



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

ML4Jets, 07.11.2023

Diffusion Models

for LHC event generation

SPONSORED BY THE



Federal Ministry
of Education
and Research

Anja Butter, Nathan Huetsch, Sofia Palacios Schweitzer, Tilman Plehn, Peter Sorrenson, Jonas Spinner
arXiv: 2305.10475



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

ML4Jets, 07.11.2023

Uncertainty-Aware Diffusion Models

for LHC event generation

SPONSORED BY THE



Federal Ministry
of Education
and Research


Anja Butter, Nathan Huetsch, Sofia Palacios Schweitzer, Tilman Plehn, Peter Sorrenson, Jonas Spinner
arXiv: 2305.10475

Why Event Generation?

Vast amount of data collected
by collider experiments

Standard Model is probed


Theoretical predictions
(simulation) need to match
experimental statistics



Why ML Event Generation?

After high luminosity runs $\rightarrow \sim 20$ times as much data

Theoretical predictions need to be even more precise (include higher correction terms)




Why ML Event Generation?

After high luminosity runs $\rightarrow \sim 20$ times as much data


Theoretical predictions needs to be even more precise (include higher correction terms)

But: Currently computationally expensive



Raising Awareness for Uncertainties

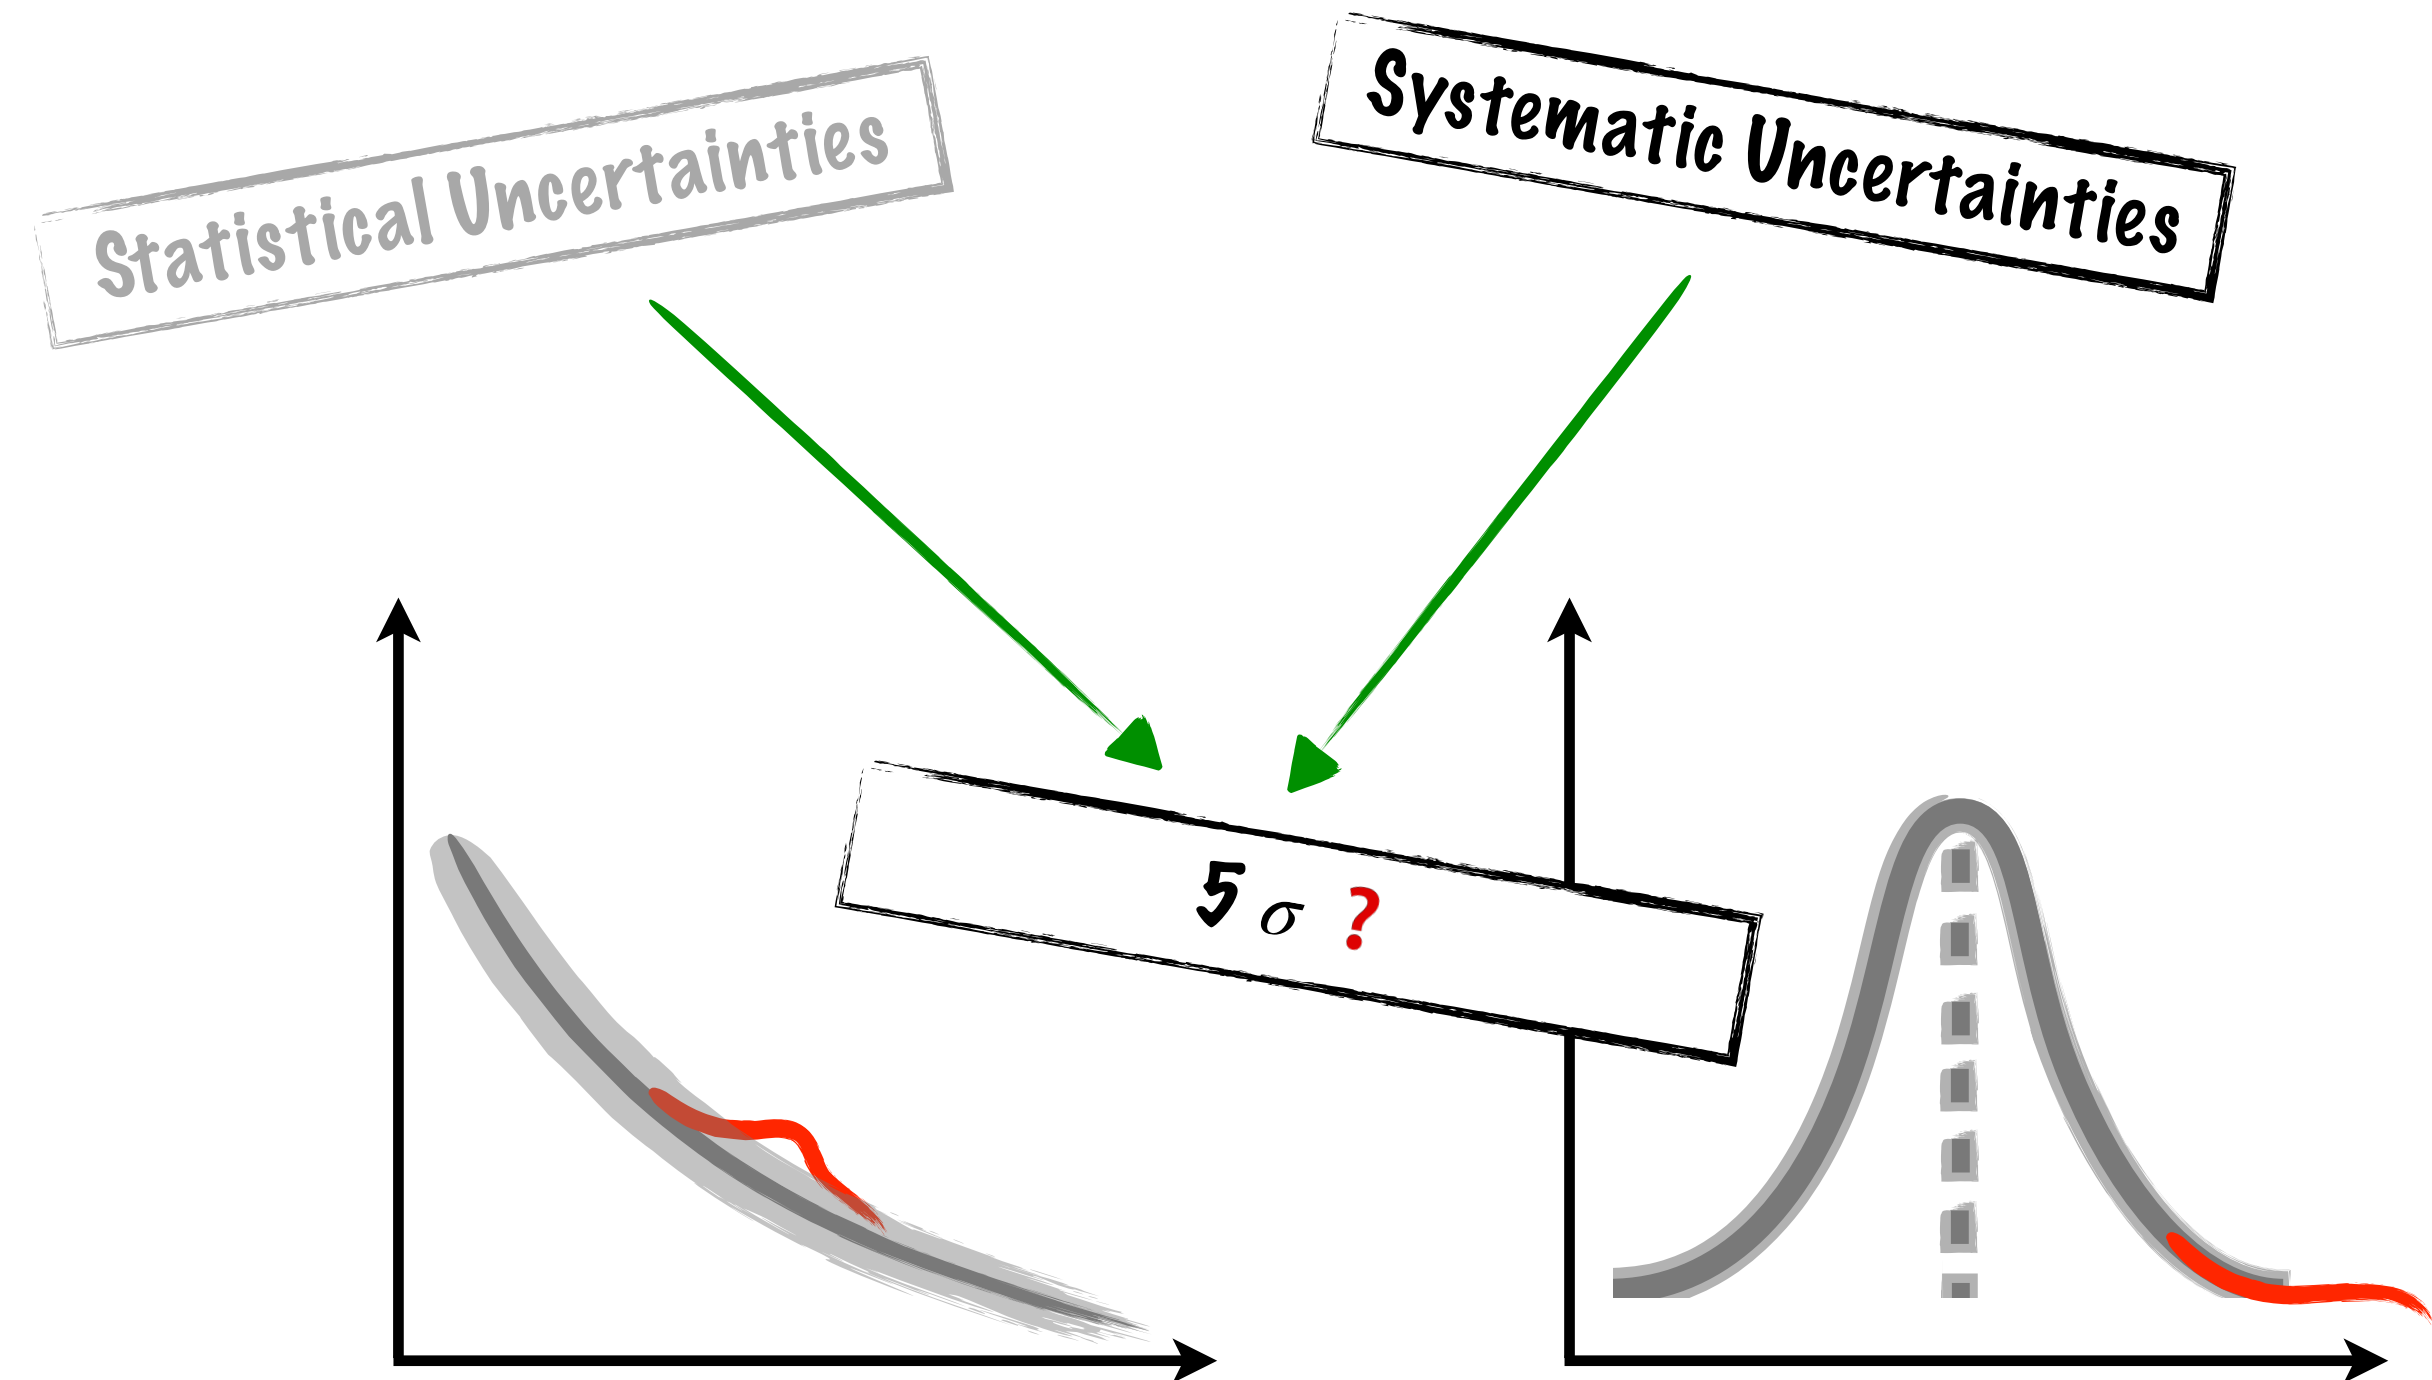
Being precise = estimating uncertainties




Raising Awareness for Uncertainties

Being precise = estimating uncertainties

How can we account for network uncertainties?




What about uncertainties?




Once training is done: \mathbf{W}_1 , \mathbf{W}_2 , \mathbf{W}_3 fixed (*"Network output is deterministic"*)

What about uncertainties?




Bayesianization: We draw each entry from \mathbf{W}_1 , \mathbf{W}_2 , \mathbf{W}_3 from distribution $q(\theta | \mu_\phi, \sigma_\phi)$

What about uncertainties?



