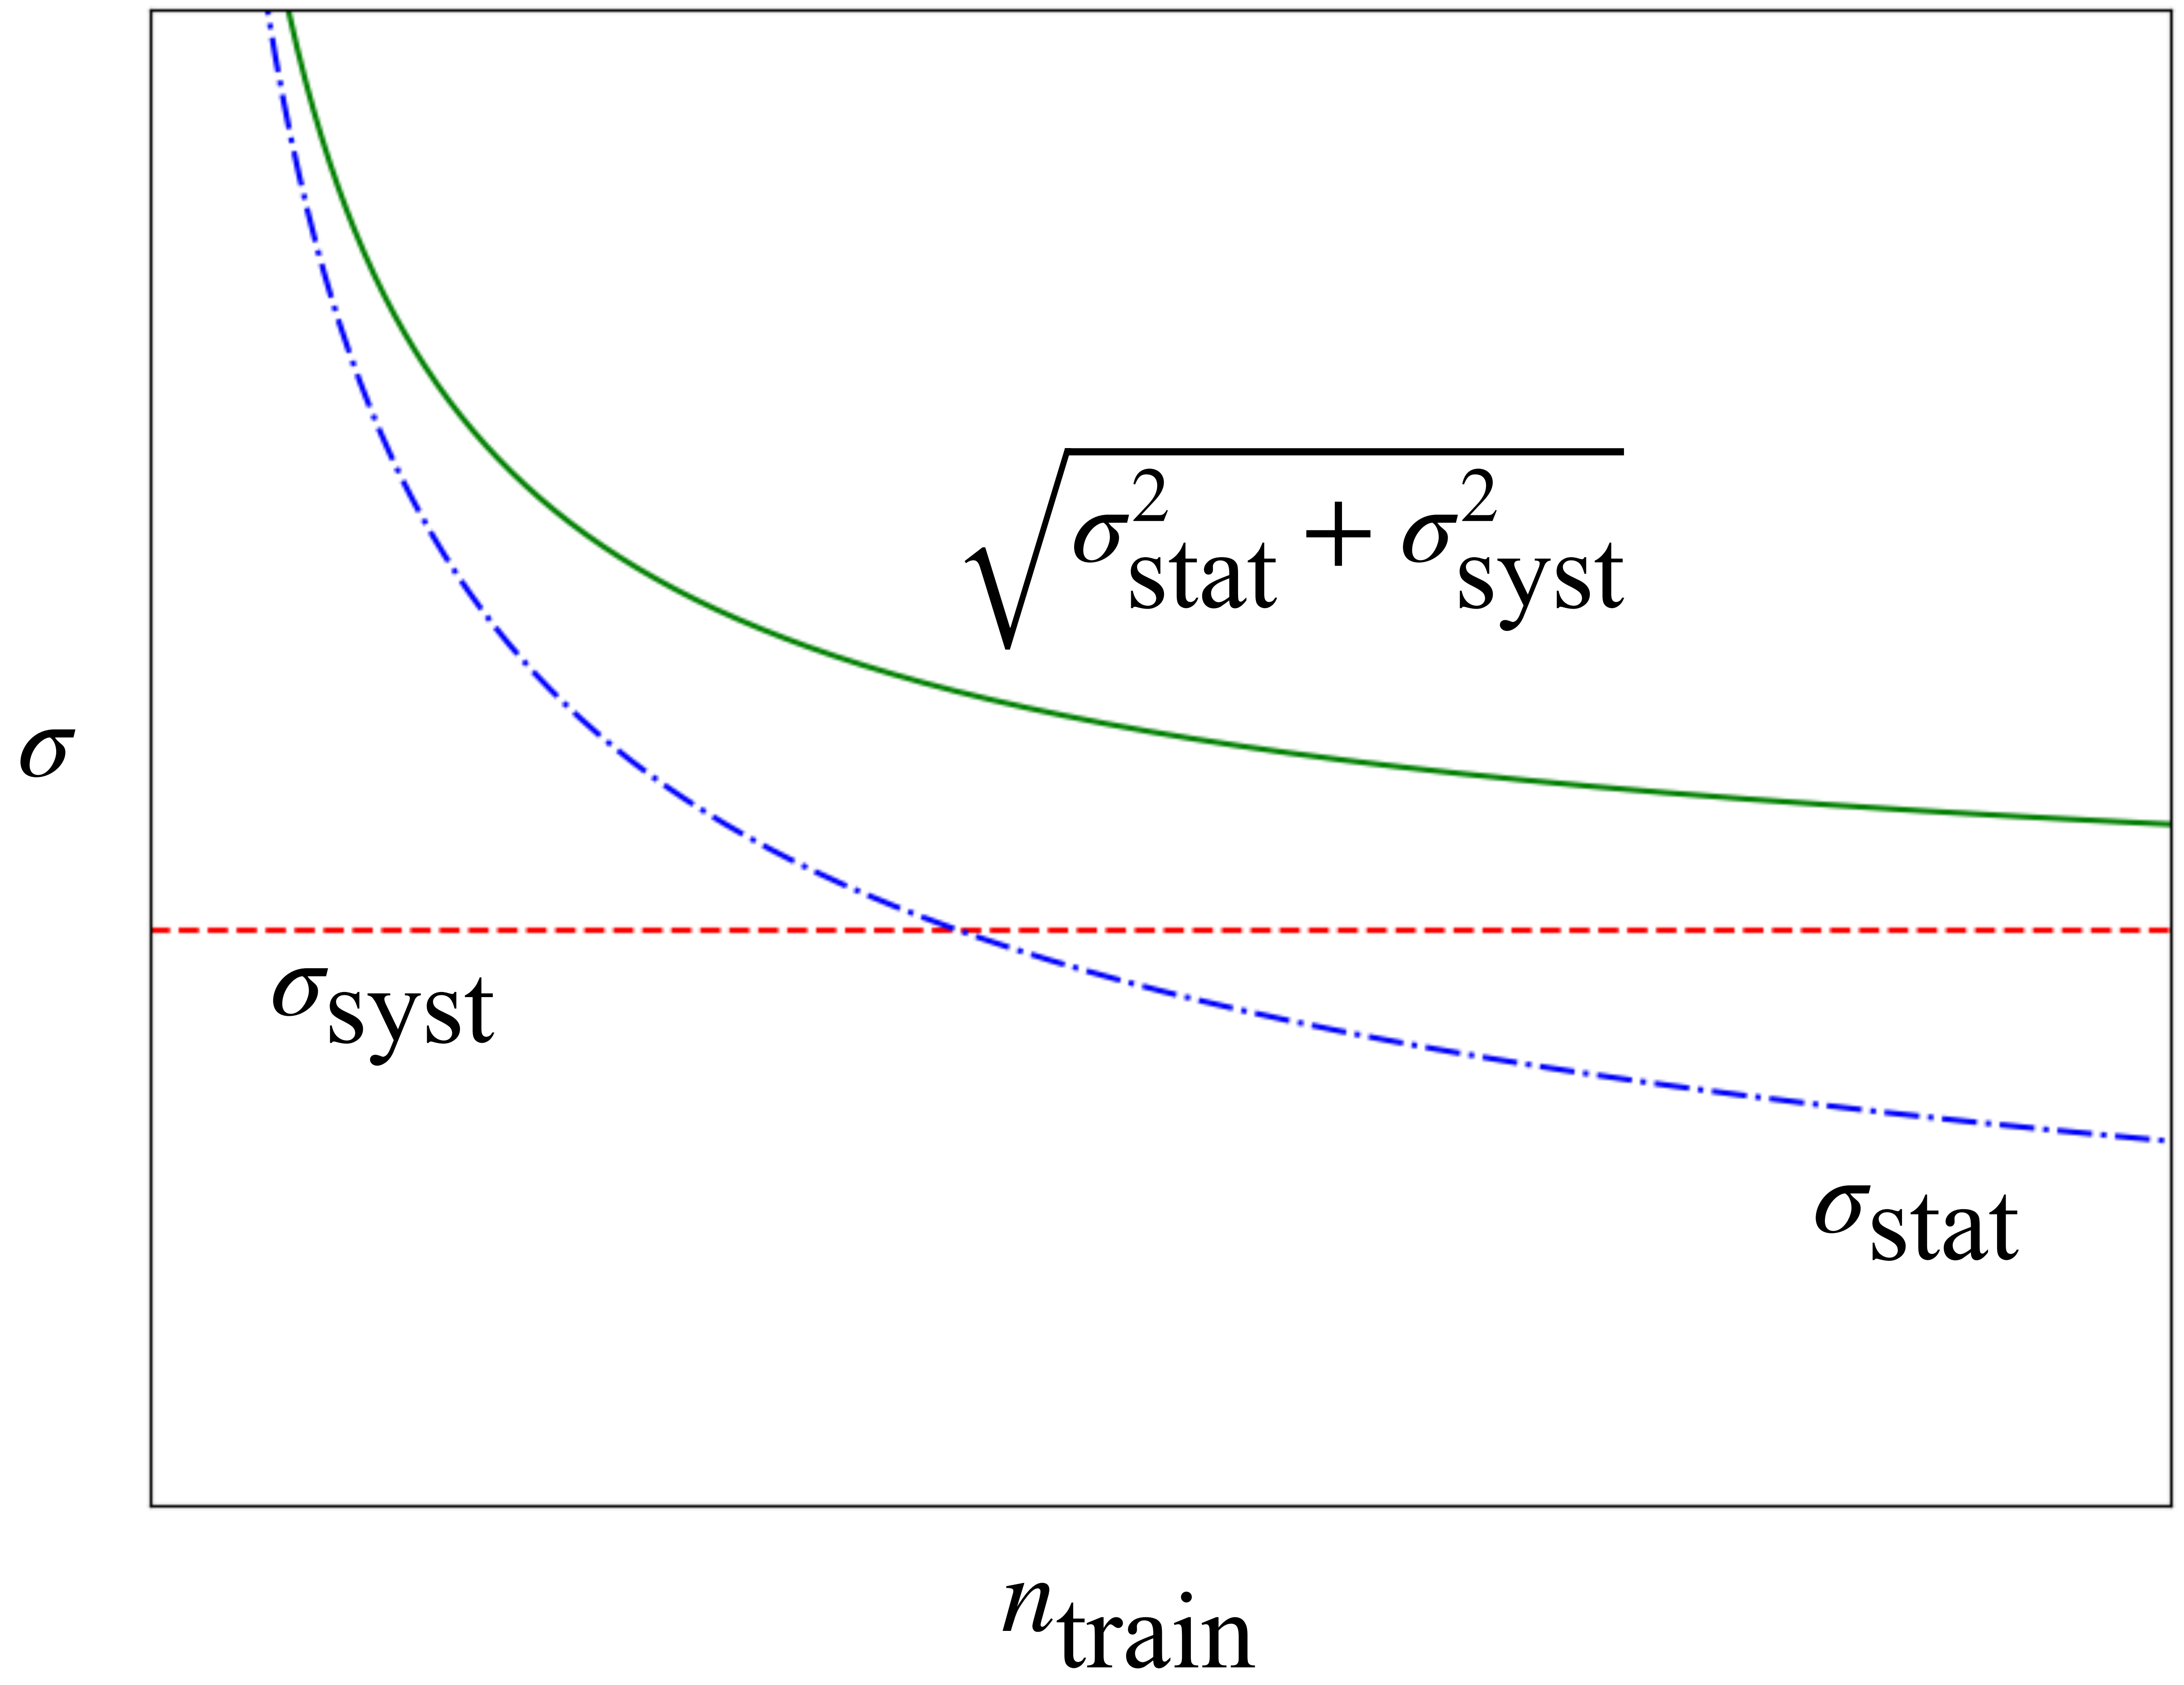

Quantifying Amplification

$$M\left[D_{\mathrm{gen}}^{n_{\mathrm{gen}}}, p_{\mathrm{true}}\right]$$

$$= \sqrt{\sigma_{\mathrm{stat}}(n_{\mathrm{gen}})^2 + \sigma_{\mathrm{model}}(p_{\mathrm{gen}}, p_{\mathrm{true}}}$$

$$\propto \frac{1}{\sqrt{n_{\mathrm{gen}}}} \qquad \propto \frac{1}{\sqrt{n_{\mathrm{train}}}}$$

