

# Learning Multimodal Latent Generative Models with Energy-Based Prior

Shiyu Yuan<sup>1✉</sup>, Jiali Cui<sup>2</sup>, Hanao Li<sup>2</sup>, and Tian Han<sup>2</sup>

<sup>1</sup> School of Systems and Enterprises,

<sup>2</sup> Department of Computer Science,

Stevens Institute of Technology, Hoboken NJ 07310, USA,

{syuan14, jcui7, hli136, than6}@stevens.edu

**Abstract.** Multimodal generative models have recently gained significant attention for their ability to learn representations across various modalities, enhancing joint and cross-generation coherence. However, most existing works use standard Gaussian or Laplacian distributions as priors, which may struggle to capture the diverse information inherent in multiple data types due to their unimodal and less informative nature. Energy-based models (EBMs), known for their expressiveness and flexibility across various tasks, have yet to be thoroughly explored in the context of multimodal generative models. In this paper, we propose a novel framework that integrates the multimodal latent generative model with the EBM. Both models can be trained jointly through a variational scheme. This approach results in a more expressive and informative prior, better-capturing of information across multiple modalities. Our experiments validate the proposed model, demonstrating its superior generation coherence.

**Keywords:** EBM · Multimodal latent generative model

## 1 Introduction

Generative model (GM) has made remarkable progress in generating high-quality image [12, 30], text [8, 18], and video [9], and recently, the multimodal GM [1, 7, 16, 19, 21, 27, 38] has garnered significant interest for providing a powerful framework that intrigues popular applications of *cross generation*, such as the image-to-text [17], text-to-image [21, 26]. However, these prior advances often focus on cross-modality learning by modeling the conditional dependency from one to the other while ignoring learning meaningful semantic representations shared across multimodalities. Learning a shared representation can play a critical role in enabling the downstream tasks [14], such as *joint generation*, thus representing an active ongoing research area.

To tackle the challenge of learning data representations, various methods of latent generative models [15, 29, 33] have been explored. In particular, the latent variable generative model consists of low-dimensional latent variables and a generation network, where the latent variables are learned to capture the data representation, and thus, the generation network can construct the high-dimensional

data by these learned latent variables. For multimodality learning, the multimodal VAEs [28, 35], have recently been developed. Given a set of modalities, these models primarily focus on encoding (inferring) latent variables from different modalities and fusing them into a single latent space. Specifically, MVAE [35] and MMVAE [28] factorize *joint posterior* that infers independent latent variables from different modalities and fuses them into one latent space. Such a single latent space thus needs to represent the data representation of multiple modalities and aims to capture their shared representations. However, these multimodal VAEs only consider less informative Gaussian or Laplacian prior to modeling the latent distribution, which can be limited in expressivity for complex data representations [3, 4, 25], resulting in an ineffectively learned latent generative model.

The energy-based model (EBM), on the other hand, is shown to be expressive and known to be powerful in modelling complex data distribution [5, 6]. In high-dimensional data space, it can be difficult for EBM learning as it typically involves Markov Chain Monte Carlo (MCMC) sampling for the EBM density [2, 11], while for latent space, the EBM can be formulated as the EBM prior [25, 36], reducing the burden of EBM prior sampling. Specifically, the EBM prior can be represented as an exponential tilting of the less informative reference distribution, where the energy function serves as a correction of reference distribution (e.g., Gaussian or Laplacian), rendering a more expressive prior distribution. However, most existing EBM works only consider modelling a single modality, and for multimodality, the latent generative model with the EBM prior still remains under-developed.

In this paper, we intend to explore the EBM prior to the more challenging multimodal learning task. In particular, we present a joint framework that is capable of leveraging the expressivity of EBM prior for modelling the shared semantic latent space from different multimodalities. With a set of modalities, we factorize a latent generative model with the EBM prior, in which the latent representation extracted from different modalities can be well-captured by the EBM prior. Compared to the uni-modal Gaussian or Laplacian prior, the EBM prior can be multi-modal and render more modelling capacity for complex representation learning, which in turn improves the generative power of the whole model and thus maintains semantic coherence for generated samples. Learning such EBM prior typically requires samples from the EBM prior and generator posterior, which is usually achieved by performing MCMC sampling for the EBM prior and generator posterior distribution. However, for the multimodal learning task, MCMC posterior sampling can be time-consuming as it may involve an additional inner-loop for computing the gradient through multiple generation networks from different modalities.

To ensure efficient posterior sampling and facilitate EBM prior learning, we employ the variational learning scheme and introduce an inference model [15, 28, 35] to approximate the generator posterior. With the inference model matching with the generator posterior, we can directly sample from the inference model and thus circumvent the burden of conducting MCMC posterior sampling. The

EBM prior sampling can be achieved by MCMC sampling because of the low-dimensional latent space and, more importantly, the lightweight energy function used. We demonstrate the proposed method can render superior performance in various benchmarks.

Our contributions can be summarized as follows:

- We propose the energy-based prior model for multimodal latent generative models to capture complex shared information within multiple modalities.
- We develop the variational training scheme where the generation model, inference model, and energy-based prior can be jointly and effectively learned.
- We conduct various experiments and ablation studies and demonstrate superior performance compared to Laplacian prior baselines.

## 2 Related Works

In this section, we present the background of multimodal variational autoencoders and energy-based models in detail.