Bayesianization: We draw each entry from $\mathbf{W}_1, \mathbf{W}_2, \mathbf{W}_3$ from distribution $q(\theta | \mu_\phi, \sigma_\phi)$

What about uncertainties?




$$\langle \vec{z} \rangle = \frac{1}{N} \sum_i \vec{z}_i$$

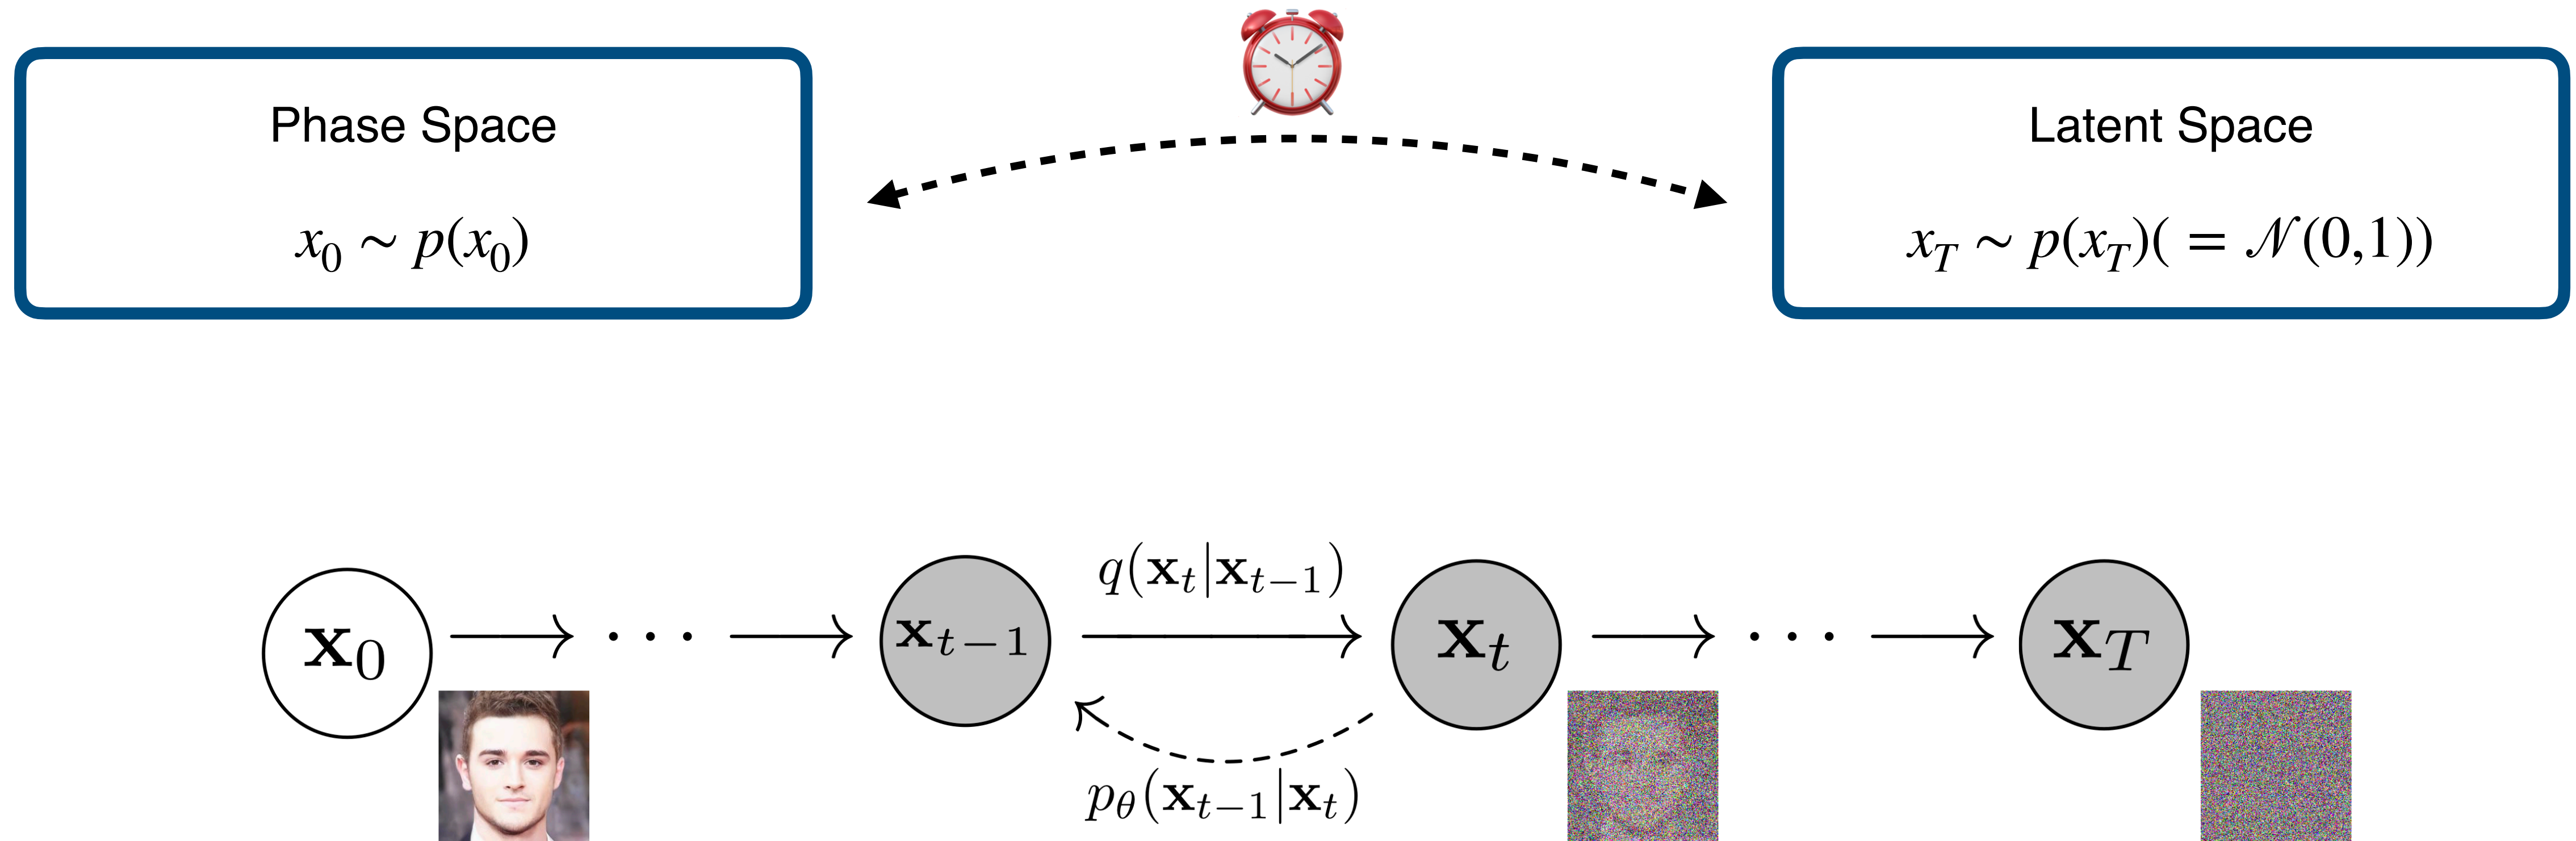
$$\sigma_{pred}^2 = \frac{1}{N} \sum_{i=1}^N (\langle \vec{z} \rangle - \vec{z}_i)^2$$

How to Bayesianize


1. Replace each linear layer with a *Bayesian* layer
2. Add additional regularisation term to likelihood loss



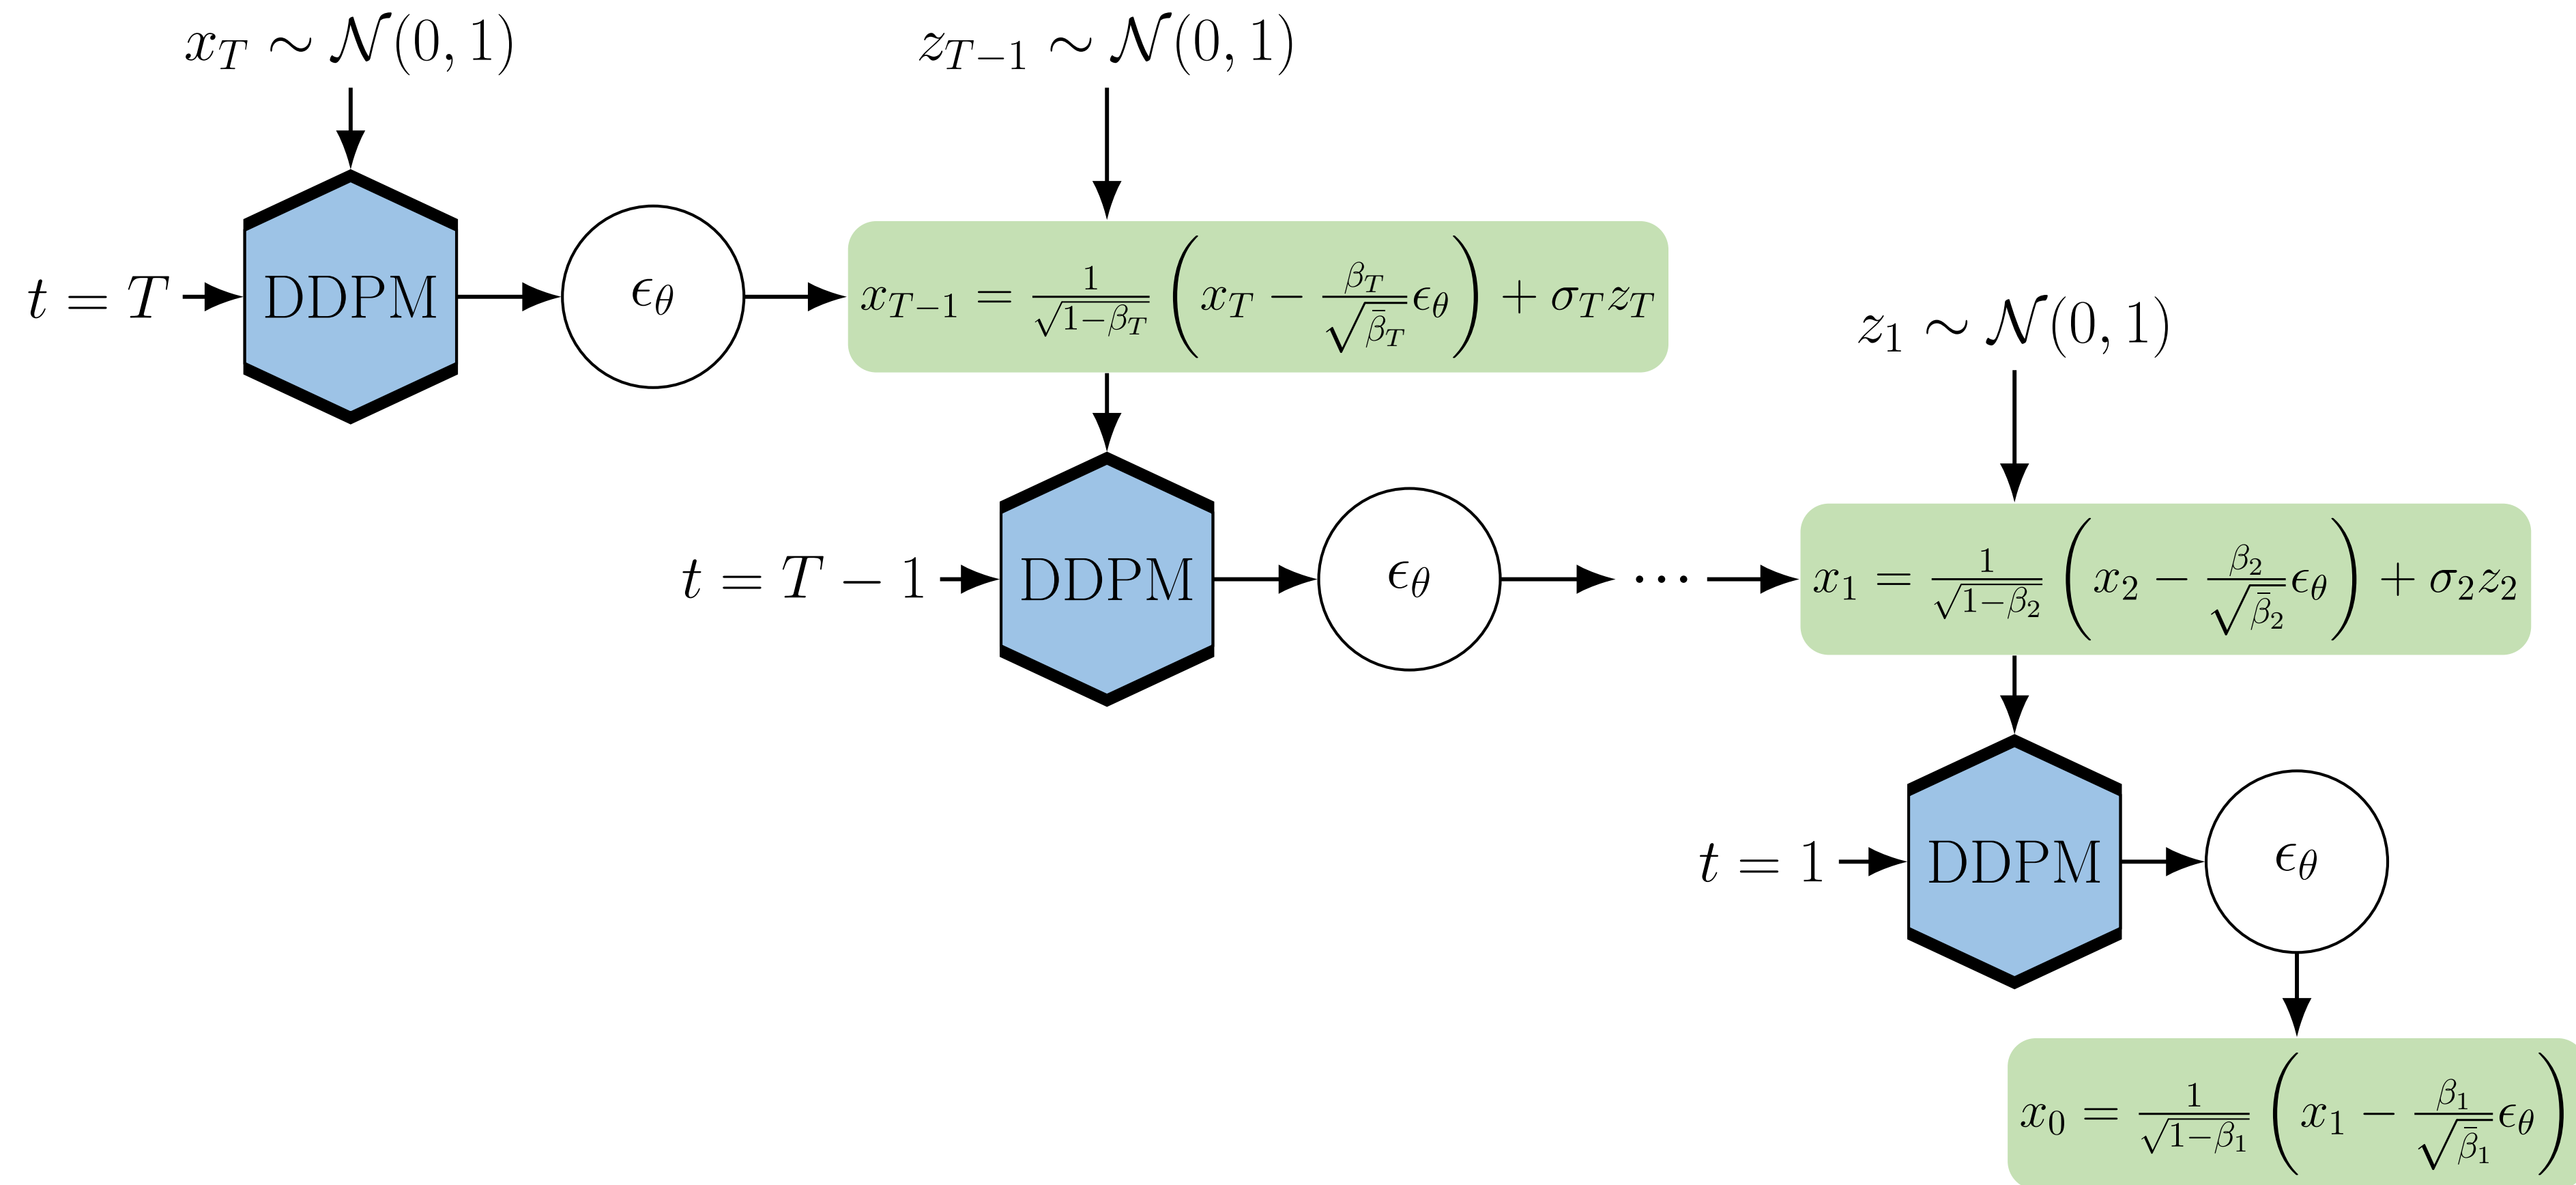
Diffusion Models (DDPM)




