

Learning Multimodal Latent Space with EBM Prior and MCMC Inference

Shiyu Yuan¹, Carlo Lipizzi¹, Tian Han²

School of Systems and Enterprises¹, Department of Computer Science²
Stevens Institute of Technology

{syuan14, clipizzi, than6}@stevens.edu

Abstract

Multimodal generative models are crucial for various applications. We propose an approach that combines an expressive energy-based model (EBM) prior with Markov Chain Monte Carlo (MCMC) inference in the latent space for multimodal generation. The EBM prior acts as an informative guide, while MCMC inference, specifically through short-run Langevin dynamics, brings the posterior distribution closer to its true form. This method not only provides an expressive prior to better capture the complexity of multimodality but also improves the learning of shared latent variables for more coherent generation across modalities. Our proposed method is supported by empirical experiments, underscoring the effectiveness of our EBM prior with MCMC inference in enhancing cross-modal and joint generative tasks in multimodal contexts.

1. Introduction

Multimodal generative models are important because of their ability to interpret, integrate, and synthesize information from diverse inputs. In these models, shared latent variables play a crucial role in integrating features from diverse modalities into a unified and informative representation for downstream generative tasks. Recent works have explored multimodal generation through denoising-based networks [Bao et al. \(2023\)](#); [Ho et al. \(2020\)](#); [Ramesh et al. \(2022\)](#) or by learning representations via modality alignment, as seen in [Radford et al. \(2021\)](#). However, the former approaches often lack a shared representation of different modalities, while the latter may not support generation tasks

Variational Autoencoder (VAE)-based models [Kingma and Welling \(2013\)](#) can achieve both objectives: learning a shared representation through a latent aggregation mechanism [Shi et al. \(2019\)](#); [Wu and Goodman \(2018\)](#) and generating data using top-down generators. However, this approach inherits the limitations of traditional VAE models, notably the reliance on uni-modal priors that are not informative enough to capture the complexity of multimodality.

To tackle the problem of non-informative prior, we pro-

pose a joint training scheme for multimodal generation that employs an EBM prior with MCMC inference. This approach leverages an expressive prior to better capture multimodal data complexity. Additionally, the use of MCMC inference with Langevin dynamics improves the learning process of EBM. In summary, our contributions are as follows:

1. We propose the use of an EBM prior to replace the uni-modal prior in multimodal generation, enhancing the capture of multimodal data complexity.
2. We employ MCMC inference to more accurately approximate the true posterior, as compared with variational inference, which improve EBM learning.
3. We conduct empirical experiments on multimodal datasets to validate our proposed EBM prior with MCMC inference, demonstrating improvements in the multimodal generative model both visually and numerically.

2. Related Work

2.1. Multimodal Generative Models

In the learning of multimodal generative models, two fundamental challenges arise: one is obtaining a shared representation that captures the common knowledge among modalities, and the other is cross-modal generation, which involves translating between modalities [Suzuki and Matsuo \(2022\)](#). VAE-based multimodal generative models [Hwang et al. \(2021\)](#); [Palumbo et al. \(2023\)](#); [Shi et al. \(2019\)](#); [Sutter et al. \(2020, 2021\)](#); [Wu and Goodman \(2018\)](#) have achieved good performance in learning such shared information and performing cross-modal generation, but they still face the non-informative prior limitation.

2.2. Expressive Prior

Due to the complexity of data distributions, recent works seek to utilize expressive priors to represent prior knowledge in generative models, such as hierarchical priors [Cui et al. \(2023\)](#); [Vahdat and Kautz \(2020\)](#), flow-based priors [Xie et al. \(2023\)](#), and energy-based priors [Pang et al. \(2020\)](#). However, such expressive priors are rarely discussed in the

context of multimodal generation.