**Multimodal Variational Autoencoders:** Multimodal Variational Autoencoders (VAEs) [28, 35] are based upon the standard VAE and have become the building blocks of all later works. Specifically, in these works, multiple pairs of encoder-decoder for the multi-modalities are learned, with the joint posterior obtained through the particular design of *product-of-experts* or *mixture-of-experts*. Inspired by such idea, mmJSD[31], MoPoE[32], MVTCAE[13], and MMVAE+[24] have explored VAE-based methods further to tackle the multimodal learning problem. mmJSD [31] adopted dynamic prior and modality-specific latent subsets, MVTCAE [13] used multi-view correlation shared representation built upon product of expert [35], MoPoE [32] includes subsets based on both product and mixture of expert [28, 35], and recent baseline MMVAE+ [24] incorporates modality-specific priors that are based on the mixture of expert[28]. However, these previous works used less informative uni-modal prior (e.g., Gaussian or Laplacian), which can be ineffective in capturing the complex latent representation shared across different modalities. To tackle this challenge, we propose learning latent space energy-based prior model, which can be more expressive in capturing the shared latent representation and rendering better synthesis across multi-modalities.

**Energy-Based Models:** The energy-based model (EBM) [6, 37] offers a flexible framework for approximating complex data distributions and is shown to be expressive in capturing the data regularities. In addition to data space EBMs, [25] proposes to apply EBM on latent space, which is capable of improving the generative power of the whole model. Learning such latent space EBM requires MCMC posterior and prior sampling, for which the MCMC posterior sampling can be computationally expensive as it requires inner loops of computing the backward gradient of the generation model. Compared to LEBM [25] that only studies one (single-) modality, our work targets the challenging multimodal learning problem

by learning latent space EBM to capture the shared latent representations. For multimodal learning problems, MCMC posterior sampling can be more difficult as it involves multiple generation networks during the learning. To alleviate the burden of MCMC posterior sampling, we further develop a variational learning scheme to facilitate efficient EBM sampling and learning. We show that the expressive EBM prior can be useful in the challenging multi-modal learning task.

### 3 Methodology

In this paper, we present a novel framework for effectively and efficiently modelling the multimodalities. In particular, we study learning the expressive EBM prior to capturing the shared and complex information across different modalities, for which the less informative Gaussian or Laplacian prior model can be limited in expressivity to effectively model. To facilitate efficiency for our EBM learning and sampling, we further develop a variational learning scheme that incorporates both the generator and inference model for jointly learning with the proposed EBM prior.

#### 3.1 Energy-based Prior for Multimodalities

Let  $\mathbf{z}$  be the latent variable and  $\mathbf{X}$  be the observation example that contains  $m$  modalities, i.e.,  $\mathbf{X} = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}\}$ . A joint distribution can be specified as

$$\begin{aligned} p_{\beta, \alpha}(\mathbf{X}, \mathbf{z}) &= p_{\beta}(\mathbf{X}|\mathbf{z})p_{\alpha}(\mathbf{z}) \quad \text{where} \\ p_{\beta}(\mathbf{X}|\mathbf{z}) &= p_{\beta_{(1)}}(\mathbf{x}^{(1)}|\mathbf{z})p_{\beta_{(2)}}(\mathbf{x}^{(2)}|\mathbf{z}) \cdots p_{\beta_{(m)}}(\mathbf{x}^{(m)}|\mathbf{z}) \end{aligned} \quad (1)$$

in which  $p_{\beta}(\mathbf{X}|\mathbf{z})$  ( $\beta$  collect  $\{\beta_{(1)}, \dots, \beta_{(m)}\}$ ) and  $p_{\alpha}(\mathbf{z})$  are the generation model and prior parameterized by  $\beta$  and  $\alpha$ , respectively. This factorization considers  $m$  modalities to be conditionally independently distributed while sharing the same latent space. The modality-common information is modelled by shared latent space  $p_{\alpha}(\mathbf{z})$ , and the modality-specific information is modelled by each  $p_{\beta_{(i)}}(\mathbf{x}^{(i)}|\mathbf{z})$ .

**Generation model.** The  $p_{\beta}(\mathbf{X}|\mathbf{z}) (= \prod_{i=1}^m p_{\beta_{(i)}}(\mathbf{x}^{(i)}|\mathbf{z}))$  consists of multiple generation models that seek to explain the high-dimensional  $\mathbf{x}^{(i)} \in \mathbb{R}^{D^{(i)}}$  by a shared low-dimensional latent vector  $\mathbf{z} \in \mathbb{R}^d$  ( $d < D^{(i)}$ ), i.e.,

$$\mathbf{x}^{(i)} = G_{\beta_{(i)}}(\mathbf{z}) + \epsilon \quad \text{where} \quad \epsilon \sim \mathcal{N}(0, I_{D^{(i)}}) \quad (2)$$

which implies  $p_{\beta_{(i)}}(\mathbf{x}^{(i)}|\mathbf{z}) \sim \mathcal{N}(G_{\beta_{(i)}}(\mathbf{z}), I_{D^{(i)}})$  with  $G_{\beta_{(i)}}$  being a top-down neural network that maps from  $\mathbf{z}$  to  $\mathbf{x}^{(i)}$ . We adopt such a generation model for its simplicity (also adopted in [15, 25, 36]), but it can also be other choices [29].

With a set of  $m$  multimodalities  $\mathbf{X} = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}\}$ , each generation model is designed to be modal-specific, forming the joint distribution  $p_{\beta_{(i)}, \alpha}(\mathbf{x}^{(i)}, \mathbf{z}) = p_{\beta_{(i)}}(\mathbf{x}^{(i)}|\mathbf{z})p_{\alpha}(\mathbf{z})$ . If  $\mathbf{z} \sim p_{\alpha}(\mathbf{z})$  can successfully capture the shared information

from different modalities,  $\mathbf{x}^{(i)} \sim p_{\beta_{(i)}, \alpha}(\mathbf{x}^{(i)}, \mathbf{z})$  and  $\mathbf{x}^{(j)} \sim p_{\beta_{(j)}, \alpha}(\mathbf{x}^{(j)}, \mathbf{z})$  can be drawn with semantic coherence.

