Reweighted Flow Matching via Unbalanced OT for Label-free Long-tailed Generation

HYUNSOO SONG, MINJUNG GIM, JAEWOONG CHOI*

국가수리과학연구소
National Institute for Mathematical Sciences

ICIM Innovation Center for Industrial Mathematics

SUNG KYUN KWAN UNIVERSITY

Majority bias of Flow Matching

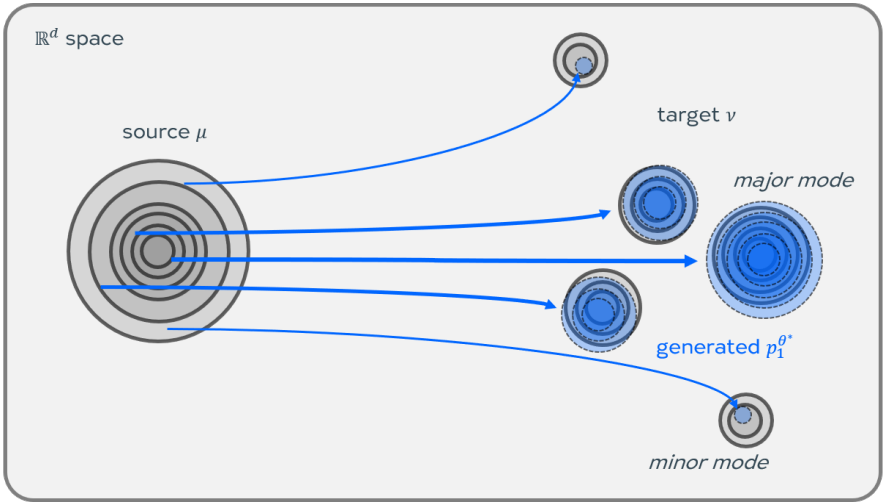

Figure 1: A conceptual illustration, which represents majority bias of flow matching

Overview of Our Approach

$$\hat{S}_\tau\left(\mathbf{x}_1^j\right) \coloneqq \frac{\hat{\pi}^u_{\tau,1}}{\hat{\nu}}\left(\mathbf{x}_1^j\right) = |B|\hat{\pi}^u_{\tau,1}\left(\mathbf{x}_1^j\right)$$

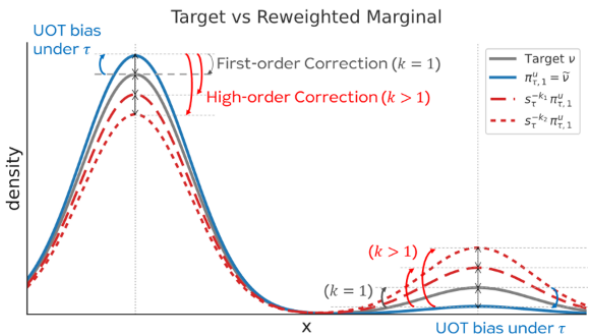

Figure 2: An illustration of our approach

Key Ideas:

- **Unbalanced Optimal Transport (UOT):** Construct the conditional vector field using mini-batch UOT, which is robust to class imbalance.
- **Label-Free Majority Score (s_τ):** We introduce a novel score, defined as the density ratio between the target distribution and the UOT marginal.
- **Inverse Reweighting:** Mitigate majority bias by reweighting the training objective with the inverse of the majority score.

Conditional Flow Matching (CFM)

Here, $\mathbf{z} = (\mathbf{x}_0, \mathbf{x}_1)$ is a pair of samples from the source and target distributions, following a joint distribution (coupling) $\pi(\mathbf{x}_0, \mathbf{x}_1)$.