2.3. MCMC-based Inference

MCMC inference enables sampling from distributions that are otherwise challenging to track directly. Several works present promising performance on generative tasks through MCMC inference, such as dual-MCMC teaching, alternating back-propagation and short-run MCMC as seen in Cui and Han (2023); Han et al. (2017); Nijkamp et al. (2019). However, these methods are rarely used in multimodal generative modeling.

3. Methodology

3.1. Preliminaries

Multimodal Generative Model Multimodal generative models aim to learn the joint distribution of multimodal data. Suppose there are M modalities; data in each modality is denoted as x^m , the entire dataset is denoted as $X = (x^1, x^2 \dots x^m)$, and the shared latent variable is denoted as z . The joint probability $p(z, X)$ can be factorized into Eqn. 1.

$$p(z, X) = p(z) \prod p(x^1|z)p(x^2|z) \dots p(x^m|z) \quad (1)$$

Most multimodal models learn $p(z, X)$ through a shared latent variable. In multimodal generative models, VAE-based models can learn such shared latents through two foundation aggregation approach: POE Wu and Goodman (2018) and MOE Shi et al. (2019), with MOE being the more commonly adopted one Hwang et al. (2021); Palumbo et al. (2023); Sutter et al. (2020, 2021). MOE averages latent variables from each modality, given by $q_{\Phi}(z|X) = \frac{1}{M} \sum_{m=1}^M q_{\phi_m}(z|x^m)$. Learning such models typically adopting ELBO as in traditional VAE models, as shown in Eqn.2.

$$\begin{aligned} L_{\text{MOE}}(\theta) &= \frac{1}{M} \sum_{m=1}^M \left[E_{q_{\phi_m}(z|x^m)} \log \frac{p_{\beta}(z, X)}{q_{\Phi}(z|X)} \right] \\ &= \frac{1}{M} \sum_{m=1}^M \left[E_{q_{\phi_m}(z|x^m)} \log p_{\beta_m}(x^m|z) \right. \\ &\quad \left. + \sum_{\substack{n=1 \\ n \neq m}}^M E_{q_{\phi_m}(z|x^m)} \log p_{\beta_n}(x^n|z) \right. \\ &\quad \left. - \text{KL}[q_{\Phi}(z|X) \| p(z)] \right] \\ &= E_{q_{\Phi}(z|X)} \log p_{\beta}(X|z) - \text{KL}[q_{\Phi}(z|X) \| p(z)] \end{aligned} \quad (2)$$

Where θ comprises (β, ϕ) , β and ϕ denote the generator and inference model parameters, respectively. One limitation is that the objective includes a uni-modal prior $p(z)$ which cannot sufficiently capture the complexity of multimodal data space. In this work, we propose using an expressive EBM prior $p_{\alpha}(z)$ to replace the non-informative prior. **EBM on Latent Space** Latent space EBM aims to learn a latent distribution with high probability by assigning low energy, as shown in Eqn.3

$$p_{\alpha}(z) = \frac{1}{Z(\alpha)} \exp[f_{\alpha}(z)] \cdot p(z) \quad (3)$$

Where $-f_{\alpha}(z)$ is the energy function, $p(z)$ is uni-modal distribution as EBM prior initialization, $Z(\alpha) = \int_z \exp[f_{\alpha}(z)] \cdot p(z) dz$ is the normalization term normally intractable. Learning an EBM prior in the latent space has shown promising performance on generative tasks Cui et al. (2023); Han et al. (2020); Pang et al. (2020); Zhang et al. (2021), but the application of EBM prior in multimodal generative models is under-explored. Moreover, EBM's unnormalized exponential distribution as in Eqn. 3 provides high flexibility in modeling the latent space and enhances its expressiveness in representing the complexity of multimodal data.

