Composition and Alignment of Diffusion Models using Constrained Learning

Shervin Khalafi[1], Ignacio Hounie[1], Dongsheng Ding[2], Alejandro Ribeiro[1]

[1] University of Pennsylvania, [2] University of Tennessee, Knoxville. Emails: shervink@seas.upenn.edu, dongshed@utk.edu

MOTIVATION AND BACKGROUND

- To improve diffusion sample quality and meet user requirements, two common methods are widely used:
  1. **Alignment:** Fine-tuning model to align it with rewards.
  2. **Composition:** Combining several pre-trained models, each emphasizing a desirable attribute of samples.
- Trade-off arises when optimizing for multiple rewards or combining multiple models, as they often represent competing properties.

We propose a **constrained optimization framework unifying alignment and composition** of diffusion models by imposing constraints on rewards and/or proximity to pre-trained models.

CONSTRAINED OPTIMIZATION FRAMEWORK

- **Product Composition (AND)**
  - Constraining the **reverse KL divergence** to multiple pre-trained models to be smaller than a common threshold.

$$\min_{p,u} \; u \quad \text{s.t. } D_{\mathrm{KL}}(p \,\|\, q_i) \le u \text{ for } i = 1, \ldots, m.$$

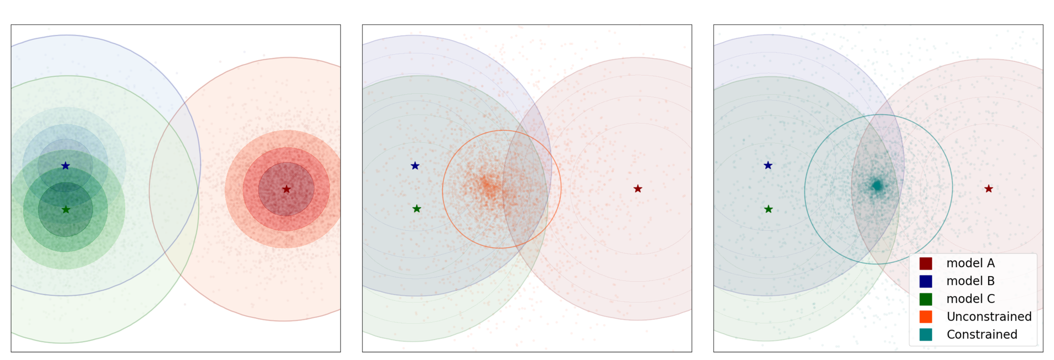

- **Reward Alignment**
  - Constraining **expected rewards** to be above user-specified (or expert-chosen) thresholds while minimizing deviation from pre-trained model.

$$\min_{p} \; D_{\mathrm{KL}}(p \,\|\, q) \quad \text{s.t. } \mathbb{E}_{x\sim p}[r_i(x)] \ge b_i \text{ for } i = 1, \ldots, m.$$

- **Mixture Composition (OR)**
  - Constraining the **forward KL divergence** to multiple data distributions to be smaller than a common threshold.
  - We study mixture composition in our NeurIPS2024 paper [Khalafi S, Ding D, Ribeiro A. Constrained diffusion models via dual training. NeurIPS, 2024].

$$\min_{p,u} \; u \quad \text{s.t. } D_{\mathrm{KL}}(q_i \,\|\, p) \le u \text{ for } i = 1, \ldots, m.$$