Diffusion Models (DDPM)



Diffusion Models (DDPM)




How to Bayesianize - DDPM



Likelihood Loss

$$\mathcal{L}_{DDPM} = -\log p_{\theta} \approx \frac{1}{2\sigma_t^2} \frac{\beta_t^2}{(1 - \beta_t)\bar{\beta}_t} |\epsilon(t) - \epsilon_{\theta}(t)|^2$$

How to Bayesianize - DDPM



Likelihood Loss


$$\mathcal{L}_{B-DDPM} = \left\langle \frac{1}{2\sigma_t^2} \frac{\beta_t^2}{(1 - \beta_t)\bar{\beta}_t} |\epsilon(t) - \epsilon_{\theta}(t)|^2 \right\rangle_{\theta \sim q(\theta)}$$

Prior weight distribution

Learned weight distribution

$$\mathcal{D}_{KL}(q(\theta), p(\theta))$$

How to Bayesianize - DDPM



Likelihood Loss

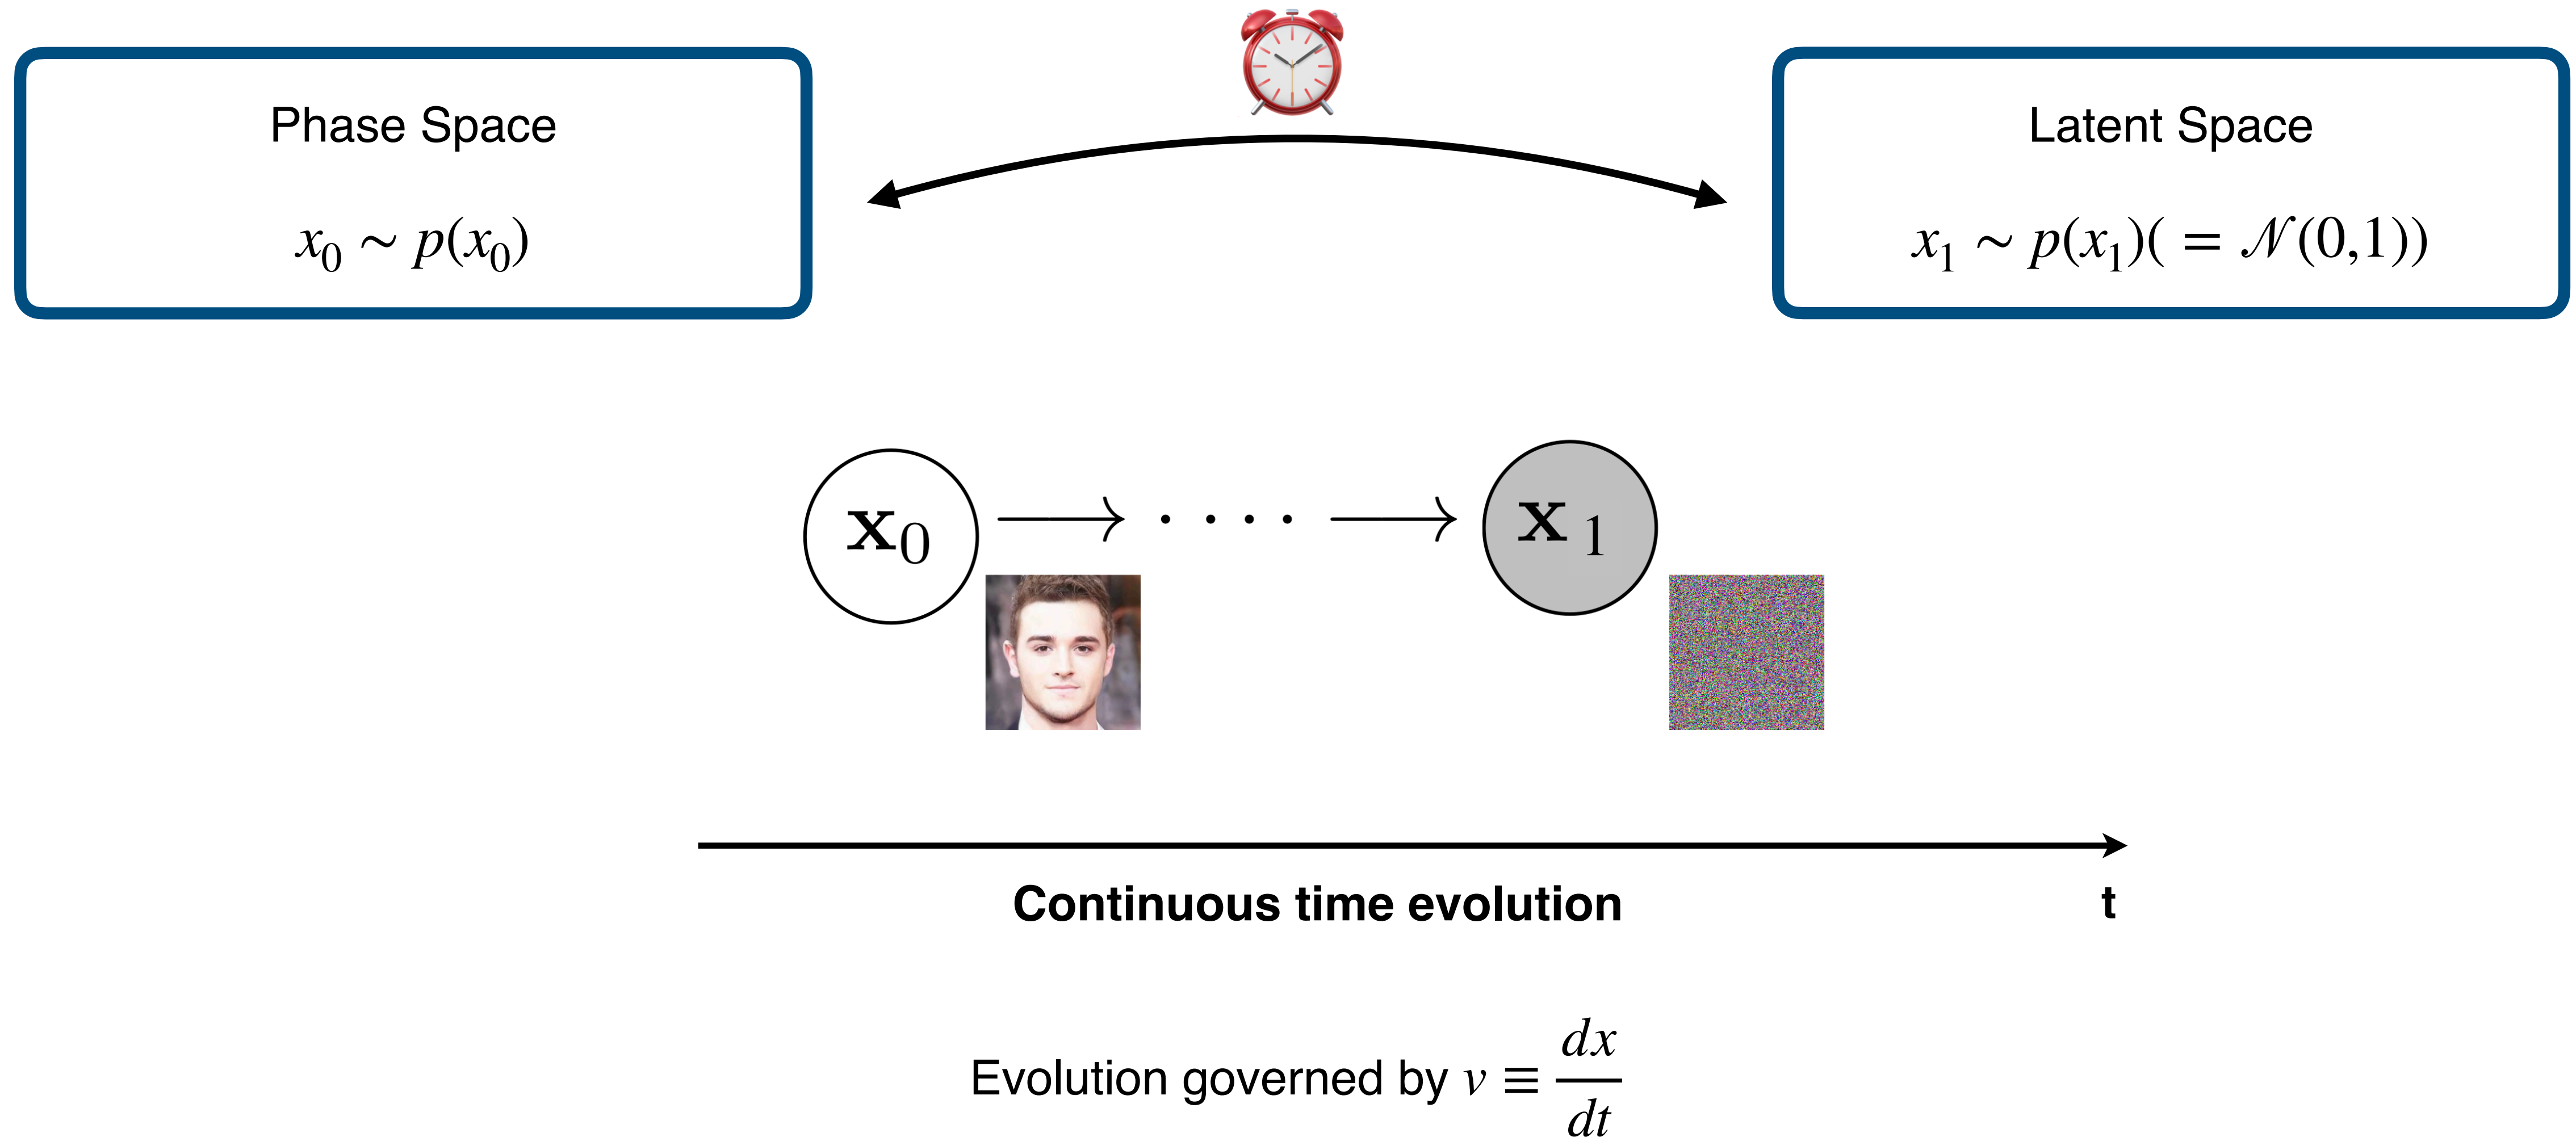
$$\mathcal{L}_{B-DDPM} = T \cdot \left\langle \frac{1}{2\sigma_t^2} \frac{\beta_t^2}{(1 - \beta_t)\bar{\beta}_t} |\epsilon(t) - \epsilon_\theta(t)|^2 \right\rangle_{\theta \sim q(\theta), t \sim \mathcal{U}(0, T)}$$

Prior weight distribution


Learned weight distribution

$$\mathcal{D}_{KL}(q(\theta), p(\theta))$$


Diffusion Models (CFM)



Diffusion Models (CFM)



How to Bayesianize - CFM




No Likelihood Loss

$$\mathcal{L}_{CFM} = (v_\theta - (x_1 - x_0))^2$$

+ ? ? ?