3.2. Method

Due to the non-informative prior in Eqn.2 and the expressiveness of the EBM prior, we propose a model that follows the MOE aggregation framework and is jointly learned with EBM prior using MLE. The objective of the joint learning model MOE-EBM is as follows in Eqn. 4.

$$\begin{aligned} L_{\text{MOE-EBM}}(\theta) &= E_{q_{\Phi}(z|X)} \log p_{\beta}(X|z) - \text{KL}[q_{\Phi}(z|X) \| p_{\alpha}(z)] \\ &= E_{q_{\Phi}(z|X)} \left[\log p_{\beta}(X|z) - \log \frac{q_{\Phi}(z|X)}{p_{\alpha}(z)} \right] \\ &= E_{q_{\Phi}(z|X)} [\log p_{\beta}(X|z) - \log q_{\Phi}(z|X) + \log p_{\alpha}(z)] \end{aligned} \quad (4)$$

VAE Learning By taking derivative of Eqn. 4, we can obtain gradient with respect to θ as shown in Eqn. 5, where $\theta = (\beta, \phi)$. By replacing $p_{\alpha}(z)$ with the equation in Eqn. 3, we obtain a refined objective that includes the ELBO and an additional energy term $f_{\alpha}(z)$. When training the VAE part, we consider the term $Z(\alpha)$ as constant since it does not involve sampling from the expectation of $q_{\Phi}(z|X)$.

$$\begin{aligned} L'_{\text{MOE-EBM}}(\theta) &= -E_{q_{\Phi}(z|X)} \left[\frac{\partial}{\partial \theta} (\log p_{\beta}(X|z) + \log q_{\Phi}(z|X) - \log p_{\alpha}(z)) \right] \\ &= -E_{q_{\Phi}(z|X)} \left[\underbrace{\frac{\partial}{\partial \theta} (\log p_{\beta}(X|z) + \log \frac{q_{\Phi}(z|X)}{p(z)})}_{\text{ELBO}} \right] \\ &\quad + E_{q_{\Phi}(z|X)} \left[\frac{\partial}{\partial \theta} f_{\alpha}(z) \right] \end{aligned} \quad (5)$$

EBM Learning For the EBM prior model, as shown in Eqn. 3, we have

$$\log p_{\alpha}(z) = f_{\alpha}(z) - \log Z(\alpha) + \log p(z) \quad (6)$$

where the derivative of Eqn. 6 is as follows:

$$\frac{\partial}{\partial \alpha} \log p_{\alpha}(z) = \frac{\partial}{\partial \alpha} f_{\alpha}(z) - E_{p_{\alpha}(z)} \frac{\partial}{\partial \alpha} f_{\alpha}(z) \quad (7)$$

According to Eqn. 4 and Eqn. 7, the learning gradient for α is as follows:

$$\begin{aligned} L'_{\text{MOE-EBM}}(\alpha) &= E_{q_{\Phi}(z|X)} \frac{\partial}{\partial \alpha} \log p_{\alpha}(z) \\ &= E_{q_{\Phi}(z|X)} \frac{\partial}{\partial \alpha} f_{\alpha}(z) - E_{p_{\alpha}(z)} \frac{\partial}{\partial \alpha} f_{\alpha}(z) \end{aligned} \quad (8)$$

MCMC Inference with Langevin Dynamics From Eqn. 8, we notice that learning EBM requires sampling z from two expectations: $E_{q_{\Phi}(z|X)}$ and $E_{p_{\alpha}(z)}$, which can be achieved through MCMC sampling, such as Langevin dynamics (LD), as shown in Eqn. 9.

$$z_{\tau+1} = z_{\tau} - \frac{s^2}{2} \frac{\partial}{\partial z} [\log \pi(z_{\tau})] + s \cdot \epsilon_{\tau} \quad (9)$$