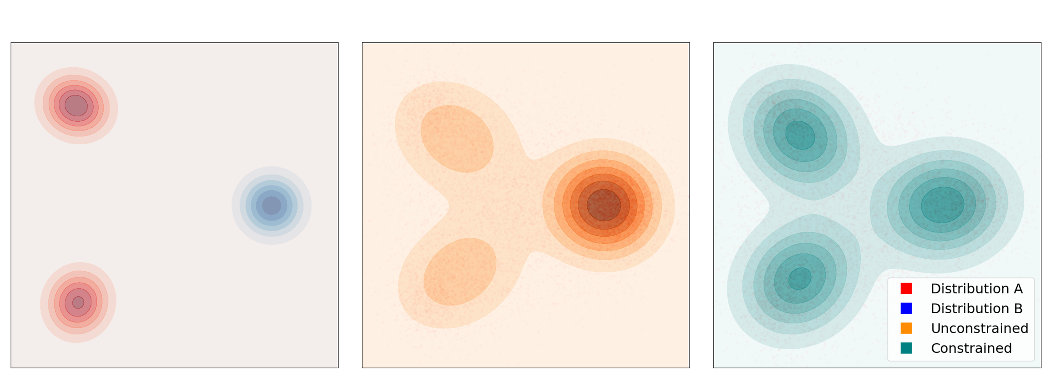

REWARD ALIGNMENT OF DIFFUSION MODELS

- Theorem 1 allows us to characterize the solution to the alignment problem in distribution space:

Theorem 1. Reward Alignment
The constrained alignment problem can be written as:

$$\min_{p\in\mathcal{P}} D_{\mathrm{KL}}\left(p \,\|\, q_{\mathrm{rw}}^{(\lambda^\star)}\right)$$

where $q_{\mathrm{rw}}^{(\lambda)} \propto q \cdot e^{\lambda^\top r}$.

  - The optimal dual variable $\lambda^\star$ ensures we can achieve the smallest possible objective while satisfying the constraints.
- From generic distributions to score-based diffusion models:

$$\min_{s_p} D_{\mathrm{KL}}\big(p_{0:T}(\cdot; s_p) \,\|\, q_{0:T}(\cdot; s_q)\big) \quad \text{s.t. } \mathbb{E}_{x_0\sim p_0(\cdot;s_p)}[r_i(x_0)] \ge b_i \quad \text{for } i = 1, \ldots, m.$$

  - The **path-wise** KL between the joint distributions $p_{0:T}$, $q_{0:T}$ over the whole backward process is straightforward:

$$D_{\mathrm{KL}}(p_{0:T}(\cdot; s_p) \,\|\, q_{0:T}(\cdot; s_q)) = \sum_{t=1}^{T} \mathbb{E}_{x_t\sim p_t(\cdot;s_p)}\left[\frac{1}{2\sigma_t^2}\|s_p(x_t,t) - s_q(x_t,t)\|^2\right]$$

PRODUCT COMPOSITION OF DIFFUSION MODELS

- Theorem 2 characterizes the solution of product composition:

Theorem 2. Product Composition
The Product Composition problem can be written as:

$$\min_{p\in\mathcal{P}} D_{\mathrm{KL}}\left(p \,\|\, q_{\mathrm{AND}}^{(\lambda^\star)}\right)$$

where $q_{\mathrm{AND}}^{(\lambda)} \propto \prod_{i=1}^{m} (q_i)^{\frac{\lambda_i}{\lambda^\top 1}}$.

  - The composed model's support is the **intersection of the individual supports**, with $\lambda^\star$ preventing excessive deviation from any one distribution.
- From generic distributions to score-based diffusion models:

$$\min_{u,\, s_p} u \quad \text{s.t. } D_{\mathrm{KL}}(p_0(x_0; s_p) \,\|\, q_0(x_0; s_q^i)) \le u, \; i = 1, \ldots, m.$$

  - In constrained composition, we care about the **point-wise** KL between marginal distributions at time $t = 0$: p_0 and q_0.

MIXTURE COMPOSITION OF DIFFUSION MODELS

- Theorem 3 allows us to characterize the solution to the mixture composition in distribution space:

Theorem 3. Mixture Composition
The Mixture Composition problem can be written as:

$$\min_{p\in\mathcal{P}} D_{\mathrm{KL}}\left(q_{\mathrm{mix}}^{(\lambda^\star)} \,\|\, p\right)$$

where $q_{\mathrm{mix}}^{(\lambda)} = \sum_{i=1}^{m} \frac{\lambda_i}{\lambda^\top 1} q^i(\cdot)$.

  - We can characterize the optimal dual variable $\lambda^\star$:

$$\lambda^\star = \underset{\lambda\ge 0}{\operatorname{argmax}} \; H(q_{\mathrm{mix}}^{(\lambda^\star)})$$

where maximizing entropy yields the most diverse mixture.

