Constraint-Led Generative Modeling

A unified framework for training, alignment, and composition of diffusion models

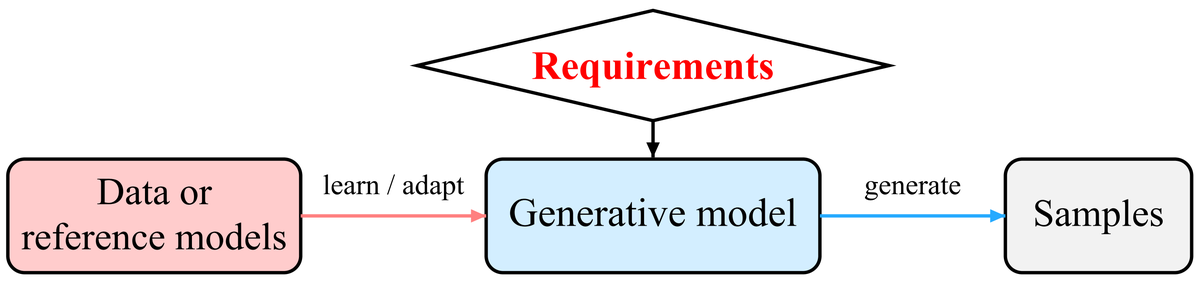
Dongsheng Ding (dongshed@utk.edu)

THE UNIVERSITY OF TENNESSEE KNOXVILLE

Explicit requirements determine how models are trained, aligned, and combined.

Context

Requirement-driven generative modeling



Requirements: fidelity, target properties, physical consistency.

How can we combine generative models while satisfying competing requirements?

Constrained generation

$$\begin{aligned} \text{maximize} \quad & \text{expected scientific utility} \\ \text{subject to} \quad & \text{expected property score} \ \geq \ \text{target} \\ & \text{distance to reference models} \ \leq \ \text{tolerance} \end{aligned}$$

Challenge: **Nonconvex model optimization**

Unified framework

Constrained distribution optimization

$$\begin{aligned} \underset{p}{\text{minimize}} \quad & D_{\text{diff}}\left(p; q\right) \\ \text{subject to} \quad & D_{\text{diff}}\left(p; q^i\right) \ \leq \ \epsilon_i \quad \text{for } i = 1, \ldots, m \\ & \mathbb{E}_{x \sim p(\cdot)}\left[\, r_j(x) \,\right] \ \geq \ b_j \quad \text{for } j = 1, \ldots, n \end{aligned}$$

- p – optimized distribution
- q, q^i – data/reference distributions
- $D_{\text{diff}}\left(p; q\right)$ – divergence from reference q
- $r_j(x)$ – jth reward

settings	objective	constraints
composition	$D_{\text{KL}}\left(p(x_0) \,\|\, q(x_0)\right)$	$D_{\text{KL}}\left(p(x_0) \,\|\, q^i(x_0)\right) \leq \epsilon_i$
alignment	$D_{\text{KL}}\left(p(x_{0:T}) \,\|\, q(x_{0:T})\right)$	$\mathbb{E}_{x_0 \sim p(\cdot)}\left[\, r_i(x_0) \,\right] \geq b_i$
training	$D_{\text{KL}}\left(q(x_{0:T}) \,\|\, p(x_{0:T})\right)$	$D_{\text{KL}}\left(q^i(x_{0:T}) \,\|\, p(x_{0:T})\right) \leq \epsilon_i$

Property: **Infinite-dimensional, convex**

Lagrangian-based dual approach

$$\begin{aligned} L(p, \lambda) \ = \ & D_{\text{diff}}\left(p; q\right) + \sum_{i=1}^{m} \lambda_i \left(D_{\text{diff}}\left(p; q^i\right) - \epsilon_i\right) \\ & + \sum_{j=1}^{n} \lambda_{m+j} \left(b_j - \mathbb{E}_{x \sim p(\cdot)}\left[\, r_j(x) \,\right]\right) \end{aligned}$$

- $\lambda \geq 0$ – Lagrangian multiplier or dual variable
- $D(\lambda) = \text{minimize}_p \, L(p; \lambda)$ – dual function
- Convex strong duality yields **optimal recovery:**

$$\lambda^\star \ \in \ \underset{\lambda \geq 0}{\arg\max} \ D(\lambda) \qquad\qquad p^\star \ \in \ \underset{p}{\arg\min} \ L(p, \lambda^\star)$$

Challenge: **Strong duality may fail after parameterization.**

parameterized model p_θ, w/ parameter θ

- Local strong concavity of dual function establishes:

Claim: Dual-based approach finds an optimal model, up to an approximation error.

Collaboration: Seeking scientific applications with reusable generative models and explicit requirements.