Where s is the step size, ϵ_{τ} represents Gaussian noise ($\epsilon_{\tau} \sim \mathcal{N}(0, I_d)$), and τ is the time step in LD. When sampling from EBM prior, $\pi(z_{\tau}) = p_{\alpha}(z)$ where $p_{\alpha}(z)$ is initialized from a simple reference distribution. In this work, we use Laplacian distribution as initialization. When sampling from $q_{\Phi}(z|X)$, $\pi(z_{\tau}) = q_{\Phi}(z|X)$ where $q_{\Phi}(z|X)$ is initialized from a variational inferred posterior. $M_{\Phi}^z(\cdot)$ denotes the Markov transition kernel of finite step LD that samples z from $q_{\Phi}(z|X)$, indicating the marginal distribution of z obtained by running $M_{\Phi}^z q_{\Phi}(z|X) = \int_{z'} M_{\Phi}^z(z') q_{\Phi}(z'|X) dz'$ initialized from $q_{\Phi}(z|X)$.

The improved version of learning EBM with MCMC sampling on both the prior and posterior has the following refined format:

$$\begin{aligned} L'_{\text{MOE-EBM}}(\alpha) \\ = E_{M_{\Phi}^z q_{\Phi}(z|X)} \frac{\partial}{\partial \alpha} f_{\alpha}(z) - E_{p_{\alpha}(z)} \frac{\partial}{\partial \alpha} f_{\alpha}(z) \end{aligned} \quad (10)$$

Because we initialize $p_{\alpha}(z)$ from a non-informative Laplace(0, I_d), but initialize $M_{\Phi}^z q_{\Phi}(z|X)$ from a relatively informative variational inferred posterior $q_{\Phi}(z|X)$, in LD for both the prior and posterior, we set different time steps τ and step numbers s to better learn EBM.

MOE with Modality Prior To validate our proposed MOE-EBM and compare it with recent MOE-based multimodal generative baselines, we adopt the recent variants of MOE Palumbo et al. (2023), which model both shared and modality-specific priors in separate latent subspaces. The detailed design of such a latent space can be referred to in Palumbo et al. (2023). We test our proposed model with this latent subspace to investigate its effectiveness within the MOE variant framework.

4. Experiment

4.1. Dataset and Experiment Settings

To evaluate our model, we use PolyMNIST Sutter et al. (2021) to numerically and visually assess the effectiveness of the EBM prior with MCMC inference. A detailed description of PolyMNIST can be found in Sutter et al. (2021). Quantitatively, we measure generative coherence Shi et al. (2019) to investigate consistency in generations. Additionally, we assess perceptual performance using FID. We test our model on MOE with a modality-specific prior: MMVAE+ Palumbo et al. (2023). Our results are also compared to other baselines built within the MOE framework,

including MMVAE Shi et al. (2019), mmJSD Sutter et al. (2020), and MoPoE Sutter et al. (2021).

4.2. EBM Prior with MCMC Inference

Figures 1 and 2 show joint and cross generation across different frameworks, with digits highlighted as the shared information among modalities. Both joint and cross-generation demonstrate improvements through visual comparisons with other MOE-based multimodal generative models. Quantitative comparisons in Table 1 further validate the effectiveness of our proposed MOE-EBM.