PATH-WISE VS POINT-WISE KL DIVERGENCE

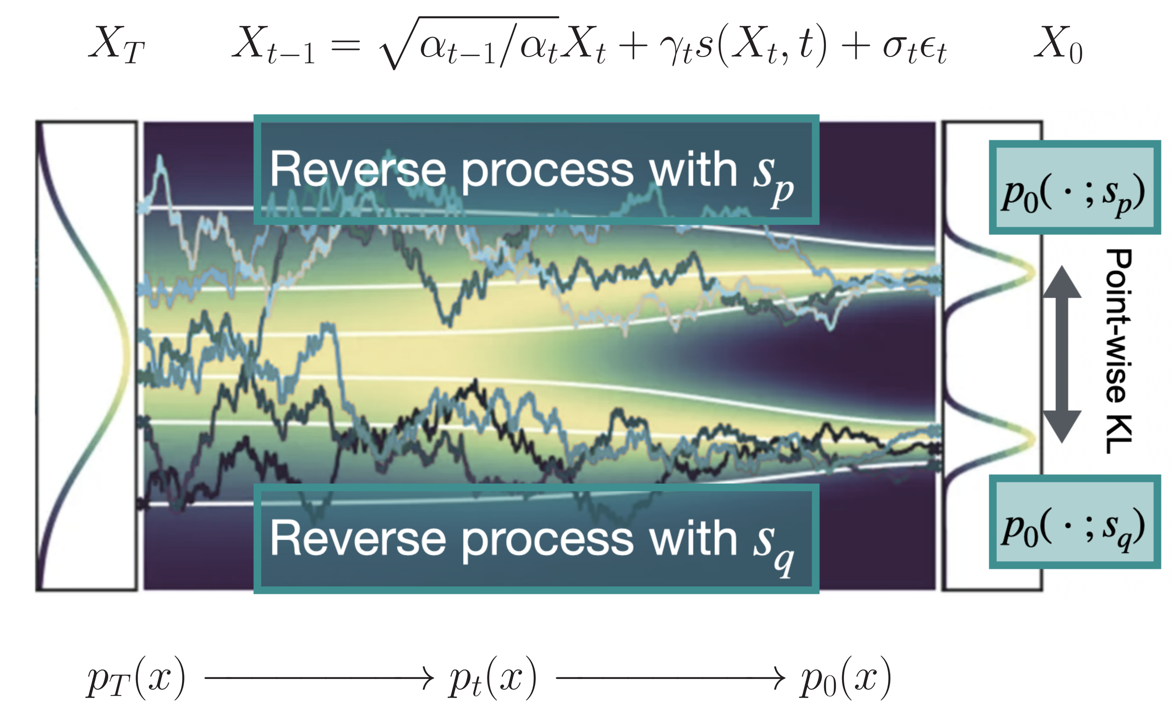

- We compute the point-wise KL between sample distributions of two diffusion models via the following lemma:

Lemma 1. Point-wise KL Computation
Consider two backward diffusion processes with score functions $s_p(x,t)$ and $s_q(x,t)$. Then:

$$D_{\mathrm{KL}}\big(p_0(\cdot; s_p) \,\|\, p_0(\cdot; s_q)\big) = \sum_{t=0}^{T} \tilde{\omega}_t \, \mathbb{E}_{x_t\sim p_t(\cdot;s_p)}\left[\left\|s_p(x_t,t) - s_q(x_t,t)\right\|_2^2\right] + \epsilon_T$$

where $\tilde{\omega}_t$ is a time-dependent constant and $\epsilon_T = \mathcal{O}(1/T)$ is a discretization error.

PRIMAL DUAL ALGORITHM

- All three problems can be shown to exhibit **strong duality**: the optimal values of the primal and dual problems are equal, which motivates our primal-dual algorithm.
- For alignment, define the **path-wise** KL-based Lagrangian as:

$$L_{\mathrm{ALI}}(p,\lambda) = D_{\mathrm{KL}}\big(p_{0:T}(\cdot) \,\|\, q_{0:T}(\cdot)\big) - \lambda^\top(\mathbb{E}_{x\sim p}[r(x)] - b).$$

1. **Primal Step:** Minimize the Lagrangian for fixed dual variable $\lambda = \lambda^{(n)}$.
$$p^{(n+1)} \in \underset{p\in P}{\operatorname{argmin}} \; L_{\mathrm{ALI}}(p, \lambda^{(n)}).$$
2. **Dual Step:** Evaluate the constraints with $p^{(n+1)}$ and update the dual variable:
$$\lambda^{(n+1)} = \left[\lambda^{(n)} + \eta\left(\mathbb{E}_{x\sim p^{(n+1)}}[r(x)] - b\right)\right]_+.$$

- For composition, the algorithm is similar but minimizing the Lagrangian is trickier: computing the **point-wise** KL requires s_p to be a valid score function of a diffusion process.