How to Bayesianize - CFM



No Likelihood Loss

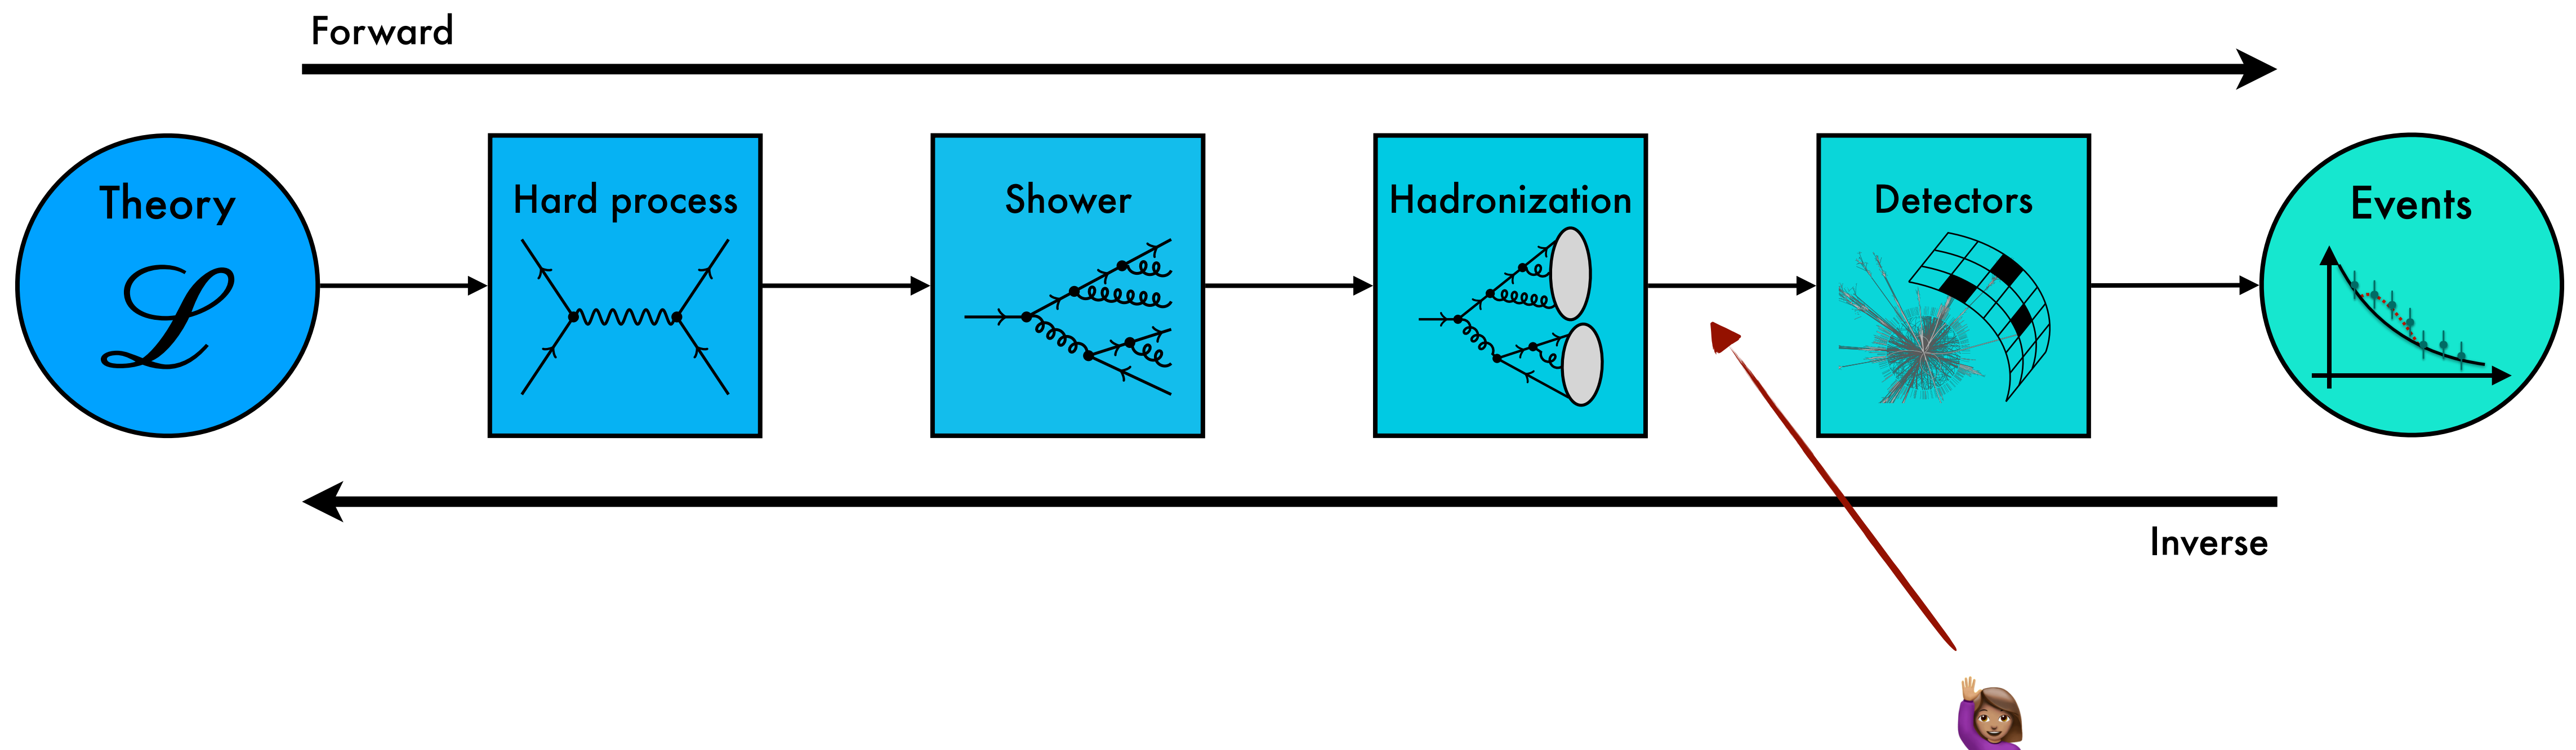
$$\mathcal{L}_{CFM} = \left\langle (v_\theta - (x_1 - x_0))^2 \right\rangle_{\theta \sim q(\theta)}$$