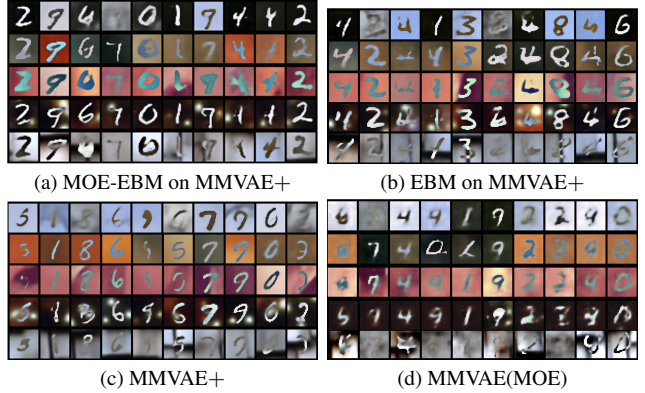


Figure 1. Joint Generation



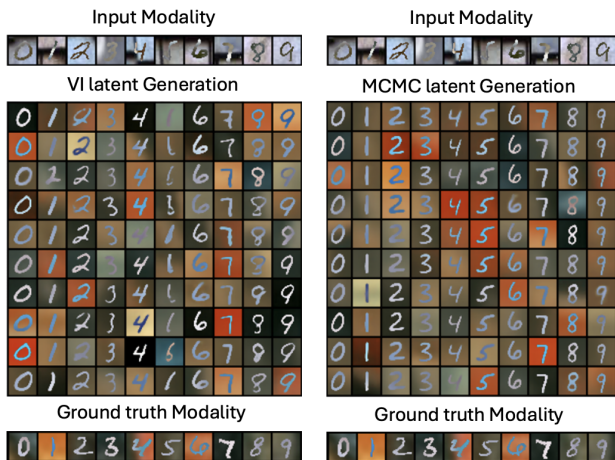
Figure 2. Cross Generation: from right to left are EBM on MMVAE(MOE), MMVAE+, EBM on MMVAE+, MOE-EBM on MMVAE+

4.3. Generation Comparison between Variational Inference and MCMC Inference


MCMC inference learned through EBM can be closely to the true posterior compare with variational inferred posterior. We explore the generation quality before and after LD

Model	Joint Coherence (\uparrow)	Cross Coherence (\uparrow)	Joint FID (\downarrow)	Cross FID (\downarrow)
MMVAE*	0.232	0.844	164.71	150.83
mmJSD*	0.060	0.778	180.55	222.09
MoPoE*	0.141	0.720	107.11	178.27
MMVAE+*	0.344	0.869	96.01	92.81
Ours(MOE-EBM: pre LD)	NA	0.885	NA	94.72
Ours(MOE-EBM: post LD)	0.574	0.943	98.23	90.32

Table 1. Generation Coherence and FID : We present generation results of our proposed MOE-EBM with before LD (variational inference, denoted as pre LD) and after LD (MCMC inference, denoted as post LD), and compare with other MOE-based multimodality generative models (* are results referred from Palumbo et al. (2023).)



(a) Variational Inference vs. MCMC Inference Generation



(b) Generation in Markov transition with LD

Figure 3. (a) Comparative visualization of generation quality before and after LD refinement. (b) Generation improvement during Markov transitions using LD.

in Figure 3a and visualize the changes in generation quality during LD refinement in Figure 3. To quantitatively validate that MCMC inference closely approximates the true posterior, we present the generation coherence before and after LD in Table 1.

5. Ablation Studies

To investigate the effectiveness of an EBM prior with MCMC inference in multimodal generation, we conduct two ablation studies: one incorporating an EBM prior with MOE, and the other incorporating an EBM prior with MOE variants, learning with a modality-specific prior. Notably, neither ablation involved MCMC inference. We present the results for each setting in comparison with our MOE-EBM framework in Tables 2 and 3. We observe that using only the EBM prior, generation coherence showed non-trivial improvements compared to the corresponding baselines. This indicates that the EBM prior can better capture shared information in the complex multimodal data space. Furthermore, MCMC inference directly benefits cross-modal generation, as validated by the ablation results.

Model	Joint Coh(\uparrow)	Cross Coh (\uparrow)
EBM-MMVAE	0.340	0.856
EBM-MMVAE+	0.531	0.877
MOE-EBM	0.574	0.943

Table 2. Ablation Results: Generation Coherence (Coh refer to Coherence due to space limitation)

Model	Joint FID (\downarrow)	Cross FID (\downarrow)
EBM-MMVAE	129.66	152.01
EBM-MMVAE+	100.65	95.37
MOE-EBM	98.23	90.32

Table 3. Ablation Results: FID

6. Future work

We plan to focus on two main avenues for future research. First, we will explore additional multimodal datasets, particularly those with high-resolution real images. Besides assessing generative coherence and perceptual performance, we aim to evaluate our model on various analytical tasks, including latent space analysis and mutual information analysis.