**Energy-based prior.** To effectively capture the semantic information shared across the multimodalities, we intend to learn an expressive prior model. In particular, we study learning the energy-based prior model defined as

$$p_{\alpha}(\mathbf{z}) = \frac{1}{\mathbb{Z}(\alpha)} \exp[f_{\alpha}(\mathbf{z})] p_0(\mathbf{z}) \quad (3)$$

where  $\mathbb{Z}(\alpha) (= \int_{\mathbf{z}} \exp[f_{\alpha}(\mathbf{z})] p_0(\mathbf{z}) d\mathbf{z})$  is the normalizing constant or partition function,  $f_{\alpha}(\cdot)$  is the energy function parameterized with  $\alpha$ , and  $p_0(\mathbf{z})$  is the referenced distribution usually assumed to be standard Laplacian [24, 28, 31]. Such EBM prior generative model has seen success in modelling the data distribution of *single* modality [25, 36] while in this paper, we intend to explore its effectiveness in modelling the multimodalities.

For multimodalities, capturing the shared content across modalities may serve as the key ingredient toward generating semantic coherent samples. [28, 31, 32, 35] adopt different frameworks for multimodal learning but only consider the less informative, uni-modal Laplacian prior, which in turn limits the model expressivity, leading to an ineffectively learned model. The proposed latent generative model, on the other hand, is learned with the EBM prior, which is known to be powerful in capturing the data regularity and complex distribution. With well-captured shared information for multimodalities, generated samples of each modality thus can maintain strong semantic coherence across different modalities.

### 3.2 Learning and Sampling

**Maximum likelihood estimation.** Given  $n$  observed examples  $\{\mathbf{X}_1, \dots, \mathbf{X}_n\}$  with each  $\mathbf{X}_i$  containing  $m$  modalities  $\mathbf{X}_i = \{\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(m)}\}$ , the generator model (Eqn. 1) can be learned by maximum likelihood estimation (MLE) as

$$\begin{aligned} L(\theta) &= \sum_{i=1}^n \log p_{\theta}(\mathbf{X}_i) = \sum_{i=1}^n \log \int_{\mathbf{z}} p_{\beta}(\mathbf{X}_i | \mathbf{z}) p_{\alpha}(\mathbf{z}) d\mathbf{z} \\ &= \sum_{i=1}^n \log \int_{\mathbf{z}} p_{\beta_{(1)}}(\mathbf{x}_i^{(1)} | \mathbf{z}) \cdots p_{\beta_{(m)}}(\mathbf{x}_i^{(m)} | \mathbf{z}) p_{\alpha}(\mathbf{z}) d\mathbf{z} \end{aligned} \quad (4)$$

where  $\theta$  collect learning parameters  $(\beta, \alpha)$ . With a large number of  $n$ , maximizing Eqn. 4 is equivalent to minimizing the KL-divergence, i.e.,  $\text{KL}(p_{\text{data}}(\mathbf{X}) || p_{\theta}(\mathbf{X}))$ , and learning  $\theta$  can be done by computing the gradient as

$$\frac{\partial}{\partial \theta} L(\theta) = \mathbb{E}_{p_{\theta}(\mathbf{z} | \mathbf{X})} \left[ \frac{\partial}{\partial \theta} \log p_{\theta}(\mathbf{X}, \mathbf{z}) \right] \quad (5)$$

which requires samples from the generator posterior [10]. To obtain the posterior samples, it can be typically achieved by performing MCMC sampling for

$p_\theta(\mathbf{z}|\mathbf{X}) \propto p_\beta(\mathbf{X}|\mathbf{z})p_\alpha(\mathbf{z})$ . However, for the proposed model on the multimodalities task,  $p_\beta(\mathbf{X}|\mathbf{z})$  consists of  $m$  of generation models of each modality, which makes the MCMC posterior sampling inefficient.

**Variational learning scheme.** To ensure efficient learning and posterior sampling, we introduce the inference model  $q_\phi(\mathbf{z}|\mathbf{X})$  as an approximation model for the generator posterior. To facilitate learning with the generation model for multimodalities, our inference model is defined to be

$$q_\phi(\mathbf{z}|\mathbf{X}) = \frac{1}{m} \sum_{i=1}^m q_{\phi_{(i)}}(\mathbf{z}|\mathbf{x}^{(i)}) \quad (6)$$

where  $q_{\phi_{(i)}}(\mathbf{z}|\mathbf{x}^{(i)}) \sim \mathcal{N}(u_{\phi_{(i)}}(\mathbf{x}^{(i)}), V_{\phi_{(i)}}(\mathbf{x}^{(i)}))$ . Such an inference model serves as the *mixture of experts* and is also adopted in [28]. Specifically, [28] parameterize pairs of generation and inference model to be modal-specific, i.e.,  $\beta = (\beta_1, \dots, \beta_m)$  and  $\phi = (\phi_1, \dots, \phi_m)$ , such that each  $p_{\beta_i}(\mathbf{x}^{(i)}|\mathbf{z})$  and  $q_{\phi_i}(\mathbf{z}|\mathbf{x}^{(i)})$  only focus on one modality as a pair of *decoder* and *encoder*. We adopt such parameterization for its effectiveness.

