

MULTI-MODAL LATENT DIFFUSION

Anonymous authors

Paper under double-blind review

ABSTRACT

Multi-modal data-sets are ubiquitous in modern applications, and multi-modal Variational Autoencoders are a popular family of models that aim to learn a joint representation of the different modalities. However, existing approaches suffer from a coherence–quality tradeoff, where models with good generation quality lack generative coherence across modalities, and vice versa. We discuss the limitations underlying the unsatisfactory performance of existing methods, to motivate the need for a different approach. We propose a novel method that uses a set of independently trained, uni-modal, deterministic autoencoders. Individual latent variables are concatenated into a common latent space, which is fed to a masked diffusion model to enable generative modeling. We also introduce a new multi-time training method to learn the conditional score network for multi-modal diffusion. Our methodology substantially outperforms competitors in both generation quality and coherence, as shown through an extensive experimental campaign.

1 INTRODUCTION

Multi-modal generative modelling is a crucial area of research in machine learning that aims to develop models capable of generating data according to multiple modalities, such as images, text, audio, and more. This is important because real-world observations are often captured in various forms, and combining multiple modalities describing the same information can be an invaluable asset. For instance, images and text can provide complementary information in describing an object, audio and video can capture different aspects of a scene. Multi-modal generative models can also help in tasks such as data augmentation (He et al., 2023; Azizi et al., 2023; Sariyildiz et al., 2023), missing modality imputation (Antelmi et al., 2019; Da Silva–Filarder et al., 2021; Zhang et al., 2023; Tran et al., 2017), and conditional generation (Huang et al., 2022; Lee et al., 2019b).

Multi-modal models have flourished over the past years and have seen a tremendous interest from academia and industry, especially in the content creation sector. Whereas most recent approaches focus on specialization, by considering text as primary input to be associated mainly to images (Rombach et al., 2022; Saharia et al., 2022; Ramesh et al., 2022; Tao et al., 2022; Wu et al., 2022; Nichol et al., 2022; Chang et al., 2023) and videos (Blattmann et al., 2023; Hong et al., 2023; Singer et al., 2022), in this work we target an established literature whose scope is more general, and in which all modalities are considered equally important. A large body of work rely on extensions of the Variational Autoencoder (VAE) (Kingma & Welling, 2014) to the multi-modal domain: initially interested in learning joint latent representation of multi-modal data, such works have mostly focused on generative modeling. Multi-modal generative models aim at *high-quality* data generation, as well as generative *coherence* across all modalities. These objectives apply to both joint generation of new data, and to conditional generation of missing modalities, given a disjoint set of available modalities.

In short, multi-modal VAES rely on combinations of uni-modal VAES, and the design space consists mainly in the way the uni-modal latent variables are combined, to construct the joint posterior distribution. Early work such as Wu & Goodman (2018) adopt a product of experts approach, whereas others Shi et al. (2019) consider a mixture of expert approach. Product-based models achieve high generative quality, but suffer in terms of both joint and conditional coherence. This was found to be due to experts mis-calibration issues (Shi et al., 2019; Sutter et al., 2021). On the other hand, mixture-based models produce coherent but qualitatively poor samples. A first attempt to address the so called **coherence–quality tradeoff** (Daunhawer et al., 2022) is represented by the mixture of product of experts approach (Sutter et al., 2021). However recent comparative studies (Daunhawer et al., 2022) show that none of the existing approaches fulfill both the generative quality

and coherence criteria. A variety of techniques aim at finding a better operating point, such as contrastive learning techniques (Shi et al., 2021), hierarchical schemes (Vasco et al., 2022), total correlation based calibration of single modality encoders (Hwang et al., 2021), or different training objectives Sutter et al. (2020). More recently, the work in (Palumbo et al., 2023) considers explicitly separated shared and private latent spaces to overcome the aforementioned limitations.

By expanding on results presented in (Daunhawer et al., 2022), in Section 2 we further investigate the tradeoff between generative coherence and quality, and argue that it is intrinsic to all variants of multi-modal VAEs. We indicate two root causes of such problem: latent variable collapse (Alemi et al., 2018; Dieng et al., 2019) and information loss due to mixture sub-sampling. To tackle these issues, in this work, we propose in Section 3 a new approach which uses a set of independent, uni-modal *deterministic* auto-encoders whose latent variables are simply concatenated in a joint latent variable. Joint and conditional generative capabilities are provided by an additional model that learns a probability density associated to the joint latent variable. We propose an extension of score-based diffusion models (Song et al., 2021b) to operate on the multi-modal latent space. We thus derive both forward and backward dynamics that are compatible with the multi-modal nature of the latent data. In section 4 we propose a novel method to train the multi-modal score network, such that it can both be used for joint and conditional generation. Our approach is based on a guidance mechanism, which we compare to alternatives. We label our approach Multi-modal Latent Diffusion (MLD).