References

- Fan Bao, Shen Nie, Kaiwen Xue, Chongxuan Li, Shi Pu, Yaole Wang, Gang Yue, Yue Cao, Hang Su, and Jun Zhu. One transformer fits all distributions in multi-modal diffusion at scale. *arXiv preprint arXiv:2303.06555*, 2023. [1](#)
- Jiali Cui and Tian Han. Learning energy-based model via dual-mcmc teaching. *arXiv preprint arXiv:2312.02469*, 2023. [2](#)
- Jiali Cui, Ying Nian Wu, and Tian Han. Learning hierarchical features with joint latent space energy-based prior. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2218–2227, 2023. [1](#), [2](#)
- Tian Han, Yang Lu, Song-Chun Zhu, and Ying Nian Wu. Alternating back-propagation for generator network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2017. [2](#)
- Tian Han, Erik Nijkamp, Linqi Zhou, Bo Pang, Song-Chun Zhu, and Ying Nian Wu. Joint training of variational auto-encoder and latent energy-based model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7978–7987, 2020. [2](#)
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. [1](#)
- HyeongJoo Hwang, Geon-Hyeong Kim, Seunghoon Hong, and Kee-Eung Kim. Multi-view representation learning via total correlation objective. *Advances in Neural Information Processing Systems*, 34:12194–12207, 2021. [1](#), [2](#)
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [1](#)
- Erik Nijkamp, Mitch Hill, Song-Chun Zhu, and Ying Nian Wu. Learning non-convergent non-persistent short-run mcmc toward energy-based model. *Advances in Neural Information Processing Systems*, 32, 2019. [2](#)
- Emanuele Palumbo, Imant Daunhawer, and Julia E Vogt. Mm-vaes+: Enhancing the generative quality of multimodal vaes without compromises. In *The Eleventh International Conference on Learning Representations*. OpenReview, 2023. [1](#), [2](#), [3](#), [4](#)
- Bo Pang, Tian Han, Erik Nijkamp, Song-Chun Zhu, and Ying Nian Wu. Learning latent space energy-based prior model. *Advances in Neural Information Processing Systems*, 33:21994–22008, 2020. [1](#), [2](#)
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [1](#)
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022. [1](#)
- Yuge Shi, Brooks Paige, Philip Torr, et al. Variational mixture-of-experts autoencoders for multi-modal deep generative models. *Advances in neural information processing systems*, 32, 2019. [1](#), [2](#), [3](#)
- Thomas Sutter, Imant Daunhawer, and Julia Vogt. Multimodal generative learning utilizing jensen-shannon-divergence. *Advances in neural information processing systems*, 33:6100–6110, 2020. [1](#), [2](#), [3](#)
- Thomas M Sutter, Imant Daunhawer, and Julia E Vogt. Generalized multimodal elbo. *arXiv preprint arXiv:2105.02470*, 2021. [1](#), [2](#), [3](#)
- Masahiro Suzuki and Yutaka Matsuo. A survey of multimodal deep generative models. *Advanced Robotics*, 36(5-6):261–278, 2022. [1](#)
- Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder. *Advances in neural information processing systems*, 33:19667–19679, 2020. [1](#)
- Mike Wu and Noah Goodman. Multimodal generative models for scalable weakly-supervised learning. *Advances in neural information processing systems*, 31, 2018. [1](#), [2](#)
- Jianwen Xie, Yaxuan Zhu, Yifei Xu, Dingcheng Li, and Ping Li. A tale of two latent flows: Learning latent space normalizing flow with short-run langevin flow for approximate inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10499–10509, 2023. [1](#)
- Jing Zhang, Jianwen Xie, Nick Barnes, and Ping Li. Learning generative vision transformer with energy-based latent space for saliency prediction. *Advances in Neural Information Processing Systems*, 34:15448–15463, 2021. [2](#)