With such introduced inference model  $q_\phi(\mathbf{z}|\mathbf{X})$ , a joint KL-divergence can be minimized,

$$-L(\beta, \phi, \alpha) = \text{KL}(p_{\text{data}}(\mathbf{X})q_\phi(\mathbf{z}|\mathbf{X})||p_\beta(\mathbf{X}|\mathbf{z})p_\alpha(\mathbf{z})) \quad (7)$$

which is equivalent to maximizing

$$L(\beta, \phi, \alpha) = \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{q_{\phi_{(i)}}(\mathbf{z}|\mathbf{x}^{(i)})} [\log \frac{p_{\beta, \alpha}(\mathbf{X}, \mathbf{z})}{q_\phi(\mathbf{z}|\mathbf{X})}] \quad (8)$$

Therefore, we compute the gradient  $\frac{\partial}{\partial \beta, \phi} L(\beta, \phi, \alpha)$  as,

$$\begin{aligned} \frac{\partial}{\partial \beta, \phi} \frac{1}{m} \sum_{i=1}^m [ & \mathbb{E}_{q_{\phi_{(i)}}(\mathbf{z}|\mathbf{x}^{(i)})} [\log p_{\beta_{(i)}}(\mathbf{x}^{(i)}|\mathbf{z})] + \sum_{j=1, j \neq i}^m \mathbb{E}_{q_{\phi_{(i)}}(\mathbf{z}|\mathbf{x}^{(i)})} [\log p_{\beta_{(j)}}(\mathbf{x}^{(j)}|\mathbf{z})] \\ & + \mathbb{E}_{q_{\phi_{(i)}}(\mathbf{z}|\mathbf{x}^{(i)})} [\log \frac{p_\alpha(\mathbf{z})}{q_\phi(\mathbf{z}|\mathbf{X})}] ] \end{aligned} \quad (9)$$

the gradient  $\frac{\partial}{\partial \alpha} L(\beta, \phi, \alpha)$  is computed as

$$\frac{\partial}{\partial \alpha} L(\beta, \phi, \alpha) = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{X})} \left[ \frac{\partial}{\partial \alpha} f_\alpha(\mathbf{z}) \right] - \mathbb{E}_{p_\alpha(\mathbf{z})} \left[ \frac{\partial}{\partial \alpha} f_\alpha(\mathbf{z}) \right] \quad (10)$$

where  $\mathbf{z} \sim q_\phi(\mathbf{z}|\mathbf{X})$  is inferred from the inference model (Eqn. 6), which is a fused joint posterior (i.e.,  $\mathbf{z} \sim \frac{1}{m} \sum_{i=1}^m q_{\phi_{(i)}}(\mathbf{z}|\mathbf{x}^{(i)})$  as average weighted over inferred latent vectors from all modalities).

**Sampling from EBM prior.** Learning EBM (Eqn. 10) requires samples from the EBM prior, which can be accomplished by conducting MCMC sampling, such as Langevin dynamics (LD) [20]. It iterates as

$$\mathbf{z}_{\tau+1} = \mathbf{z}_{\tau} - \frac{s^2}{2} \frac{\partial}{\partial \mathbf{z}} [\log p_{\alpha}(\mathbf{z}_{\tau})] + s \cdot \epsilon_{\tau} \quad \text{where } \epsilon_{\tau} \sim \mathcal{N}(0, I_d) \quad (11)$$

where  $s$  is the step size,  $\epsilon$  is the Gaussian noise, and  $\tau$  is the time step of Langevin dynamics. As  $s \rightarrow 0$ , and  $\tau \rightarrow \infty$ , the marginal distribution of  $\mathbf{z}$  can asymptotically converge to the target  $p_{\alpha}(\mathbf{z})$  as the stationary distribution. In this work, we employ Laplacian as the initial distribution (i.e.  $\mathbf{z}_0 \sim \mathcal{N}(0, I_d)$ ) and conduct short-run Langevin dynamics, which can also provide meaningful learning signals [22, 23].

**Connection to ELBO.** The VAEs compute the evidence lower bound (ELBO) as the learning objective, which is

$$\mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{X})} [\log p_{\beta}(\mathbf{X}|\mathbf{z})] - \text{KL}(q_{\phi}(\mathbf{z}|\mathbf{X})||p_0(\mathbf{z})) \quad (12)$$

where  $p_0(\mathbf{z})$  is usually assumed to be Gaussian. Whereas, our objective  $L(\beta, \phi, \alpha)$  (Eqn. 7) can be decomposed into the form

$$L(\beta, \phi, \alpha) = \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{X})} [\log p_{\beta}(\mathbf{X}|\mathbf{z})] - \text{KL}(q_{\phi}(\mathbf{z}|\mathbf{X})||p_{\alpha}(\mathbf{z})) \quad (13)$$

which is closely related to the ELBO of VAEs.

Different from Gaussian or Laplacian prior  $p_0(\mathbf{z})$  in Eqn. 12, we consider learning the EBM prior  $p_{\alpha}(\mathbf{z})$ . With a set of  $m$  modalities, each inference model  $q_{\phi(i)}(\mathbf{z}|\mathbf{x}^{(i)})$  together forms a *mixture of experts*, and thus the Gaussian or Laplacian prior can be limited in expressivity to capture and match with the posterior distribution, while the EBM prior can be more expressive and multi-modal, leading to a well-learned shared latent space.

---

#### Algorithm 1 Learning Energy-Based Multimodal Model

---

**Input:** observation examples  $\{\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(m)}\}_{i=1}^n$ , iteration number  $T$ , Langevin steps  $K$  and step size  $s$ . iteration step  $t = 0$

**repeat**

**Posterior Sampling:** Given  $\{\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(m)}\}$ , obtain  $\{\mathbf{z}_i^{(1)}, \dots, \mathbf{z}_i^{(m)}\}$  from inference model (Eqn. 6).

**Prior Sampling:** Obtain EBM prior samples  $\mathbf{z}^-$  by performing Langevin dynamics (Eqn. 11) with  $K$  and  $s$ .

**Learn  $\phi$  and  $\beta$ :** Update learning parameter  $\phi, \beta$  with  $\{\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(m)}\}$  and  $\{\mathbf{z}_i^{(1)}, \dots, \mathbf{z}_i^{(m)}\}$  using Eqn. 9.