Lemma 2. Point-wise KL-based Lagrangian
The Lagrangian for Constrained Composition can be written as:

$$L_{\mathrm{AND}}(s,\lambda) = D_{\mathrm{KL}}\left(p_0(x_0; s) \,\|\, q_{\mathrm{AND}}^{(\lambda)}(x_0)\right) - \log Z_{\mathrm{AND}}(\lambda).$$

where $Z_{\mathrm{AND}}(\lambda)$ is a normalization factor dependent on λ.

  - With Lemma 2, we can minimize the Lagrangian using standard score matching with samples from $q_{\mathrm{AND}}^{(\lambda)}$ obtained using sampling methods like Markov Chain Monte Carlo.

NUMERICAL EXPERIMENTS

- **Alignment with Multiple Rewards:** We align pre-trained stable diffusion with multiple rewards via constraints.
  - Setting #1: Aligning with multiple aesthetic rewards.
  - Constraints **minimize deviation from pre-trained** model while **satisfying reward requirements**.

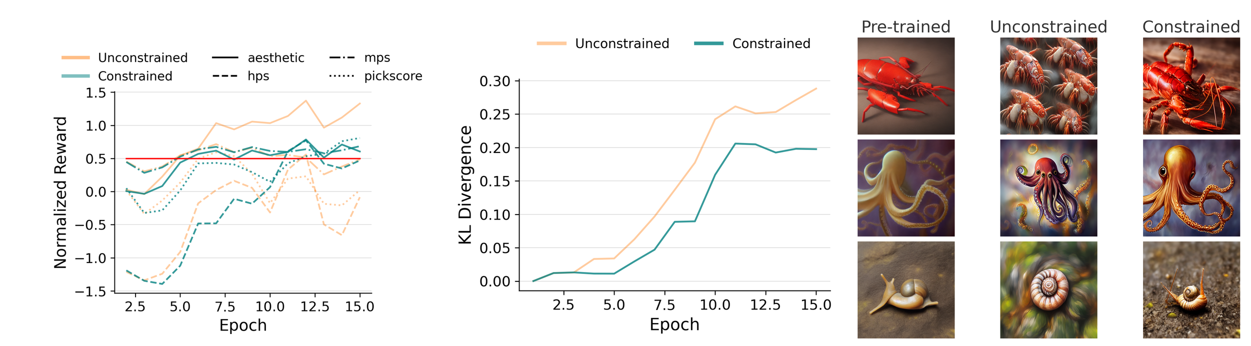

  - Setting #2: We constrain saturation and contrast when aligning with a type of aesthetic reward (MPS) to images with excessively high saturation and contrast.
  - Constraints enforce reward requirements, **preventing overfitting and imbalance** among multiple rewards.