Hyperparameter

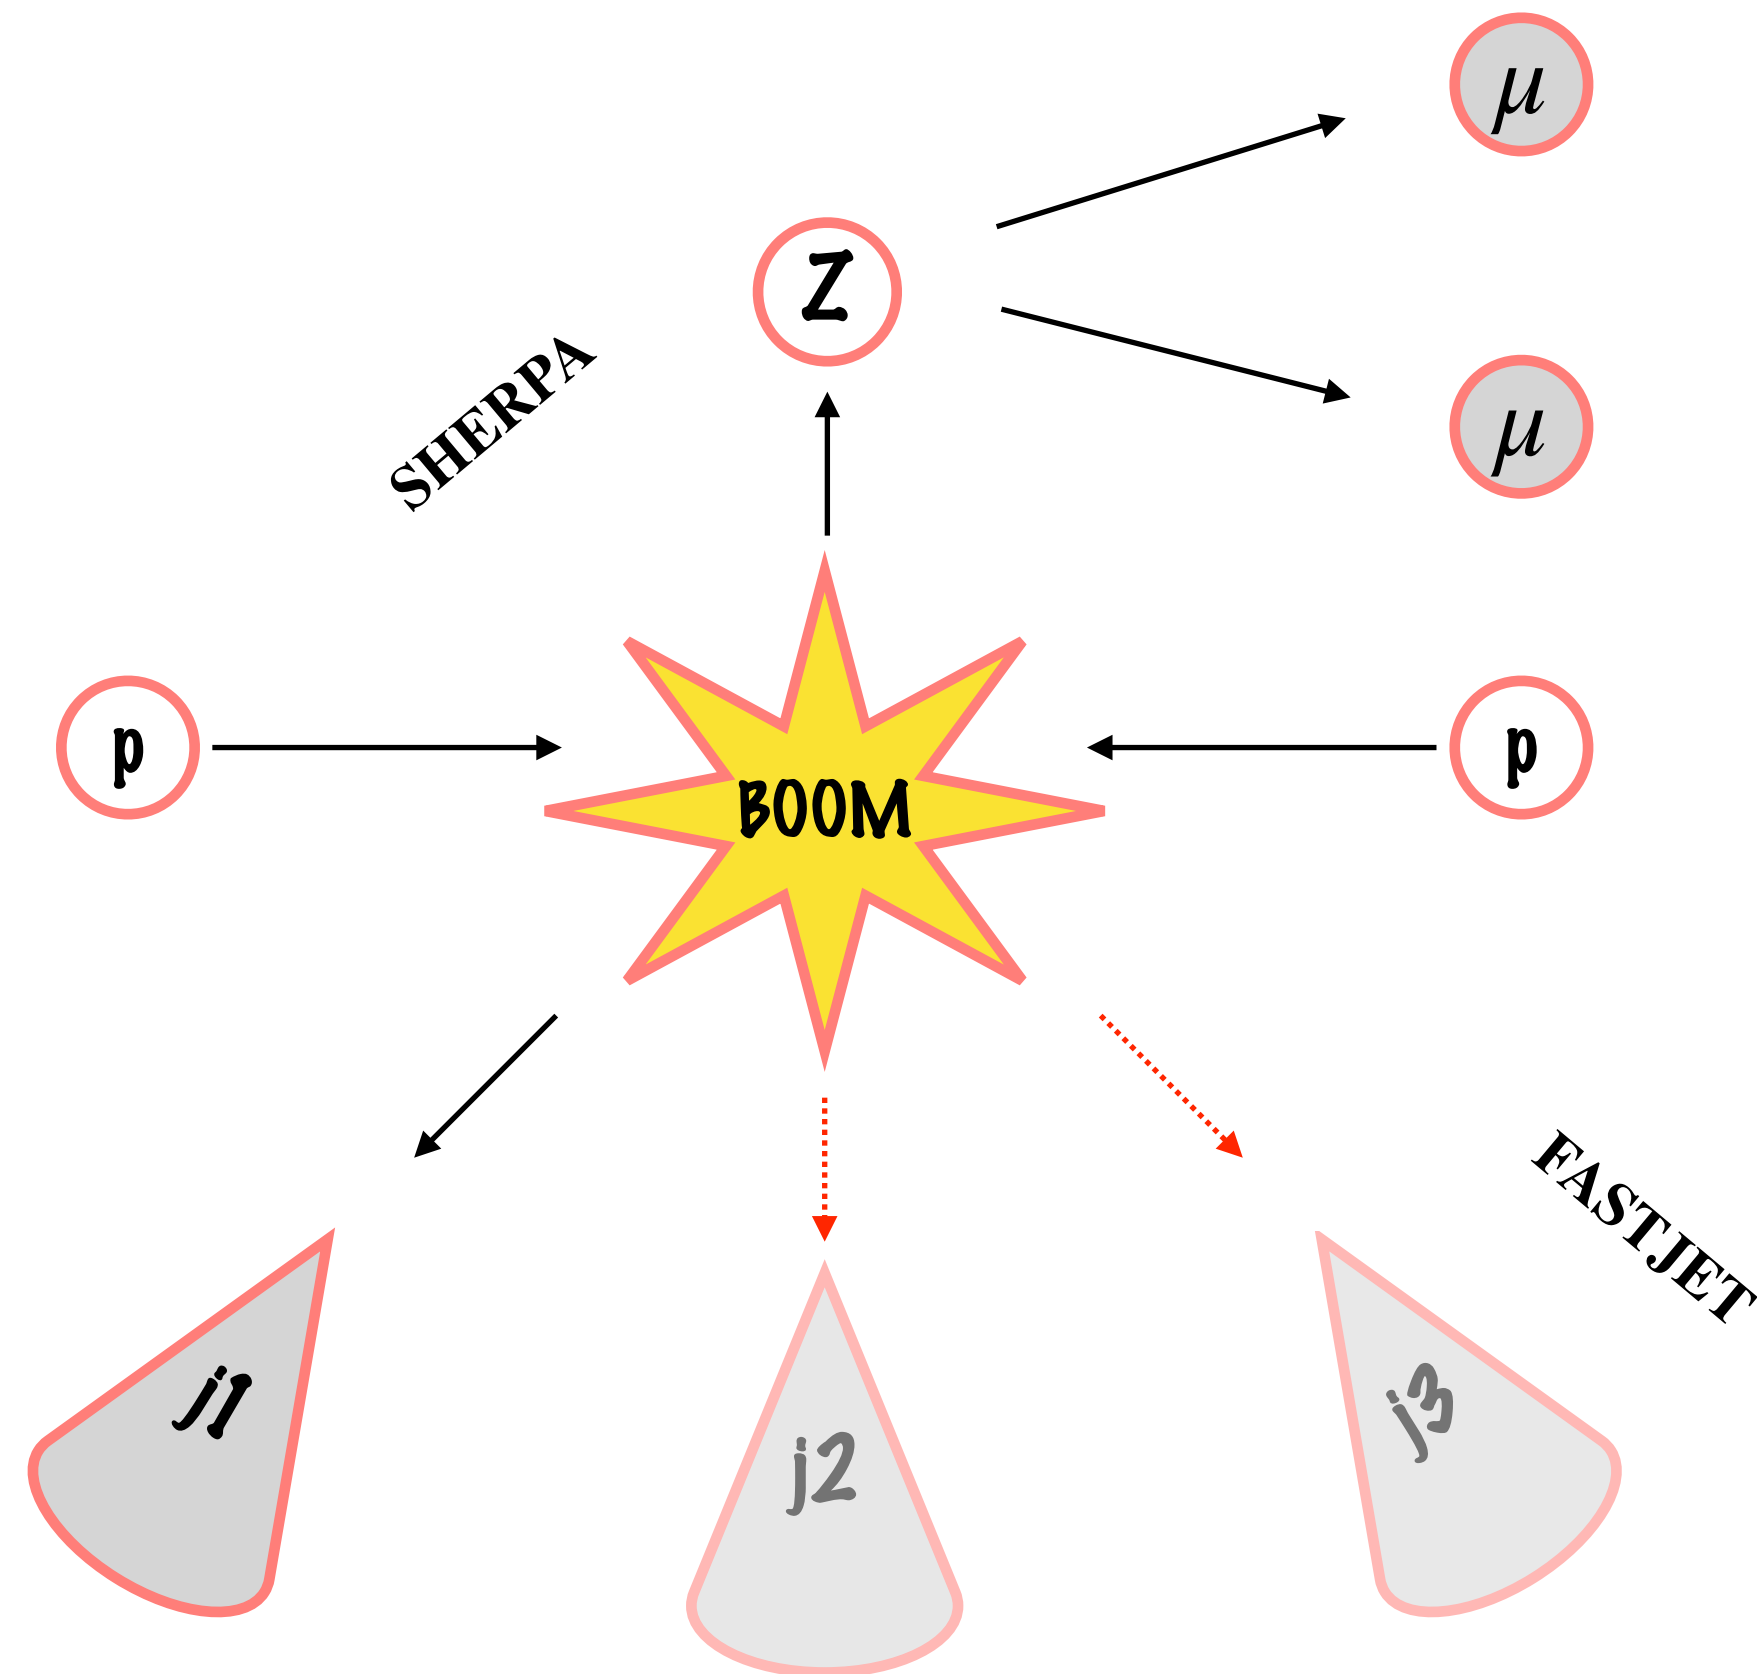
+

$$c \cdot \mathcal{D}_{KL}(q(\theta), p(\theta))$$

Choosing Test Phase-Space



Concrete Application — LHC



$Z (\rightarrow \mu\mu) + \text{jets}$:

3 - 5 final state particles (including jets)

12 - 20 dimensional phase space

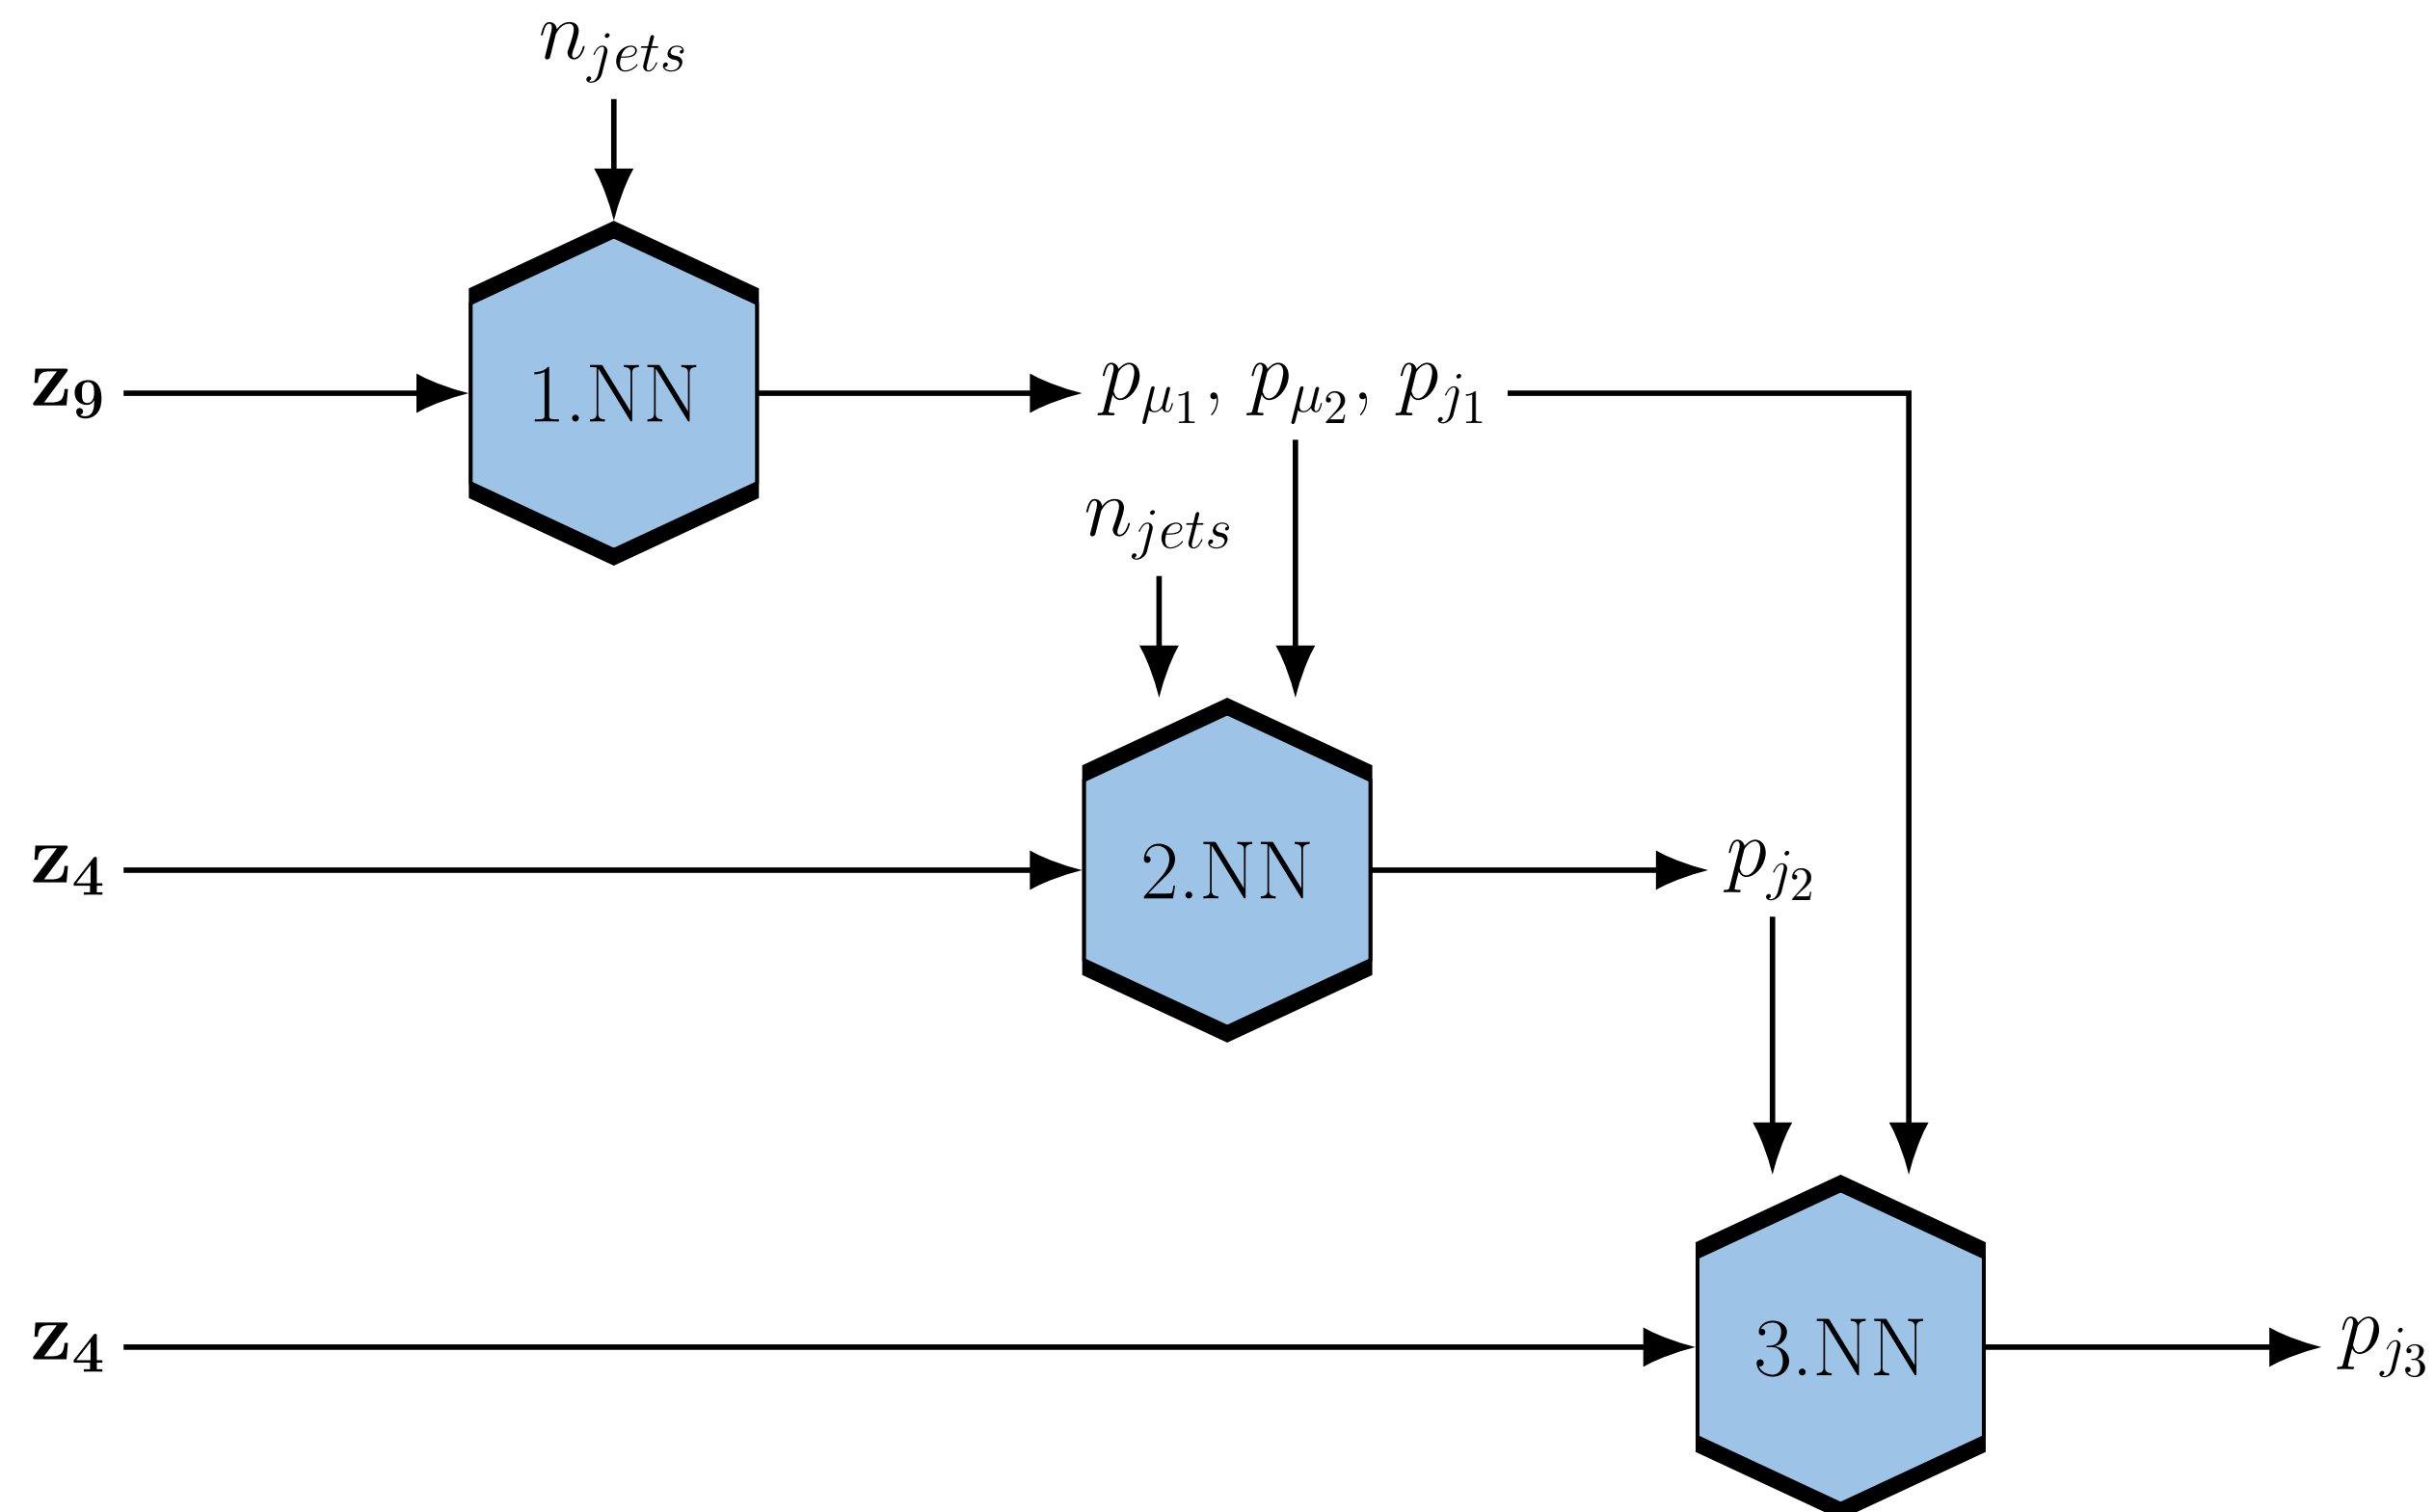
Smart preprocessing:

Global Phase Shift

Drop muon masses

→ reduces phase space to
9 - 17 dimensions

Concrete Application — LHC



$Z (\rightarrow \mu\mu) + jets$:

3 - 5 final state particles (including jets)

12 - 20 dimensional phase space

Smart preprocessing:


Global Phase Shift

Drop muon masses

→ reduces phase space to 9 - 17 dimensions


To be precise

Percent level precision (comparable to statistical uncertainty)



To be precise

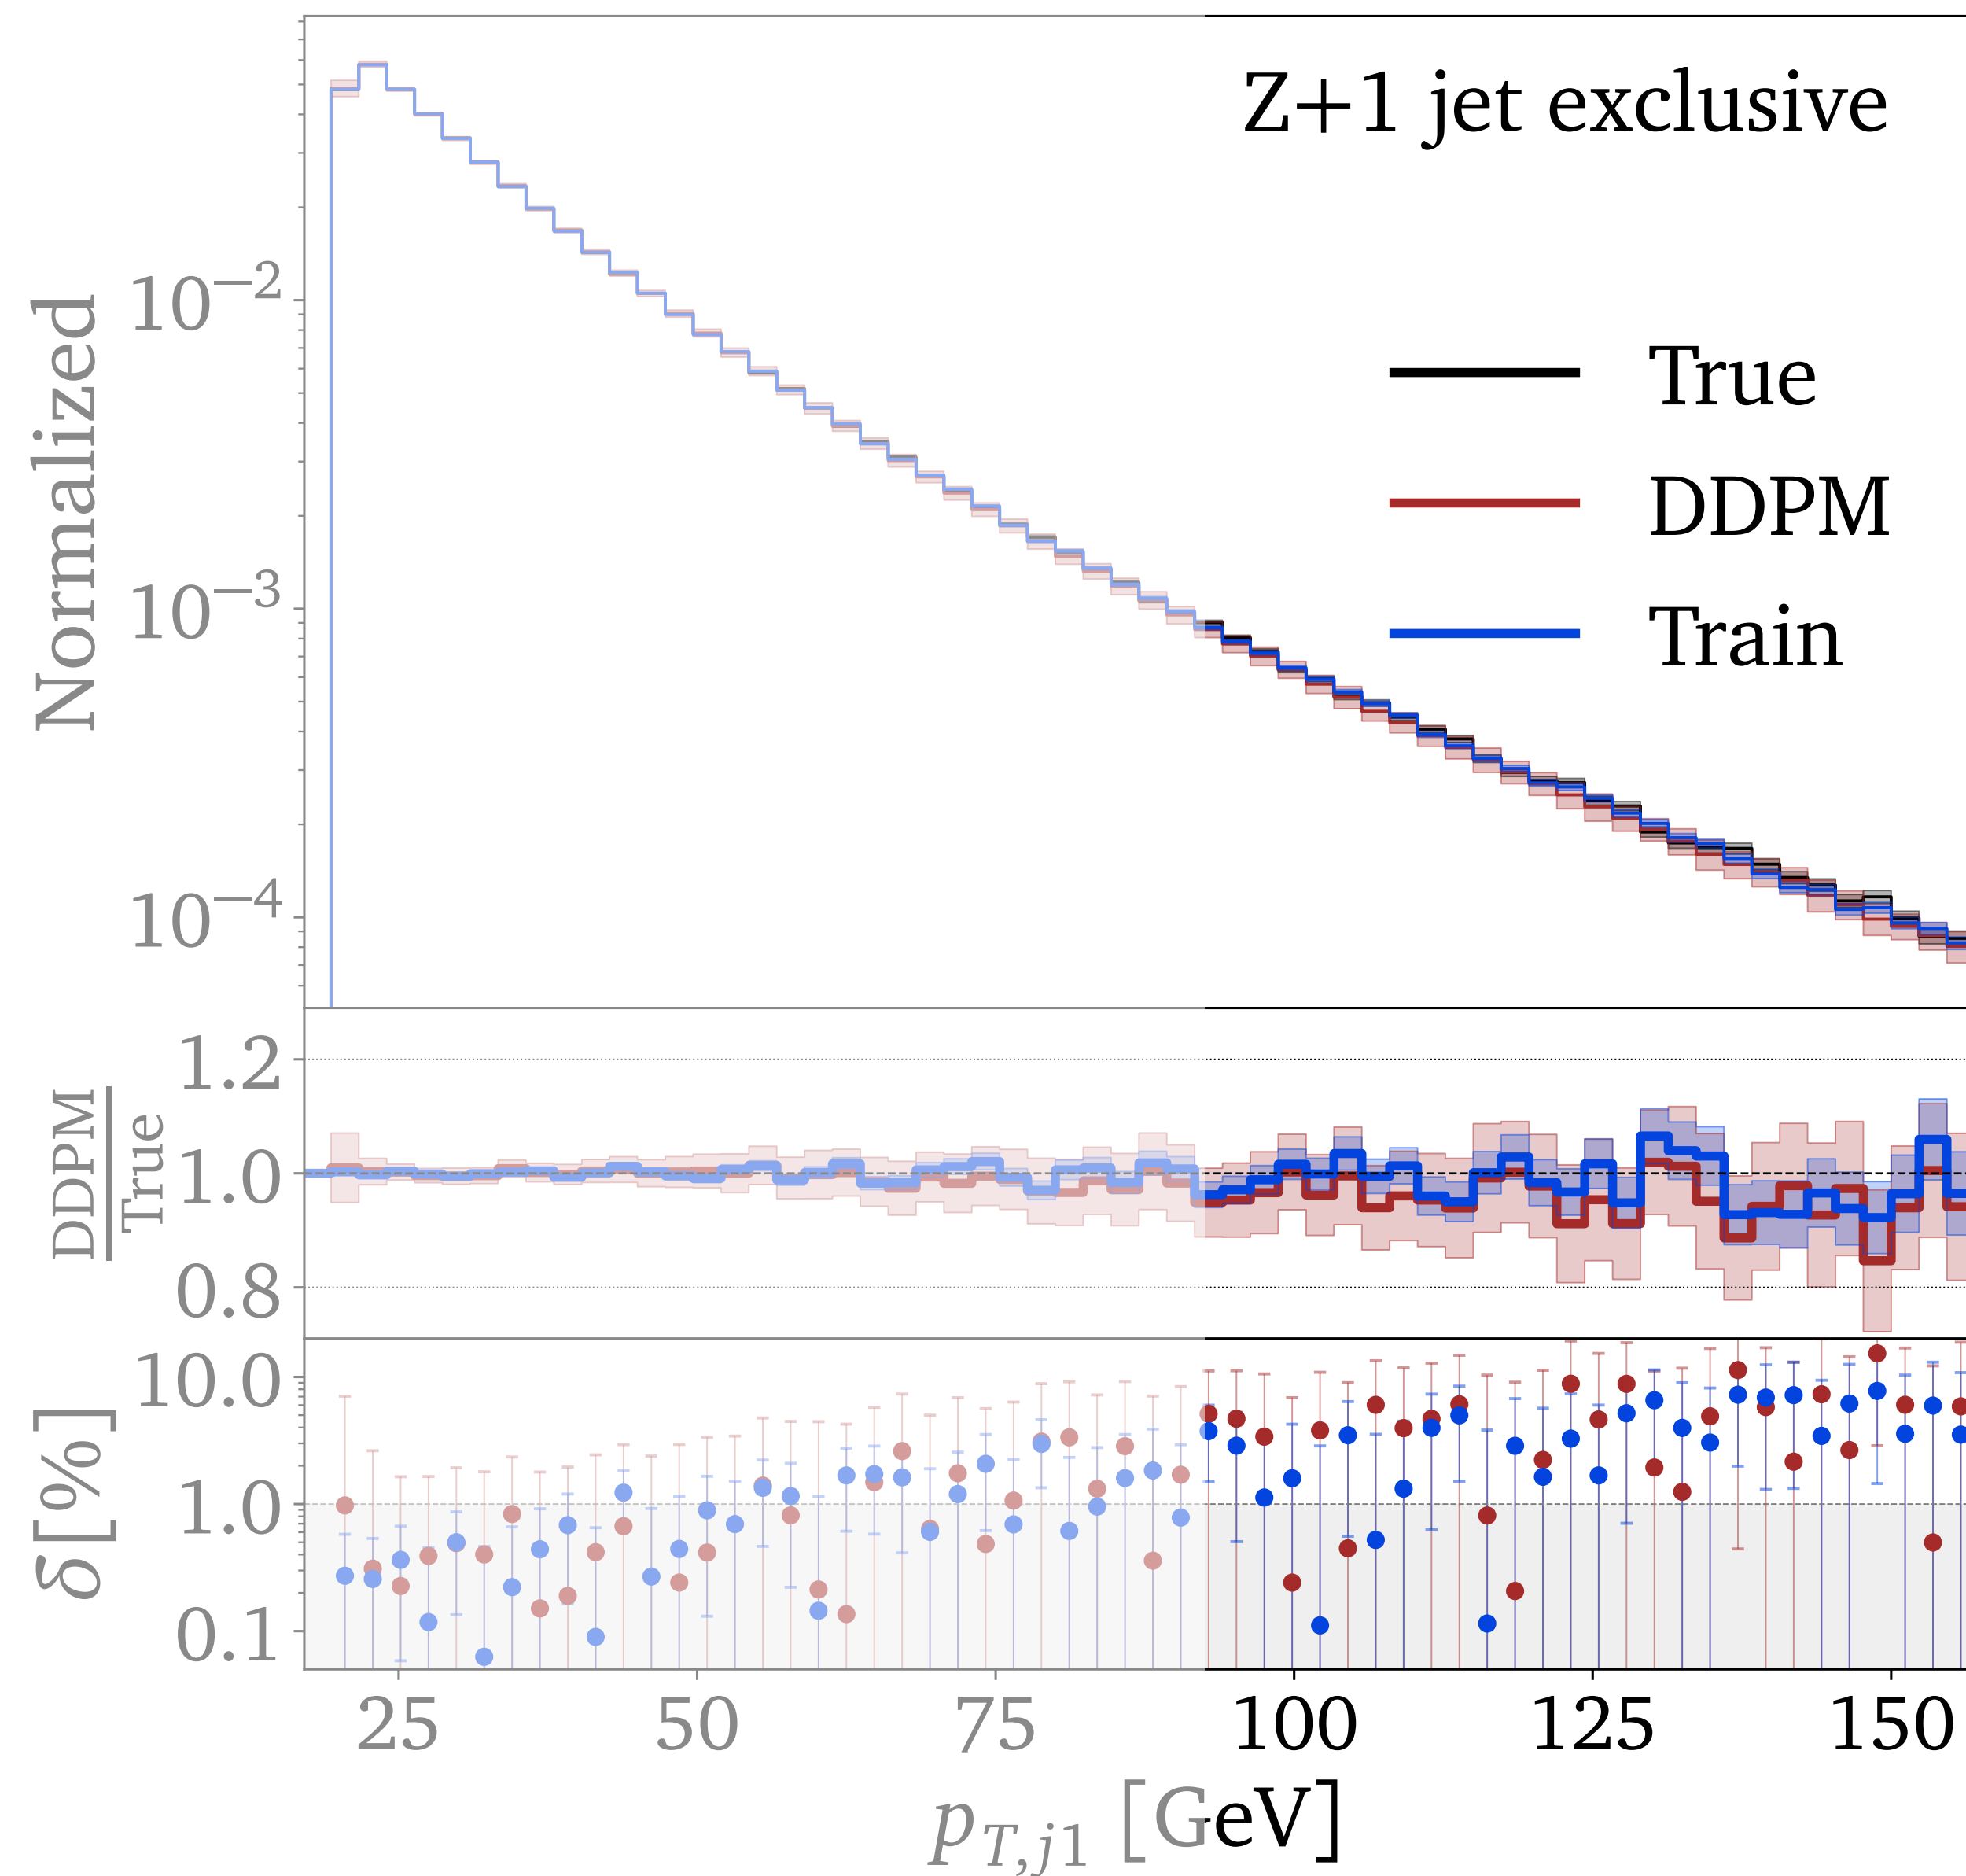
Percent level precision (comparable to statistical uncertainty)