**Learn  $\alpha$ :** Update learning parameter  $\alpha$  with inferred samples  $\{\mathbf{z}_i^{(1)}, \dots, \mathbf{z}_i^{(m)}\}$  using Eqn. 6 and  $\mathbf{z}^-$  using Eqn. 10.

    Let  $t = t + 1$ ;

**until**  $t = T$

---

## 4 Experiments

In this section, we conduct various experiments to demonstrate the expressiveness of the EBM prior in capturing shared latent information across multimodalities more effectively than less-informative unimodal priors (e.g., Gaussian or Laplacian priors). In our experiments, we follow the prior arts [25] and train our model on standard multimodal datasets, such as the PolyMNIST [32] and MNIST-SVHN [28]. We evaluate the performance of our model in terms of joint coherence (Sec. 4.1) and cross coherence (Sec. 4.2). Additionally, we demonstrate the applicability and flexibility of the proposed method, showing its capacity for generalization to various factorizations (Sec. 4.3), including the incorporation of model-specific factors [24]. And we visualize the Markov transition on image synthesis and corresponding generated text applied to Caltech UCSD Birds (CUB) dataset [34]. Furthermore, we conduct ablation studies (Sec. 4.4) to gain a deeper understanding of our approach. We provided our code in <https://github.com/syyuan2021/Learning-Multimodal-%Latent-Generative-Models-with-EBM>

**Baseline Method.** For comparisons, our direct baselines include MVAE[35] and MMVAE [28], and we also compare with recent mmJSD [31], MoPoE [32], MVTACE[13] and MMVAE+ [24] that are developed based on the foundation methods of *product-of-expert* [35] and *mixture-of-expert* [28]. For a fair comparison, we adopt the same generation and inference network structures, as well as the pre-trained classifiers from [32] for PolyMNIST, and follow [28]’s implementation to train the classifier for MNIST-SVHN.

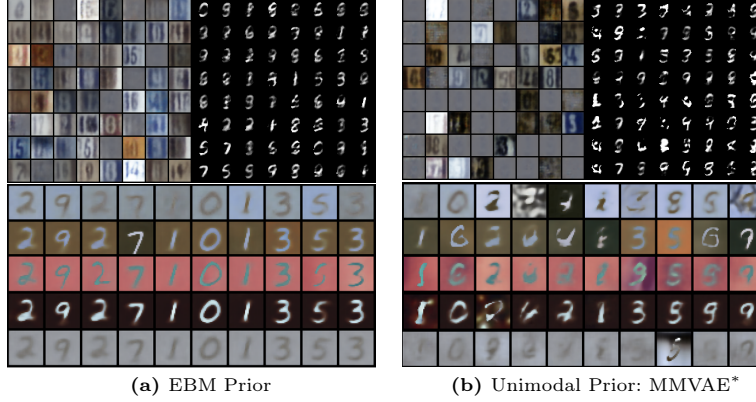
### 4.1 Joint Coherence

With a set of  $m$  modalities, we assess our model in generating *unconditional* image synthesis that maintains strong coherence across different modalities. If the EBM prior is well-learned, it should be capable of sampling the latent variables that capture semantic representations and thus generating the image synthesis semantically consistent between multimodalities. We measure such coherence by computing the classification accuracy of generated images of each modality with corresponding classifiers. Higher accuracy indicates better consistency between predicted categories of multimodal synthesis.

We report the quantitative results in Tab. 1, where our EBM prior shows superior performance compared to baseline models. Both numerically and visually, the EBM prior successfully captures complex latent representations, whereas the unimodal prior (Laplacian) is less expressive and limited in modeling meaningful semantic representations. Synthesis of joint generation can be found in Fig. 1.

**Table 1:** Accuracy of Joint Coherence and Cross Coherence for PolyMNIST and MNIST-SVHN datasets. The PolyMNIST results for MVAE are referenced from [24], and MNIST-SVHN results for MVAE are referenced from [28]. MMVAE\* refers to *mixture of expert* without important sampling.

Model	Joint Coherence		Cross Coherence		
	PolyMNIST	MNIST-SVHN	PolyMNIST	MNIST-SVHN	
				M->S	S->M
Ours	<b>0.735</b> ↑	<b>0.412</b> ↑	<b>0.857</b> ↑	<b>0.195</b> ↑	<b>0.657</b> ↑
MVAE	0.080	0.127	0.298	0.095	0.093
MMVAE*	0.232	0.215	0.844	0.169	0.523



**Fig. 1:** Qualitative Results of Joint Generation on MNIST-SVHN (top) and PolyMNIST (bottom))

## 4.2 Cross Coherence

Next, we evaluate our model in cross-modal generation. Cross-modal generation is to generate image synthesis from input modality  $\mathbf{x}^{(i)}$  to target modality  $\mathbf{x}^{(j)}$  ( $i \neq j$ ). Specifically, given input  $\mathbf{x}^{(i)}$ , we first obtain inferred latent vector  $\mathbf{z} \sim q_\phi(\mathbf{z}|\mathbf{x}^{(i)})$  and then generate synthesis via generation model  $\mathbf{x}^{(j)} \sim p_{\beta_j}(\mathbf{x}^{(j)}|\mathbf{z})$ . With our EBM prior, the inference and generation model can be learned to integrate both the modality bias and the energy-based refinement and thus the latent vectors inferred from  $\mathbf{x}^{(i)}$  carry the semantic representation shared with  $\mathbf{x}^{(j)}$ , leading to coherent input  $\mathbf{x}^{(i)}$  and output  $\mathbf{x}^{(j)}$ . To measure the coherence, we compute the classification accuracy (same classifiers used in Sec. 4.1) of the predicted category of  $\mathbf{x}^{(j)}$  and true category of  $\mathbf{x}^{(i)}$ . Higher accuracy means both the  $\mathbf{x}^{(i)}$  and  $\mathbf{x}^{(j)}$  share higher similarity in semantic features (e.g., digit classes in PolyMNIST). The results can be found in Tab. 1, in which the proposed method can render competitive performance compared to the baseline methods, suggesting the effectiveness of our EBM prior in multimodalities learning.