The (general) Conditional Flow Matching Loss, $\mathcal{L}_{CFM}$, is defined as:

$$\mathcal{L}_{CFM}(\theta) = \mathbb{E}_{t\sim\mathcal{U}[0,1],\mathbf{z}\sim\pi(\mathbf{z}),\mathbf{x}_t\sim p_t(\mathbf{x}_t|\mathbf{z})}||\mathbf{v}_\theta(t,\mathbf{x}_t) - \mathbf{u}_{t|\mathbf{z}}(\mathbf{x}_t|\mathbf{z})||_2^2$$

A simple, common choice for the conditional probability path and vector field is:

$$p_t(\mathbf{x}_t|\mathbf{z}) = \mathcal{N}(\mathbf{x}_t|t\mathbf{x}_1 + (1-t)\mathbf{x}_0, \sigma^2 I)$$
$$\mathbf{u}_t(\mathbf{x}_t|\mathbf{z}) = \mathbf{x}_1 - \mathbf{x}_0$$

Optimal Transport (OT) and OT-CFM

To improve upon simple independent couplings, recent works use mini-batch **Optimal Transport (OT)** to pair source ($\mathbf{x}_0$) and target ($\mathbf{x}_1$) samples.

OT finds the most cost-efficient plan π to transport mass from a source distribution μ to a target ν. The Kantorovich formulation is:

$$C_{OT}(\mu, \nu) := \inf_{\pi \in \Pi(\mu,\nu)} \mathbb{E}_{(\mathbf{x}_0,\mathbf{x}_1)\sim\pi}[c(\mathbf{x}_0, \mathbf{x}_1)]$$

where $c(\mathbf{x}, \mathbf{y})$ is the transport cost, e.g., $||\mathbf{x}_0 - \mathbf{x}_1||^2$.

The OT-CFM loss is defined as:

$$\mathcal{L}_{OT-CFM}(\theta) = \mathbb{E}_{t\sim\mathcal{U}[0,1],\mathbf{z}\sim\pi^o(\mathbf{z}),\mathbf{x}_t\sim p_t(\mathbf{x}_t|\mathbf{z})}||\mathbf{v}_\theta(t, \mathbf{x}_t) - \mathbf{u}_{t|\mathbf{z}}(\mathbf{x}_t|\mathbf{z})||_2^2,$$

where $\mathbf{z} = (\mathbf{x}_0, \mathbf{x}_1)$ and $\pi^o = \arg\inf_{\pi\in\Pi(\mu,\nu)} \mathbb{E}_{(\mathbf{x}_0,\mathbf{x}_1)\sim\pi}[c(\mathbf{x}_0, \mathbf{x}_1)]$.

A key property of OT is the **exact marginal constraint**:

- $\Pi(\mu, \nu)$ is the set of all joint distributions $\pi(\mathbf{x}_0, \mathbf{x}_1)$ whose marginals are precisely μ and ν.
- This strict constraint makes OT sensitive to outliers and class imbalance, which is a problem for long-tailed data.

Unbalanced Optimal Transport (UOT)

Our method is built upon **Unbalanced Optimal Transport (UOT)**, which is more suitable for imbalanced distributions.

UOT relaxes the exact marginal constraints by adding penalty terms that allow the marginals of the transport plan (π_0, π_1) to differ from the source and target distributions (μ, ν).

$$C_{UOT}(\mu, \nu) = \inf_{\pi} \left[\int c(\mathbf{x}_0, \mathbf{x}_1) d\pi + \tau_1 D_1(\pi_0 || \mu) + \tau_2 D_2(\pi_1 || \nu) \right]$$

- D_1 and D_2 are f-divergences that penalize deviations from the marginals.
- $\tau_1, \tau_2 > 0$ control how strictly the marginals must be matched. (It will be equal to OT when $\tau_1 = \infty, \tau_2 = \infty$.)
- This relaxation provides robustness to outliers and allows for more flexible transport between distributions of unequal mass, as seen in long-tailed settings.

UOT-CFM

The UOT-CFM loss is defined as:

$$\mathcal{L}_{UOT-CFM}(\theta) = \mathbb{E}_{t\sim\mathcal{U}[0,1],\mathbf{z}\sim\pi^u_\tau(\mathbf{z}),\mathbf{x}_t\sim p_t(\mathbf{x}_t|\mathbf{z})}||\mathbf{v}_\theta(t,\mathbf{x}_t) - \mathbf{u}_{t|\mathbf{z}}(\mathbf{x}_t|\mathbf{z})||_2^2.$$

- Training with this loss (UOT-CFM) yields a **biased** target marginal p_1, which is not the true target ν.
- The resulting marginal p_1 converges to the UOT coupling's target marginal, $\pi^u_{\tau,1}$.