To be precise

Percent level precision (comparable to statistical uncertainty)

Uncertainty well-defined

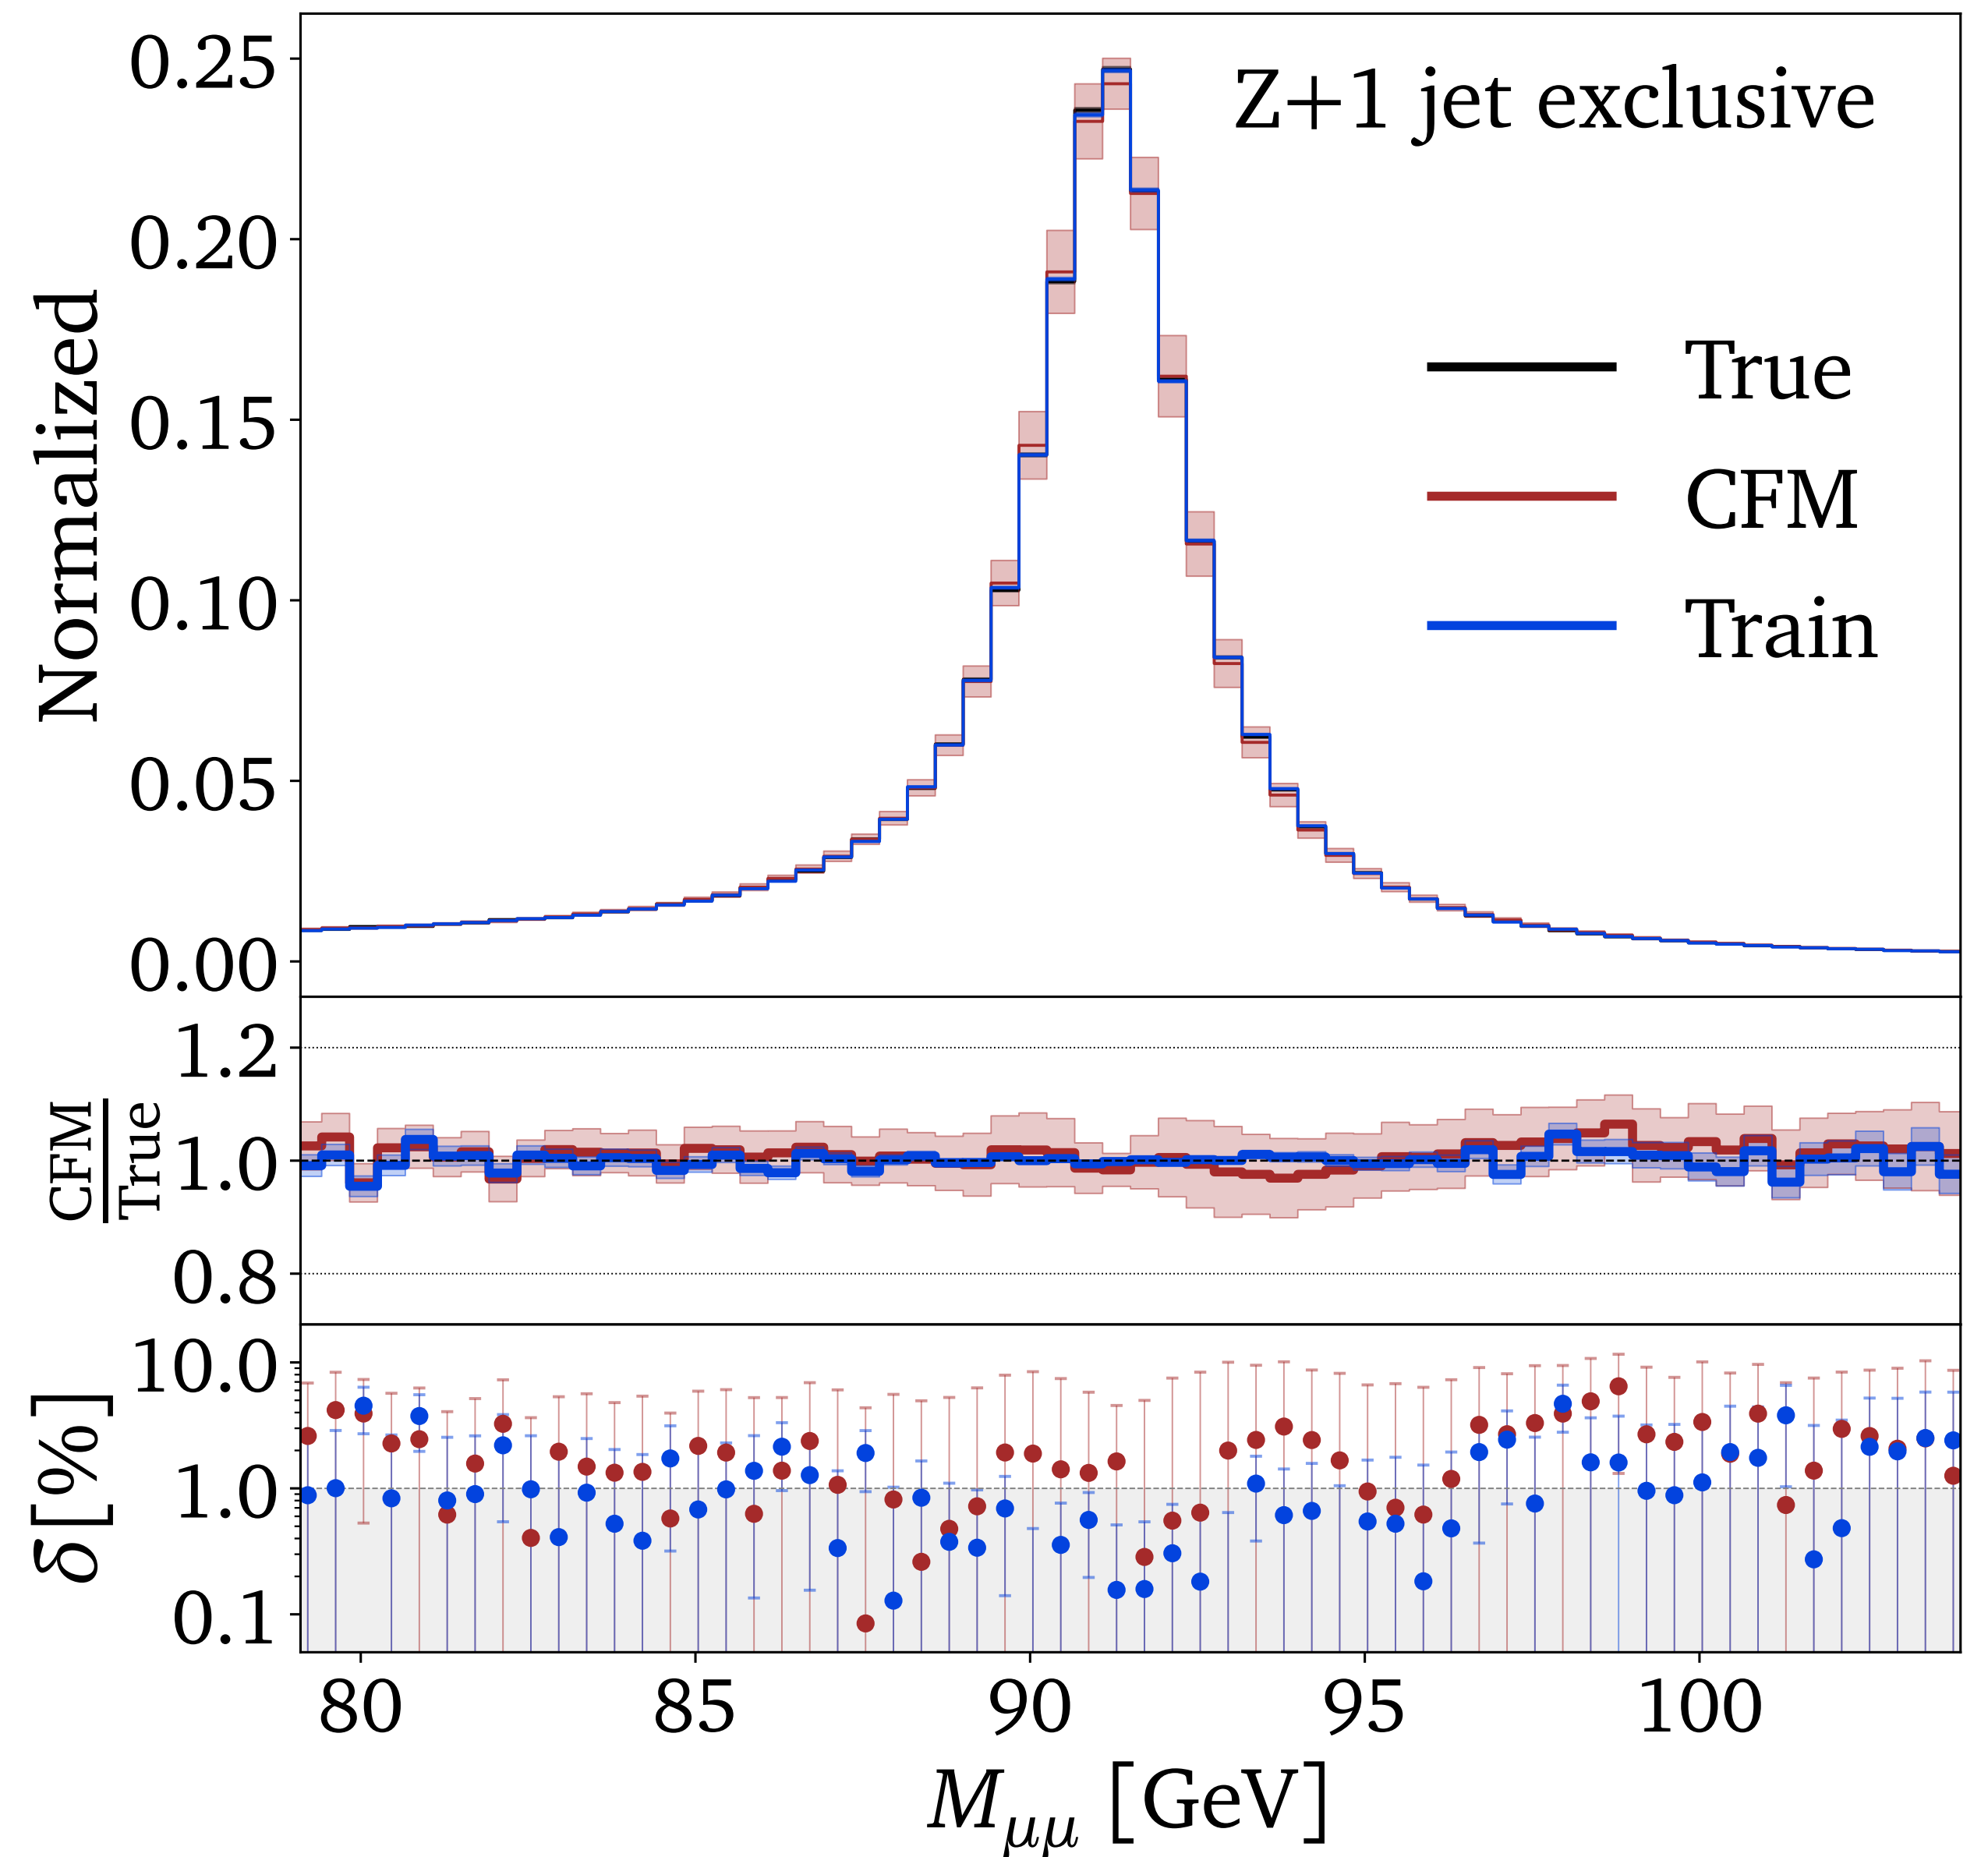


To be precise

Percent level precision (comparable to statistical uncertainty)

Uncertainty well-defined


Surpasses INN precision (A. Butter et al.: arXiv:2110.13632)



To be precise

By construction:

Systematic failure modes of network
not covered by Bayesian uncertainties




And now what?

Diffusion Models can compete with current benchmark (precision-wise)

A lot of on-going research (generation speed up, precision, etc.)