### 4.3 Model Generalization

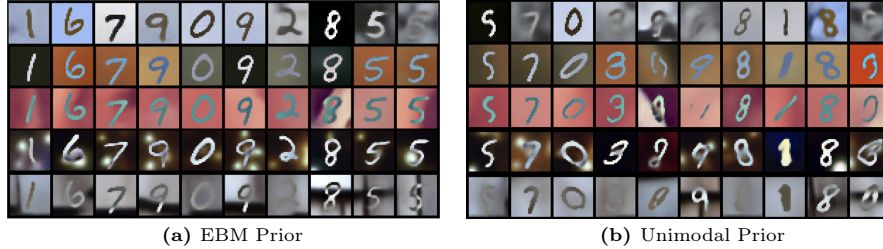
This paper studies learning a novel framework for foundation multimodal latent generative models, such as MVAE and MMVAE, which serve as the building blocks of various recent works. In this section, we highlight the applicable capability of the proposed method toward other prior advances [24]. These advances usually factorize additional *modal-specific prior* or *modal-subset* to benefit the model complexity in learning the complex multimodalities. In particular, MMVAE+ develop their latent generative model as

$$p_{\beta}(\mathbf{X}, \mathbf{z}, \mathbf{W}) = p_{\beta}(\mathbf{X}|\mathbf{z}, \mathbf{W})p_0(\mathbf{z})p_0(\mathbf{W}) \quad (14)$$

where the prior model  $p_0(\mathbf{W})$  is modal-specific (i.e.,  $\mathbf{W} = \{\mathbf{w}^{(1)}, \dots, \mathbf{w}^{(m)}\}$ ) and is introduced to improve the expressivity of the whole prior model (compared to Eqn. 1). However, the shared latent representation is still modelled by  $p_0(\mathbf{z})$ , which is assumed to be Laplacian and can be less informative.

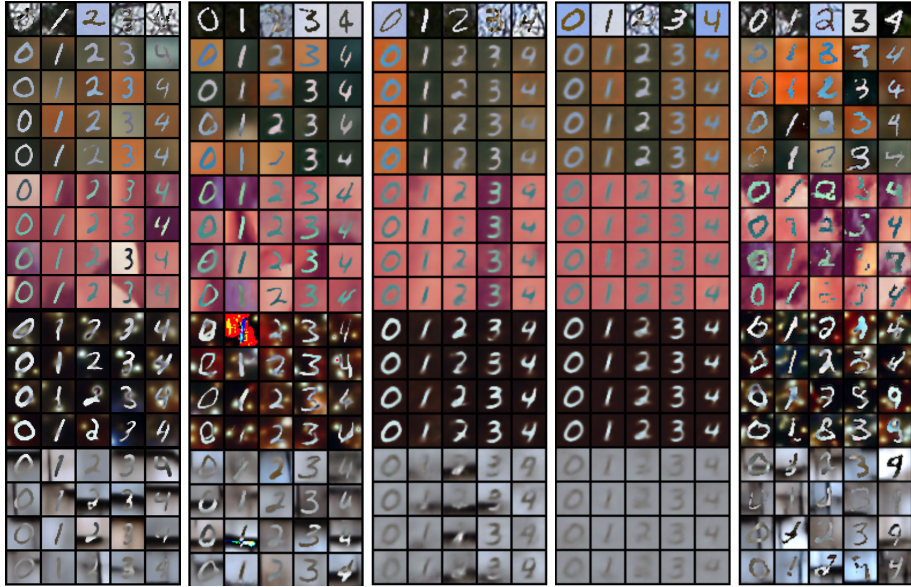
**Table 2:** Accuracy of Joint Coherence and Cross Coherence for PolyMNIST. Results of MMVAE+, MoPoE, MVTCAE, and mmJSD are reported in [24]. Ours+ refers to generalizing our model with modality-specific priors.

Model	Joint Coherence	Cross Coherence
Ours+	<b>0.848</b> ↑	<b>0.889</b> ↑
MMVAE+	0.344	0.869
MoPoE	0.141	0.720
MVTCAE	0.003	0.591
mmJSD	0.060	0.778



**Fig. 2:** Qualitative results of Joint Generation on PolyMNIST

**Generalization.** Our model is flexible and can be adapted to incorporate such a modal-specific factor, i.e.,  $p_{\beta, \alpha}(\mathbf{X}, \mathbf{z}, \mathbf{W}) = p_{\beta}(\mathbf{X}|\mathbf{z}, \mathbf{W})p_{\alpha}(\mathbf{z})p_0(\mathbf{W})$ , such that the shared latent representation is modelled by the proposed EBM prior  $p_{\alpha}(\mathbf{z})$ , which can be multi-modal and more expressive than the unimodal prior model  $p_0(\mathbf{z})$ . To examine the effectiveness, we follow MMVAE+ and train our model



**Fig. 3:** Qualitative results of Cross Generation on PolyMNIST. From left to right, cross generation results are from (a) EBM prior, (b) MMVAE+, (c) MoPoE, (d) mmJSD, and (e) MVTCAE.

on PolyMNIST with the same network structures. We report the results of joint coherence and cross coherence in Tab. 2 and the qualitative results in Fig. 2 and Fig. 3.