Training a conditional flow matching model with UOT transport plan π^u_τ yields the target marginal to be:

$$p_1 = \pi^u_{\tau,1}.$$

Proposed Method (1/2): The Majority Score

The key insight is that the optimal UOT coupling π^u naturally concentrates mass on high-density (majority) modes while down-weighting low-density (minority or outlier) modes.

This allows us to define a **label-free majority score**, $s_\tau(\mathbf{x}_1)$, as the density ratio (Radon-Nikodym derivative) between the UOT target marginal $\pi^u_{\tau,1}$ and the true target distribution ν.

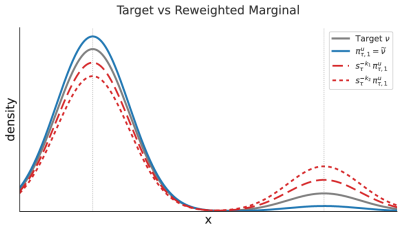

$$s_\tau(\mathbf{x}_1) := \frac{d\pi^u_{\tau,1}}{d\nu}(\mathbf{x}_1) > 0$$

- If $s_\tau(\mathbf{x}_1) > 1$, the sample $\mathbf{x}_1$ belongs to a majority region emphasized by UOT.
- If $s_\tau(\mathbf{x}_1) \ll 1$, the sample $\mathbf{x}_1$ belongs to a minority region down-weighted by UOT.

Notations

- $\mathbf{x}_1$: A **sample** from the target data distribution ν.
- ν: The **true target distribution** that the generative model aims to learn.
- $\pi^u_{\tau,1}$: The **target marginal distribution** of the optimal Unbalanced Optimal Transport (UOT) coupling.
- $\frac{d\pi^u_{\tau,1}}{d\nu}$: The **Radon-Nikodym derivative**, which defines the density ratio of the UOT marginal $\pi^u_{\tau,1}$ with respect to the target distribution ν.
- τ: A hyperparameter that controls the **strength of the marginal matching penalty** in the UOT problem.

Proposed Method (2/2): The UOT-RFM Objective

To counteract the majority bias, we introduce an inverse reweighting factor into the conditional flow matching loss, based on our majority score.

The proposed **UOT-RFM objective** with a correction order $k \geq 1$ is:

$$\mathcal{L}_{\text{ours},k}(\boldsymbol{\theta}) = \mathbb{E}_{t\sim\mathcal{U}[0,1],\mathbf{z}\sim\pi^u_\tau(\mathbf{z}),\mathbf{x}_t\sim p_t(\mathbf{x}_t|\mathbf{z})} \left[s_\tau(\mathbf{x}_1)^{-k} \|\mathbf{v}_\theta(t,\mathbf{x}_t) - \mathbf{u}_{t|\mathbf{z}}(\mathbf{x}_t|\mathbf{z})\|_2^2 \right], \quad (1)$$

where the conditional variable $\mathbf{z} = (\mathbf{x}_0, \mathbf{x}_1)$.

Theoretical Flow:

- Training with the UOT coupling alone (equivalent to $k = 0$) learns a biased distribution $p_1 = \pi^u_{\tau,1} \neq \nu$.
- With first-order correction $(k = 1)$, our method provably recovers the true target distribution ν.
- With higher-order correction $(k > 1)$, the model learns to generate from a distribution $p_1 \propto s_\tau^{-(k-1)}\nu$, further up-weighting minority samples.