Constraint-Led Model Composition

Composition of diffusion models

KHD†R, NeurIPS '25

$$\begin{aligned} \underset{p}{\text{minimize}} \quad & D_{\text{KL}}\left(p(x_0) \,\|\, q(x_0)\right) \\ \text{subject to} \quad & D_{\text{KL}}\left(p(x_0) \,\|\, q^i(x_0)\right) \ \leq \ \epsilon_i \quad \text{for } i = 1, \ldots, m \end{aligned}$$

- $q^\lambda_{\text{prod}}(\cdot) \propto (q(\cdot))^{\frac{1}{1+\mathbf{1}^\top \lambda}} \prod_i (q^i(\cdot))^{\frac{\lambda_i}{1+\mathbf{1}^\top \lambda}}$ – product distribution
- **optimal recovery:**

$$p^\star \ \in \ \underset{p}{\arg\min} \ D_{\text{KL}}\left(p(x_0) \,\|\, q^{\lambda^\star}_{\text{prod}}(x_0)\right)$$

Constraints promote closeness to each model.

	Min. CLIP (↑)	Min. BLIP (↑)
prompting	21.52	0.206
equal weights	22.18	0.203
constrained	**22.45**	**0.221**



Above: Average min. CLIP/BLIP scores over the composed concepts. **Right:** R1: 'volcano', 'pineapple'; R2: 'donut', 'turtle'; R3: 'dandelion', 'spider-web', 'cinnamon roll'.

Constraint-Led Reward Alignment

Alignment of diffusion models

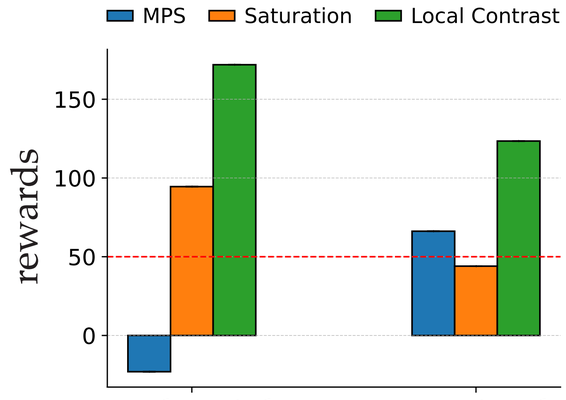
KHD†R, NeurIPS '25

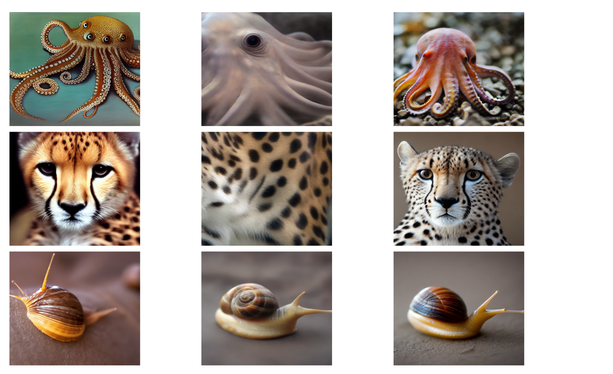
$$\begin{aligned} \underset{p}{\text{minimize}} \quad & D_{\text{KL}}\left(p(x_{0:T}) \,\|\, q(x_{0:T})\right) \\ \text{subject to} \quad & \mathbb{E}_{x_0 \sim p(\cdot)}\left[\, r_j(x_0) \,\right] \ \geq \ b_j \quad \text{for } j = 1, \ldots, n \end{aligned}$$

- $q^\lambda_{\text{rw}}(\cdot) \propto q(\cdot) \mathrm{e}^{\lambda^\top r}$ – exponentially weighted distribution
- **optimal recovery:**

$$p^\star \ \in \ \underset{p}{\arg\min} \ D_{\text{KL}}\left(p(x_{0:T}) \,\|\, q^{\lambda^\star}_{\text{rw}}(x_{0:T})\right)$$

Constraints balance multiple reward criteria.





Constraint-Led Training

Constrained diffusion models

KD†R, NeurIPS '24

$$\begin{aligned} \underset{p}{\text{minimize}} \quad & D_{\text{KL}}\left(q(x_{0:T}) \,\|\, p(x_{0:T})\right) \\ \text{subject to} \quad & D_{\text{KL}}\left(q^i(x_{0:T}) \,\|\, p(x_{0:T})\right) \ \leq \ \epsilon_i \quad \text{for } i = 1, \ldots, m \end{aligned}$$

- $q^\lambda_{\text{mix}}(x_0) = \frac{1}{1+\mathbf{1}^\top \lambda}(q(x_0) + \sum_i \lambda_i q^i(x_0))$ – mixture distribution
- **optimal recovery:**

$$p^\star \ \in \ \underset{p}{\arg\min} \ D_{\text{KL}}\left(q^{\lambda^\star}_{\text{mix}}(x_{0:T}) \,\|\, p(x_{0:T})\right)$$

Constraints improve balance across groups.