**Caltech UCSD Birds Dataset (CUB).** To examine the scalability, we further train our model on the challenging multimodal dataset Caltech UCSD Birds (CUB)[34]. The proposed EBM can be viewed as an exponential tilting of the reference distribution and thus can correct the less informative prior model toward being more expressive. In this section, we intend to demonstrate the expressivity of our EBM prior in correcting the unimodal prior model. In practice, we utilize the same network structure and train MMVAE+ on CUB, and then we learn our EBM prior with the pre-trained generation and inference model. If our EBM prior can be learned well, the quality of synthesis generated by our EBM prior samples should be better than samples of the unimodal prior.

For illustration, we visualize the Markov chain transition during Langevin dynamics on generated images and text in Fig. 4. Specifically, the Markov chain transition starts from the latent drawn from the unimodal prior  $p_0(\mathbf{z})$  and then progresses as an iterative sampler for our EBM prior  $p_\alpha(\mathbf{z})$ . It can be seen that the quality of image and text synthesis becomes better as the Langevin dynamics progresses and, more importantly, at the final step, the generated images and text render better semantic coherence, which further indicates the effectiveness of the proposed EBM prior.



**Fig. 4:** Visualization of Markov chain transition on image synthesis and corresponding generated text. **Left:** visualization of Markov chain transition where left column images are generated by *unimodal* prior sample, and right column images are generated by *EBM* prior sample. **Right:** jointly generated text corresponding to each row in left figure. The sentence at the top is generated by *unimodal* prior sample, and the sentence at the bottom is generated by *EBM* prior sample.

#### 4.4 Ablation Studies

In the previous sections, we demonstrated the effectiveness of using an EBM prior in multimodal contexts, illustrating that our model can better capture complex data representations across different modalities. To further investigate our model’s capabilities, we examine the impact of our EBM prior settings on performance through a series of experiments conducted on the PolyMNIST dataset.

**Energy-Based Model MCMC Steps:** We first examine the influence of the number of MCMC steps, denoted as  $S$ , when sampling the latent variables from EBM prior. We fixed the network architecture to consist of 4 layers with 400 hidden units each. As observed in the middle two rows of Tab. 3, a smaller number of MCMC steps results in poorer performance. However, increasing the number of MCMC steps is computationally expensive, leading to a trade-off between computation time and model performance.

**Energy-Based Model Layers:** Then we examine whether more information transformation and interaction will increase coherence scores. So we fixed hidden dimension and Langevin steps while increase EBM layer  $L$  from 4 to 6. As shown in the first two rows of Tab. 3, increasing the number of layers  $L$  leads to better coherence results.

**Energy-Based Model Complexity:** Increasing the model dimension allows the model to learn more complex and representative features. To examine scalability, we increase the dimension of hidden units  $D$  to investigate the performance impact of capacity in capturing and processing information. We maintain the number of layers at 4 and use 50 Langevin steps. As observed in the bottom two rows of Tab. 3, increasing the number of units results in better coherence performance.

**Table 3:** Ablation studies for energy-based models with different numbers of hidden units  $D$ , MCMC steps  $S$ , and network layers  $L$ .

Model	Joint Coherence	Cross Coherence
Ours (D=200, L=4, S=50)	0.580	0.854
Ours (D=200, L=6, S=50)	0.682	0.840
Ours (D=400, L=4, S=30)	0.641	0.825
Ours (D=400, L=4, S=50)	<b>0.735</b>	<b>0.857</b>
Ours (D=200, L=4, S=50)	0.580	0.854
Ours (D=400, L=4, S=50)	<b>0.735</b>	<b>0.857</b>

## 5 Conclusions and Future Work

In this paper, we present a novel framework for multimodal latent generative models with an EBM prior. This expressive and flexible prior can better represent multimodal data complexity and capture shared information among modalities. Our experiments demonstrate significantly improved coherence of synthesized samples across different modalities compared to baseline models. Our proposed model also facilitates cross-generation between modalities, as validated by experimental results.

However, the proposed method is based on a simple *mixture of experts* scheme with an EBM prior to optimize the multimodal ELBO. Approaches that provide a tighter bound on ELBO, such as importance sampling and stratified sampling with an EBM prior, have not been fully explored. Additionally, to gain a better understanding of shared information learning schemes under multimodal contexts, other expressive priors such as normalizing flow or hierarchical priors will be considered in our future research. For EBM learning, we use latent variables from the variationally inferred posterior, which is less accurate in approximating the true posterior compared to methods such as MCMC sampling. While the latter can be time-consuming, this trade-off either sacrifices generative performance or computational efficiency. This dilemma is non-trivial in multimodal generative problems. Lastly, we will explore the scalability of our EBM prior in future work by investigating its performance on realistic multimodal datasets. This will allow us to evaluate the model’s effectiveness in more complex and varied real-world scenarios, further validating its applicability and robustness.

## Acknowledgements

This work is supported in part by NSF IIS-2339604. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