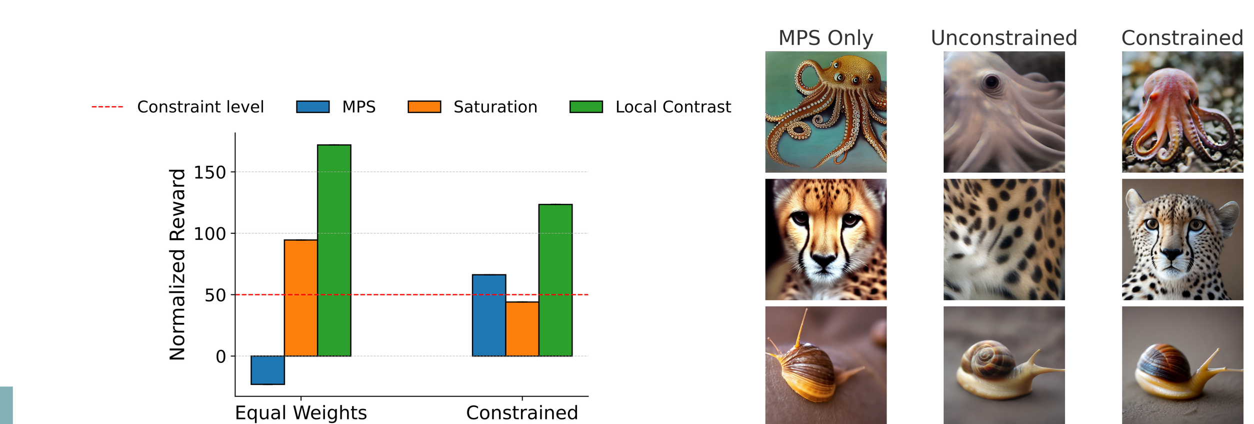

- **Composition of Multiple Pre-trained Models:** We compose multiple models each fine-tuned using a different reward function.
  - Constraints **balance KL divergence** from each model, leading to **more uniform expected rewards** for the composed model.

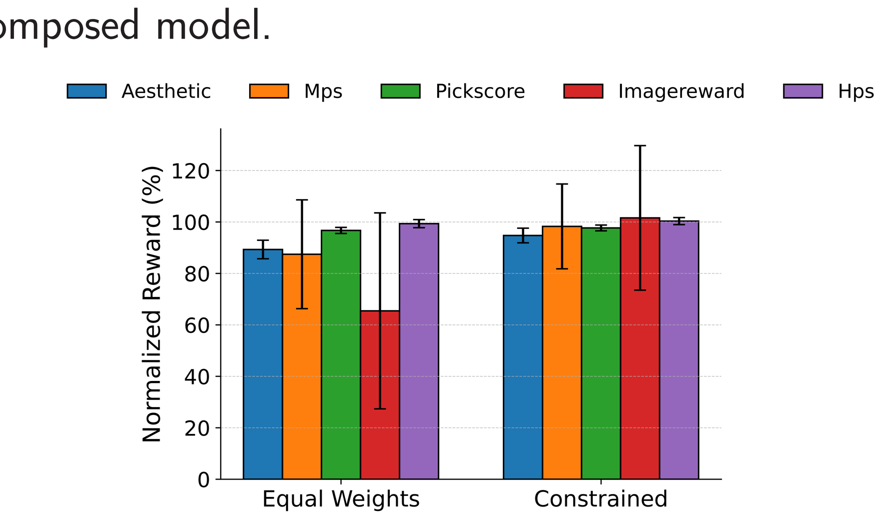

- **Concept Composition:** We compose stable diffusion conditioned on different prompts.
  - Constraints promote **closeness** to each model, thereby reducing the chances of some concepts/models being ignored in favor of others.

	Min. CLIP (↑)	Min. BLIP (↑)
Joint Prompt	21.52	0.206
Eq. Weights	22.18	0.203
Constrained	**22.45**	**0.221**

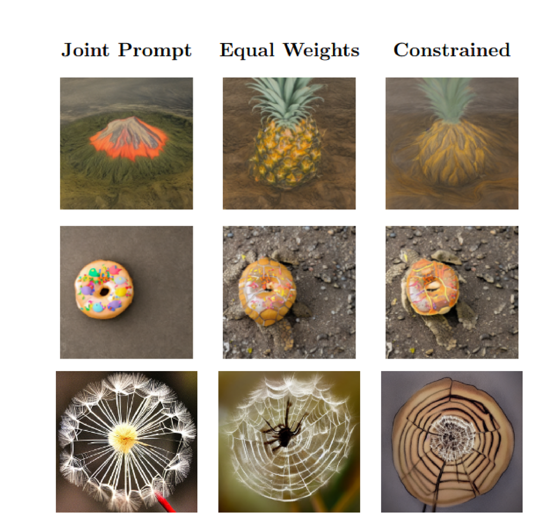

Above: Average minimum CLIP and BLIP score among the composed concepts.
Right: Examples of concept composition, (Row 1: 'volcano', 'pineapple'), (Row 2: 'donut', 'turtle'), (Row 3: 'dandelion', 'spider-web', 'cinnamon roll').