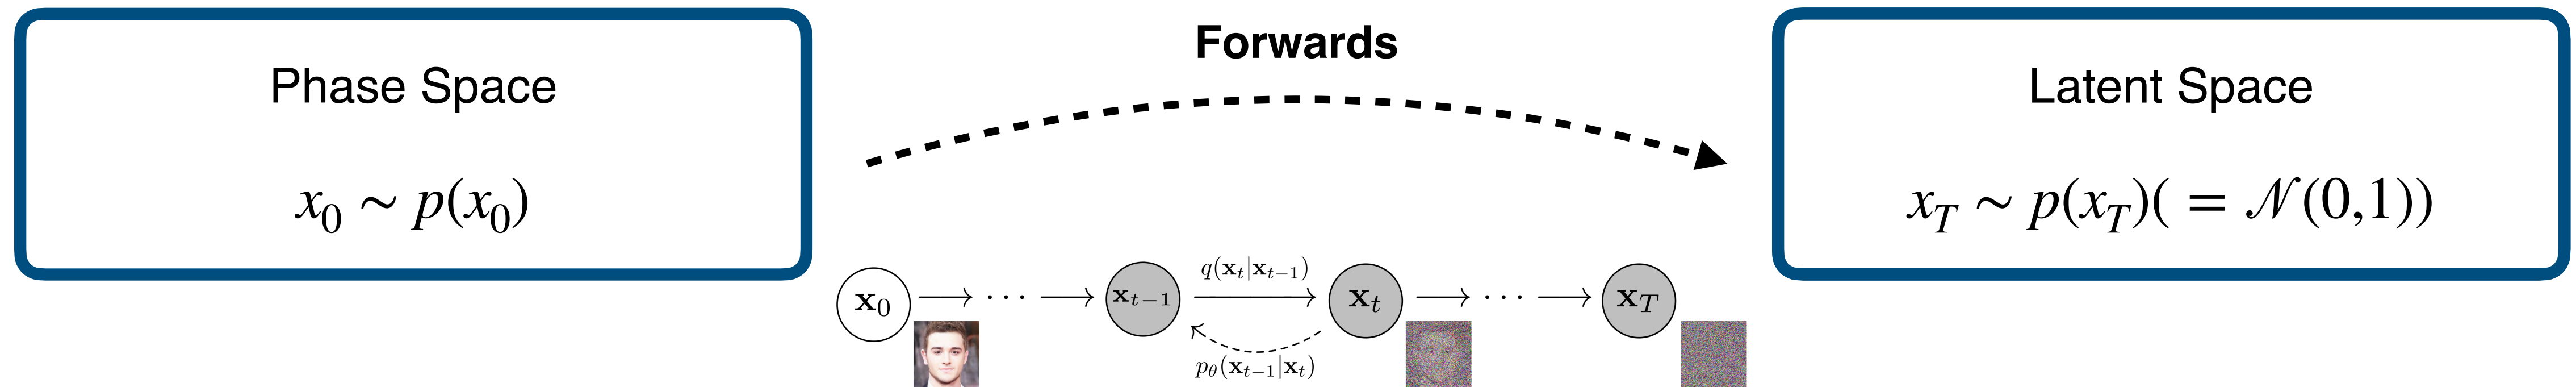
Bayesian Versions seem to estimate training uncertainty correctly



(Bayesian) Diffusion models show potential to be applied to particle physics tasks

Backup


Diffusion Models (DDPM)



$$q(x_1, \dots, x_T | x_0) = \prod_{t=1}^T q(x_t | x_{t-1})$$

$$q(x_t | x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t} x_{t-1}, \beta_t) \quad \text{where } \beta_t \text{ follows noise scheduler}$$


Diffusion Models (DDPM)



To reverse:
$$q(x_{t-1} | x_t) = \frac{q(x_t | x_{t-1})q(x_{t-1})}{q(x_t)}$$

...but we don't know $q(x_t)$ & $q(x_{t-1})$ 🙄 🙄


Diffusion Models (DDPM)



Instead, learn: $p_\theta(x_{t-1} | x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \sigma_\theta^2(x_t, t))$


$$p(x_0, \dots, x_T | \theta) = p(x_T) \prod_{t=1}^T p_\theta(x_{t-1} | x_t)$$

Diffusion Models (DDPM)




$$\mathcal{L}_{DDPM} = -\log p_\theta(x_0) \approx \frac{1}{2\sigma_t^2} \frac{\beta_t^2}{(1-\beta_t)\bar{\beta}_t} |\epsilon(t) - \epsilon_\theta(t)|^2$$

Diffusion Models (DDPM)

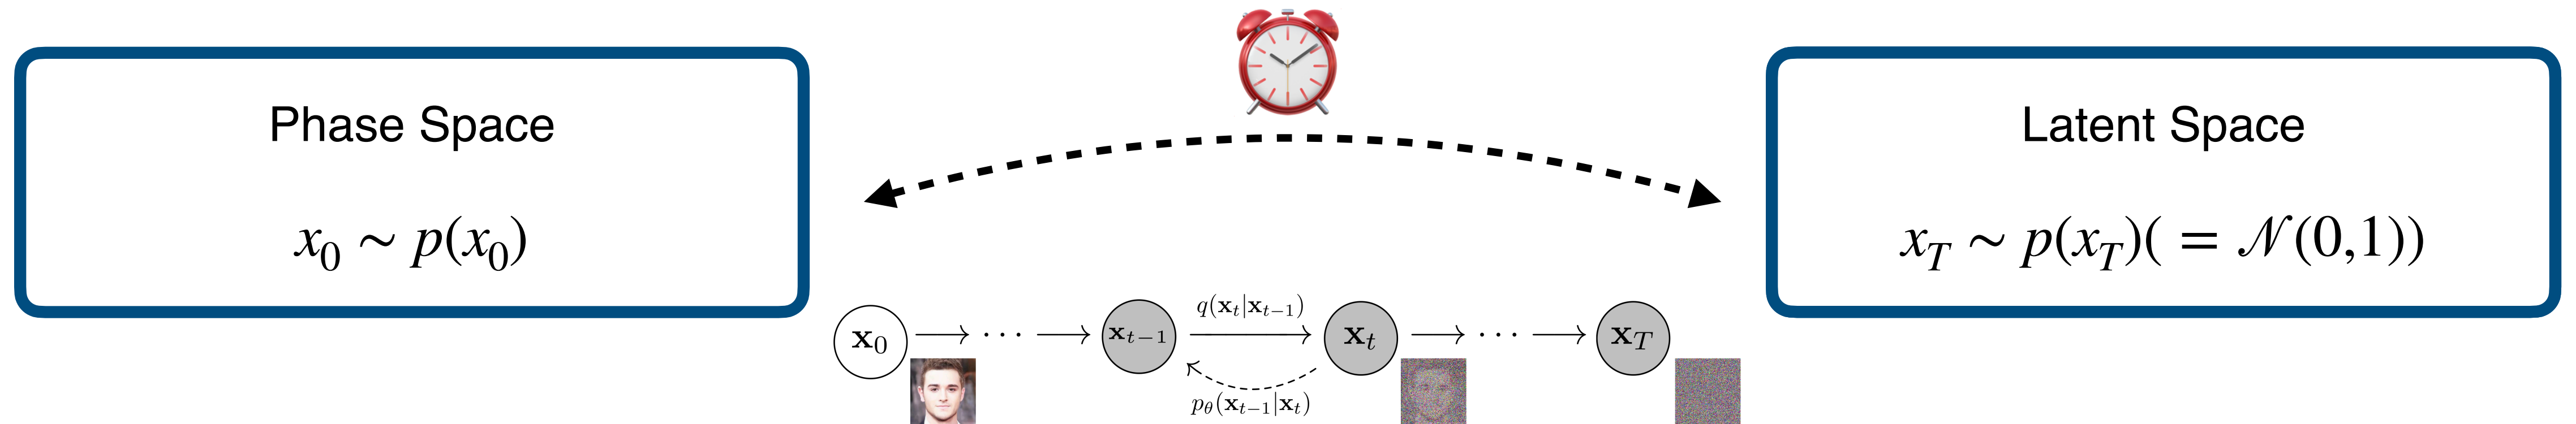


$$\mathcal{L}_{DDPM} = -\log p_\theta(x_0) \approx \frac{1}{2\sigma_t^2} \frac{\beta_t^2}{(1-\beta_t)\bar{\beta}_t} |\epsilon(t) - \epsilon_\theta(t)|^2$$

-  : Reparametrization
- Estimation
- ...

Predicted and actual noise added at time t

Diffusion Models (DDPM)




$$\mathcal{L}_{DDPM} = -\log p_\theta(x_0) \approx \frac{1}{2\sigma_t^2} \frac{\beta_t^2}{(1 - \beta_t)\bar{\beta}_t} |\epsilon(t) - \epsilon_\theta(t)|^2$$



Denoising

Diffusion Models (CFM)

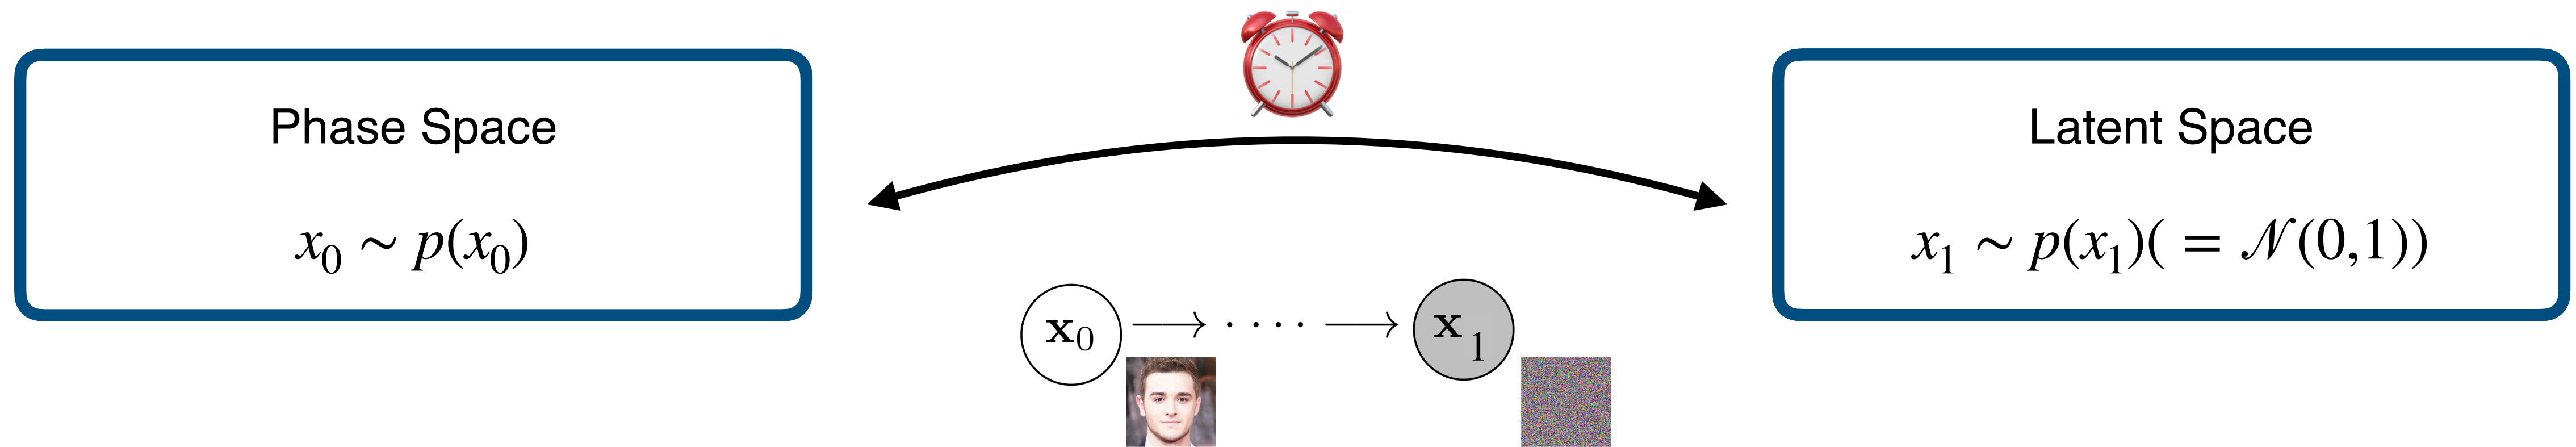


1. Connect x_0 with x_1 by a linear trajectory: $x(t | x_0) = (1 - t)x_0 + tx_1$

2. Derive: $\frac{d}{dt}x(t | x_0) = (x_1 - x_0) \equiv v_\theta(x_1, x_0)$

3. Solve for x_0 (ODE): $x_0 = x_t - \int_0^1 v_\theta(x(t), t)dt$

Diffusion Models (CFM)



1. Connect x_0 with x_1 by a linear trajectory: $x(t | x_0) = (1 - t)x_0 + tx_1$

2. Derive: $\frac{d}{dt}x(t | x_0) = (x_1 - x_0) \equiv v_\theta(x_1, x_0)$

3. Solve for x_0 (ODE): $x_0 = x_t - \int_0^1 v_\theta(x(t), t) dt$