## Bibliography

- [1] Baltrušaitis, T., Ahuja, C., Morency, L.P.: Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence* **41**(2), 423–443 (2018)
- [2] Cui, J., Han, T.: Learning energy-based model via dual-mcmc teaching. *37th Conference on Neural Information Processing Systems* (2023)
- [3] Cui, J., Wu, Y.N., Han, T.: Learning hierarchical features with joint latent space energy-based prior. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 2218–2227 (October 2023)
- [4] Cui, J., Wu, Y.N., Han, T.: Learning joint latent space ebm prior model for multi-layer generator. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 3603–3612 (June 2023)
- [5] Du, Y., Li, S., Tenenbaum, J., Mordatch, I.: Improved contrastive divergence training of energy based models. *arXiv preprint arXiv:2012.01316* (2020)
- [6] Gao, R., Song, Y., Poole, B., Wu, Y.N., Kingma, D.P.: Learning energy-based models by diffusion recovery likelihood. *arXiv preprint arXiv:2012.08125* (2020)
- [7] Gu, S., Chen, D., Bao, J., Wen, F., Zhang, B., Chen, D., Yuan, L., Guo, B.: Vector quantized diffusion model for text-to-image synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10696–10706 (2022)
- [8] Guo, J., Lu, S., Cai, H., Zhang, W., Yu, Y., Wang, J.: Long text generation via adversarial training with leaked information. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 32 (2018)
- [9] Gupta, A., Yu, L., Sohn, K., Gu, X., Hahn, M., Fei-Fei, L., Essa, I., Jiang, L., Lezama, J.: Photorealistic video generation with diffusion models. *arXiv preprint arXiv:2312.06662* (2023)
- [10] Han, T., Lu, Y., Zhu, S.C., Wu, Y.N.: Alternating back-propagation for generator network. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 31 (2017)
- [11] Han, T., Nijkamp, E., Zhou, L., Pang, B., Zhu, S.C., Wu, Y.N.: Joint training of variational auto-encoder and latent energy-based model. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7978–7987 (2020)
- [12] Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
- [13] Hwang, H., Kim, G.H., Hong, S., Kim, K.E.: Multi-view representation learning via total correlation objective. *Advances in Neural Information Processing Systems* **34**, 12194–12207 (2021)
- [14] Jiang, Q., Chen, C., Zhao, H., Chen, L., Ping, Q., Tran, S.D., Xu, Y., Zeng, B., Chilimbi, T.: Understanding and constructing latent modality structures

- in multi-modal representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7661–7671 (2023)
- [15] Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
  - [16] Li, B., Qi, X., Lukasiewicz, T., Torr, P.: Controllable text-to-image generation. *Advances in Neural Information Processing Systems* **32** (2019)
  - [17] Li, J., Selvaraju, R., Gotmare, A., Joty, S., Xiong, C., Hoi, S.C.H.: Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems* **34**, 9694–9705 (2021)
  - [18] Lin, K., Li, D., He, X., Zhang, Z., Sun, M.T.: Adversarial ranking for language generation. *Advances in neural information processing systems* **30** (2017)
  - [19] Lin, X., Bertasius, G., Wang, J., Chang, S.F., Parikh, D., Torresani, L.: Vx2text: End-to-end learning of video-based text generation from multimodal inputs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7005–7015 (2021)
  - [20] Neal, R.M., et al.: Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo* **2**(11), 2 (2011)
  - [21] Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint arXiv:2112.10741 (2021)
  - [22] Nijkamp, E., Hill, M., Han, T., Zhu, S.C., Wu, Y.N.: On the anatomy of mcmc-based maximum likelihood learning of energy-based models. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 5272–5280 (2020)
  - [23] Nijkamp, E., Hill, M., Zhu, S.C., Wu, Y.N.: Learning non-convergent non-persistent short-run mcmc toward energy-based model. *Advances in Neural Information Processing Systems* **32** (2019)
  - [24] Palumbo, E., Daunhawer, I., Vogt, J.E.: Mmvae+: Enhancing the generative quality of multimodal vaes without compromises. In: The Eleventh International Conference on Learning Representations. OpenReview (2023)
  - [25] Pang, B., Han, T., Nijkamp, E., Zhu, S.C., Wu, Y.N.: Learning latent space energy-based prior model. *Advances in Neural Information Processing Systems* **33**, 21994–22008 (2020)
  - [26] Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sasstry, G., Aspell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
  - [27] Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems* **35**, 36479–36494 (2022)

- [28] Shi, Y., Paige, B., Torr, P., et al.: Variational mixture-of-experts autoencoders for multi-modal deep generative models. *Advances in neural information processing systems* **32** (2019)
- [29] Sønderby, C.K., Raiko, T., Maaløe, L., Sønderby, S.r.K., Winther, O.: Ladder variational autoencoders. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 29. Curran Associates, Inc. (2016), <https://proceedings.neurips.cc/paper/2016/file/6ae07dcb33ec3b7c814df797cbda0f87-Paper.pdf>
- [30] Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. *International Conference on Learning Representations* (2020)
- [31] Sutter, T., Daunhawer, I., Vogt, J.: Multimodal generative learning utilizing jensen-shannon-divergence. *Advances in neural information processing systems* **33**, 6100–6110 (2020)
- [32] Sutter, T.M., Daunhawer, I., Vogt, J.E.: Generalized multimodal elbo. *International Conference on Learning Representations* (2021)
- [33] Vahdat, A., Kautz, J.: Nvae: A deep hierarchical variational autoencoder. *Advances in neural information processing systems* **33**, 19667–19679 (2020)
- [34] Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset. [authors.library.caltech.edu](http://authors.library.caltech.edu) (2011)
- [35] Wu, M., Goodman, N.: Multimodal generative models for scalable weakly-supervised learning. *Advances in neural information processing systems* **31** (2018)
- [36] Xiao, Z., Han, T.: Adaptive multi-stage density ratio estimation for learning latent space energy-based model. In: *NeurIPS* (2022), [http://papers.nips.cc/paper\\_files/paper/2022/hash/874a4d89f2d04b4bcf9a2c19545cf040-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/874a4d89f2d04b4bcf9a2c19545cf040-Abstract-Conference.html)
- [37] Xiao, Z., Kreis, K., Kautz, J., Vahdat, A.: Vaebm: A symbiosis between variational autoencoders and energy-based models. *arXiv preprint arXiv:2010.00654* (2020)
- [38] Zhang, H., Koh, J.Y., Baldrige, J., Lee, H., Yang, Y.: Cross-modal contrastive learning for text-to-image generation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 833–842 (2021)