Theorem: First-order correction ($k = 1$)

Theorem

Let π_τ^u be the optimal source-fixed UOT coupling between μ and ν with $\tau_2 = \tau > 0$ and assume that its target marginal satisfies $\nu \ll \pi_{\tau,1}^u$, i.e., ν is absolutely continuous w.r.t. $\pi_{\tau,1}^u$. Training a flow matching model with π_τ^u yields the biased distribution $p_1 = \pi_{\tau,1}^u \neq \nu$. However, applying the first-order correction (our method with $k = 1$) recovers the true target distribution ν.

$$\mathcal{L}_{\text{ours,k=1}}(\boldsymbol{\theta}) = \mathbb{E}_{t\sim\mathcal{U}[0,1],\mathbf{z}\sim\pi_\tau^u(\mathbf{z}),\mathbf{x}_t\sim p_t(\mathbf{x}_t|\mathbf{z})} \left[s_\tau(\mathbf{x}_1)^{-1} \|\mathbf{v}_\theta(t,\mathbf{x}_t) - \mathbf{u}_{t|\mathbf{z}}(\mathbf{x}_t|\mathbf{z})\|_2^2 \right]. \quad (2)$$

More generally, UOT-RFM with correction order k generates $p_1 \propto s_\tau^{-k}\pi_{\tau,1}^u = s_\tau^{-(k-1)}\nu$.

Proposed Method (3/3): Training Algorithm

Algorithm 1: Training mini-batch UOT-RFM

Input: Empirical or sampleable distributions μ, ν, bandwidth σ, batch size b, initial network $\mathbf{v}_\theta$, sinkhorn target marginal weight τ, weight power scale k.

Initialize: $\tau_1 \leftarrow \infty$ `// Source-fixed UOT`

while *Training do* **do**

  Sample batches of size b *i.i.d.* from the datasets:

  $\mathbf{x}_0 \sim \mu$; $\mathbf{x}_1 \sim \nu$

  $\pi_r^\tau \leftarrow \mathsf{UOT}(\mathbf{x}_1, \mathbf{x}_0, \tau)$ `// Source-fixed UOT` $\tau_1 = \infty, \tau_2 = \tau$

  $(\mathbf{x}_0, \mathbf{x}_1) \sim \pi_r^\tau$; $t \sim \mathcal{U}(0,1)$

  $\mu_t \leftarrow t\mathbf{x}_1 + (1-t)\mathbf{x}_0$

  $\mathbf{x} \sim \mathcal{N}(\mu_t, \sigma^2 I)$

  $\hat{s}_\tau(\mathbf{x}_1) := \frac{d\pi^u_{\tau,1}}{d\nu}(\mathbf{x}_1)$

  $\mathcal{L}_{\mathsf{ours}}(\theta) \leftarrow \hat{s}_\tau(\mathbf{x}_1)^{-k} \|\mathbf{v}_\theta(t, \mathbf{x}) - (\mathbf{x}_1 - \mathbf{x}_0)\|^2$

  $\theta \leftarrow \mathsf{Update}(\theta, \nabla_\theta \mathcal{L}_{\mathsf{ours}}(\theta))$

end

return $\mathbf{v}_\theta$

Experimental Results - Main Quantitative Evaluations

Table 1: Evaluation of marginal distribution matching under the LT→LT setting for the long-tailed benchmarks. We report FID scores (↓) on CIFAR-10-LT and CIFAR-100-LT with two imbalance ratios: $\mathcal{I} = 0.01$ and $\mathcal{I} = 0.001$.

Model	CIFAR-10-LT		CIFAR-100-LT	
	$\mathcal{I} = 0.01$	$\mathcal{I} = 0.001$	$\mathcal{I} = 0.01$	$\mathcal{I} = 0.001$
I-CFM	14.57	17.54	25.55	31.86
OT-CFM	17.31	21.26	31.34	38.37
UOT-CFM	14.25	18.13	25.33	31.83
ours	**11.17**	**13.05**	**15.29**	**18.25**

Table 2: Evaluation of marginal distribution matching under the LT→Balanced setting for the long-tailed benchmarks. We report FID scores (↓) with imbalance ratio $\mathcal{I} = 0.01$.