$$M\left[D_{\mathrm{true}}^{n_{\mathrm{equiv}}}, p_{\mathrm{true}}\right]$$

$$= \sigma_{\mathrm{stat}}(n_{\mathrm{equiv}})$$

$$\propto \frac{1}{\sqrt{n_{\mathrm{eqiv}}}}$$

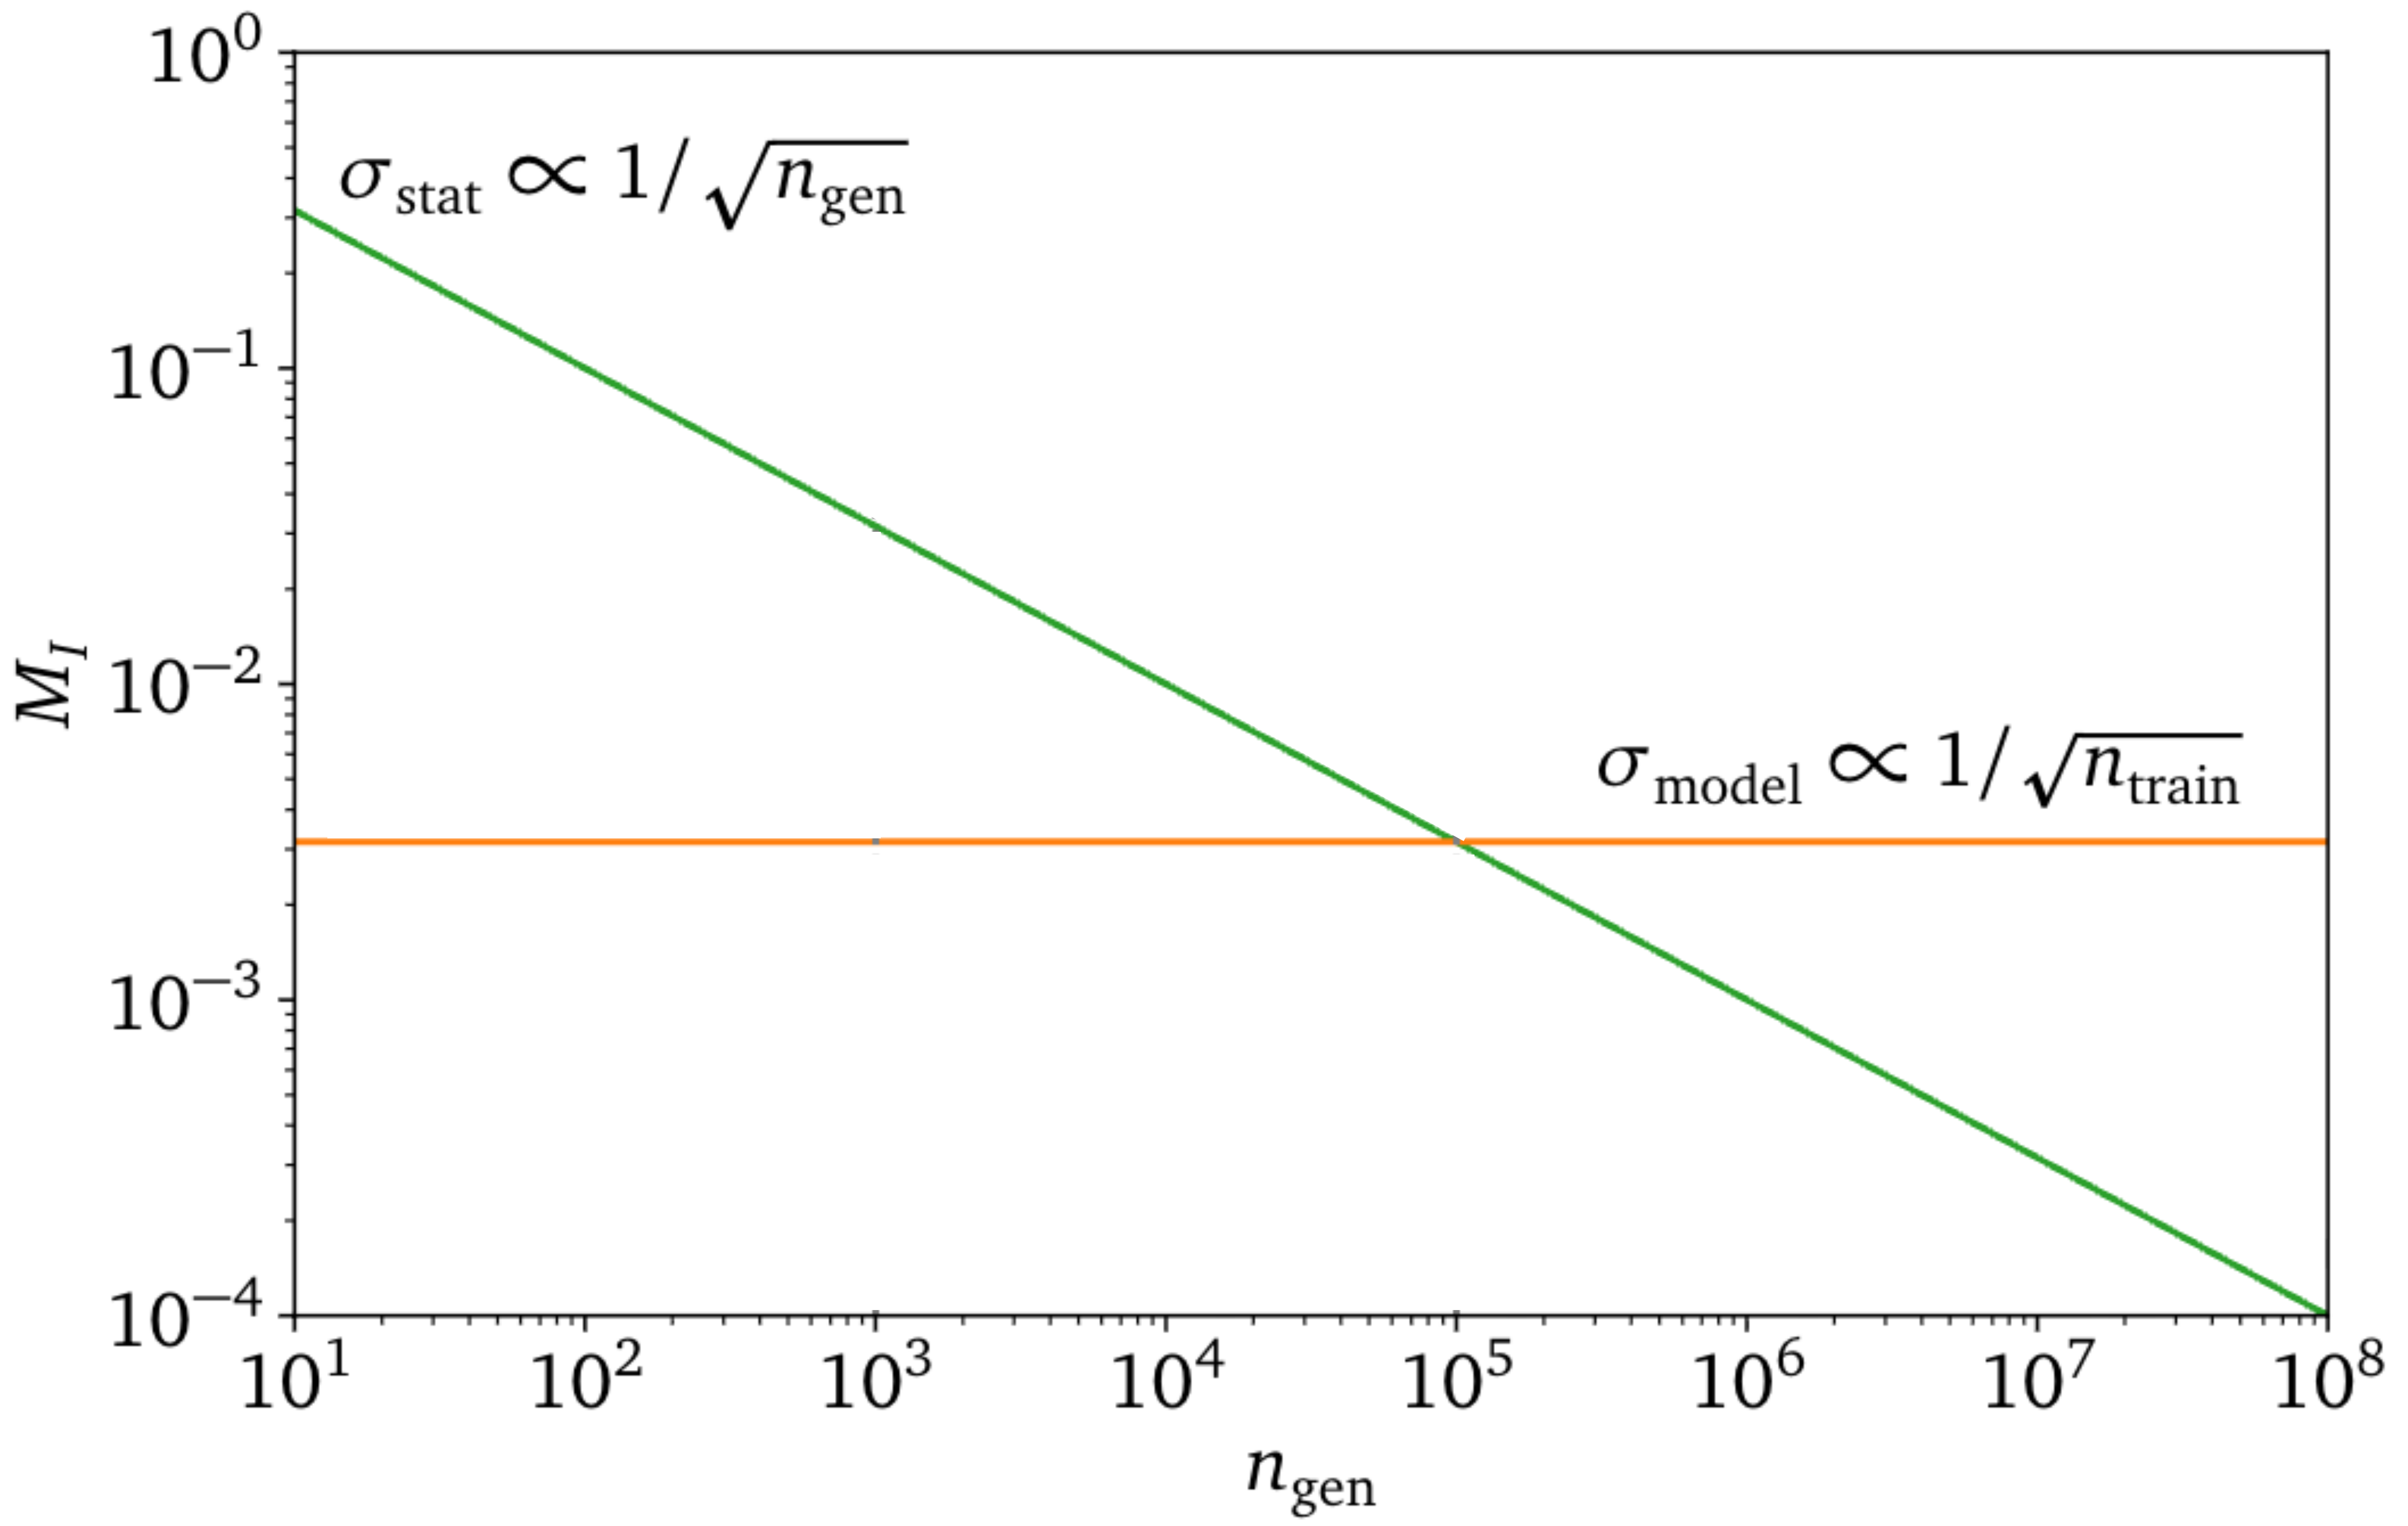

Bahl, **Diefenbacher** et al. Forecasting Generative Amplification 2509.08048

Quantifying Amplification

- How to get $M\left[D^n_{\text{true}}, p_{\text{true}}\right]$ and $M\left[D^n_{\text{gen}}, p_{\text{true}}\right]$?
  - Option 1: Extrapolate
  - Option 2: Calculate Approximately