Model	CIFAR-10-LT	CIFAR-100-LT
I-CFM	25.46	24.39
OT-CFM	27.51	29.19
UOT-CFM	24.94	24.05
ours	**24.06**	**16.83**

Experimental Results - Main Quantitative Evaluations

Table 3: Quantitative evaluation of mode coverage under the LT→LT setting with $\mathcal{I} = 0.01$.

Model	CIFAR-10-LT			CIFAR-100-LT		
	P. (↑)	R. (↑)	F1 (↑)	P. (↑)	R. (↑)	F1 (↑)
I-CFM	0.71	0.29	0.41	0.62	0.29	0.40
OT-CFM	**0.73**	0.24	0.36	**0.73**	0.24	0.36
UOT-CFM	**0.73**	0.29	0.42	**0.73**	0.27	0.40
ours	0.62	**0.41**	**0.49**	0.69	**0.32**	**0.44**

Table 4: FID scores (↓) on standard balanced benchmarks.

Model	CIFAR-10	CIFAR-100
I-CFM	3.78	6.39
OT-CFM	3.64	**6.14**
UOT-CFM	3.62	6.45
ours	**3.58**	6.54

Experimental Results - Qualitative Evaluation

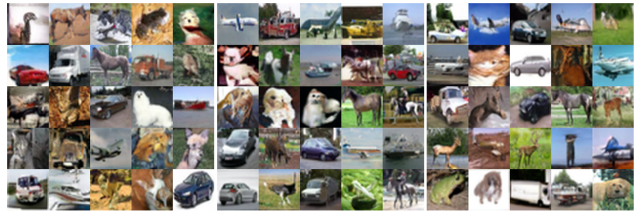

(a) I-CFM (b) OT-CFM (c) UOT-RFM

(d) I-CFM (e) OT-CFM (f) UOT-RFM

Figure 3: CIFAR image generation results: The first row shows images randomly generated from models trained on the balanced CIFAR10 dataset. The second row shows images from models trained on the CIFAR-10-LT dataset.

Experimental Results - Qualitative Evaluation



(a) I-CFM (b) OT-CFM (c) UOT-RFM

(d) I-CFM (e) OT-CFM (f) UOT-RFM

Figure 4: Tail-classes CIFAR image generation results: The classified images from models trained on the CIFAR-10-LT dataset.

Experimental Results - Distribution Estimation

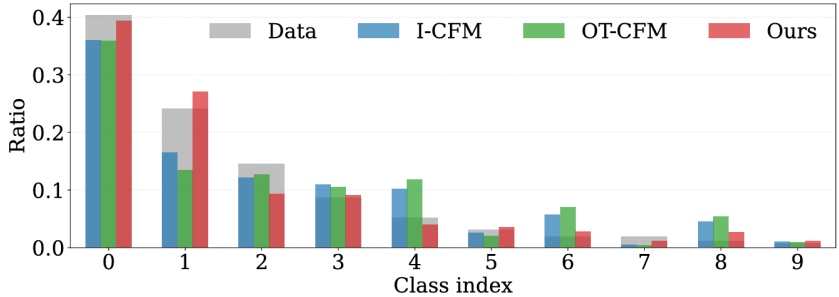

Figure 5: Generated class distribution on CIFAR-10-LT with $\mathcal{I} = 0.01$: The average Normalized Class Ratio Errors (NCREs) are: I-CFM = 0.84, OT-CFM = 1.02, and ours = 0.40.

$$\text{NCRE}_i = \frac{|r_{gen,i} - r_{data,i}|}{r_{data,i}}, \tag{3}$$

Experimental Results - Ablation Study

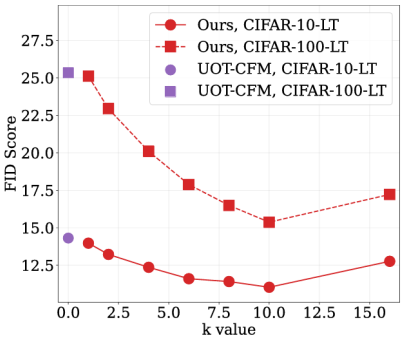

(a) Ablation study on k when $\tau = 2.0$

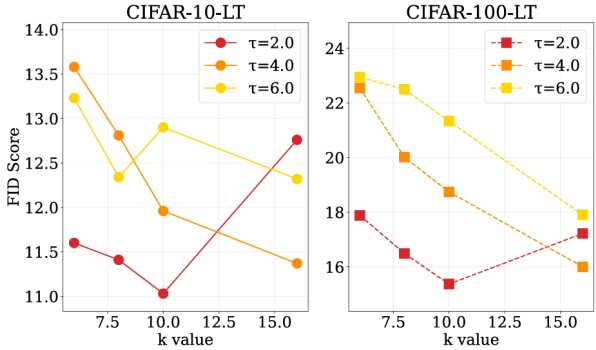

(b) Ablation study on τ across multiple k (UOT-RFM)

Figure 6: Ablation studies on the correction order k and the marginal matching strength τ

Experimental Results - NLL

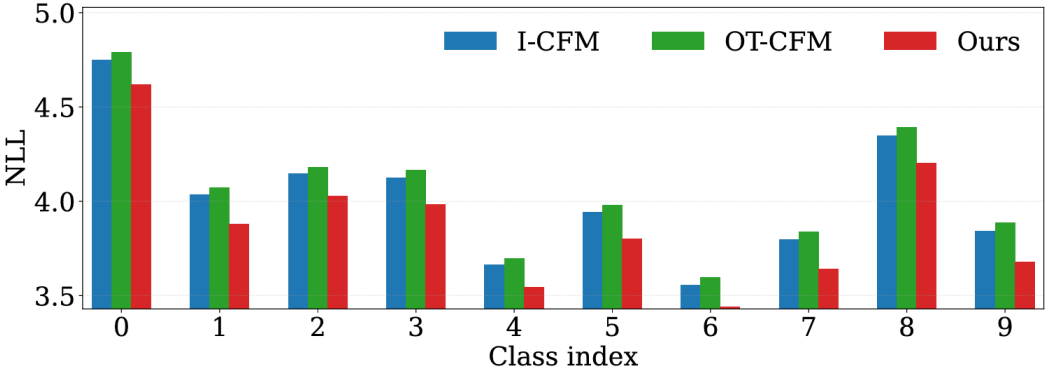

Figure 7: Class-wise negative log-likelihood of CIFAR-10-LT data samples on each model: Our method achieves an overall average NLL of 3.88, while I-CFM and OT-CFM achieve 4.02 and 4.06, respectively.

Experimental Results - ODE Sampler

Table 5: FID comparison between UOT-RFM and OT-CFM across different ODE solvers and NFEs

Method	**UOT-RFM (↓)**	**OT-CFM (↓)**
dopri5 (NFE: 152/140)	**11.16**	17.58
Euler (5 steps)	38.00	39.13
Euler (10 steps)	19.61	24.99
Euler (20 steps)	14.74	20.07
Euler (50 steps)	12.37	17.84
Euler (100 steps)	11.72	17.37
Euler (200 steps)	11.49	**17.18**
Euler (500 steps)	11.28	17.27

Experimental Results - Other datasets, Latent space

Table 6: FID comparison between OT-CLFM and UOT-RLFM(ours, $k = 10.0$) on FFHQ-64 dataset

Model	FID(↓)
OT-CLFM	14.24
UOT-RLFM ($\tau = 6.0$)	13.82
UOT-RLFM ($\tau = 2.0$)	**13.53**

Table 7: FID comparison between OT-CLFM and UOT-RLFM(ours, $k = 10.0, \tau = 2.0$) on Place-LT(256x256) dataset

Model	FID(↓)
OT-CLFM	18.69
UOT-RLFM	**18.13**

Conclusion

Contributions:

- **Label-free majority score** from the UOT coupling, which captures minority features beyond class labels
- **Provable correction**: $k = 1$ recovers the target ν; $k > 1$ further up-weights minority samples
- **Training-time** method: no additional inference cost

Potential applications:

- Generative modeling on highly imbalanced data without labels
- Data augmentation for rare samples (e.g., anomaly detection)