Our experimental evaluation of MLD in Section 5 provides compelling evidence of the superiority of our approach for multi-modal generative modeling. We compare MLD to a large variety of VAE-based alternatives, on several real-life multi-modal data-sets, in terms of generative quality and both joint and conditional coherence. Our model outperforms alternatives in all possible scenarios, even those that are notoriously difficult because modalities might be only loosely correlated. Note that recent work also explore the joint generation of multiple modalities Ruan et al. (2023); Hu et al. (2023), but such approaches are application specific, e.g. text-to-image, and essentially only target two modalities. When relevant, we compare our method to additional recent alternatives to multi-modal diffusion (Bao et al., 2023; Wesego & Rooshenas, 2023), and show superior performance of MLD.

2 LIMITATIONS OF MULTI-MODAL VAEs

In this work, we consider multi-modal VAEs (Wu & Goodman, 2018; Shi et al., 2019; Sutter et al., 2021; Palumbo et al., 2023) as the standard modeling approach to tackle both joint and conditional generation of multiple modalities. Our goal here is to motivate the need to go beyond such a standard approach, to overcome limitations that affect multi-modal VAEs, which result in a trade-off between generation quality and generative coherence (Daunhawer et al., 2022; Palumbo et al., 2023).

Consider the random variable $X = \{X^1, \dots, X^M\} \sim p_D(x^1, \dots, x^M)$, consisting in the set of M of modalities sampled from the (unknown) multi-modal data distribution p_D . We indicate the marginal distribution of a single modality by $X^i \sim p_D^i(x^i)$ and the collection of a generic subset of modalities by $X^A \sim p_D^A(x^A)$, with $X^A \stackrel{\text{def}}{=} \{X^i\}_{i \in A}$, where $A \subset \{1, \dots, M\}$ is a set of indexes. For example: given $A = \{1, 3, 5\}$, then $X^A = \{X^1, X^3, X^5\}$.

We begin by considering uni-modal VAEs as particular instances of the Markov chain $X \rightarrow Z \rightarrow \hat{X}$, where Z is a latent variable and \hat{X} is the generated variable. Models are specified by the two conditional distributions, called the encoder $Z |_{X=x} \sim q_\psi(z|x)$, and the decoder $\hat{X} |_{Z=z} \sim p_\theta(\hat{x}|z)$. Given a prior distribution $p_n(z)$, the objective is to define a generative model whose samples are distributed as closely as possible to the original data.

In the case of multi-modal VAEs, we consider the general family of Mixture of Product of Experts (MOPOE) (Sutter et al., 2021), which includes as particular cases many existing variants, such as Product of Experts (MVAE) (Wu & Goodman, 2018) and Mixture of Expert (MMVAE) (Shi et al., 2019). Formally, a collection of K arbitrary subsets of modalities $S = \{A_1, \dots, A_K\}$, along with weighting coefficients $\omega_i \geq 0$, $\sum_{i=1}^K \omega_i = 1$, define the posterior $q_\psi(z|x) = \sum_i \omega_i q_{\psi, A_i}^i(z|x^{A_i})$, with $\psi = \{\psi^1, \dots, \psi^K\}$. To lighten the notation, we use q_{ψ, A_i} in place of q_{ψ, A_i}^i noting that the various q_{ψ, A_i}^i can have both different parameters ψ^{A_i} and functional form. For example, in the MOPOE (Sutter et al., 2021) parametrization, we have: $q_{\psi, A_i}(z|x^{A_i}) = \prod_{j \in A_i} q_{\psi^j}(z|x^j)$. Our exposition is

more general and not limited to this assumption. The selection of the posterior can be understood as the result induced by the two step procedure where i) each subset of modalities A_i is encoded into specific latent variables $Y_i \sim q_{\psi^{A_i}}(\cdot | x^{A_i})$ and ii) the latent variable Z is obtained as $Z = Y_i$ with probability ω_i . Optimization is performed w.r.t. the following evidence lower bound (ELBO) (Daunhawer et al., 2022; Sutter et al., 2021):

$$\mathcal{L} = \sum_i \omega_i \int p_D(x) q_{\psi^{A_i}}(z | x^{A_i}) \log p_\theta(x|z) - \log \frac{q_{\psi^{A_i}}(z | x^{A_i})}{p_n(z)} dz dx. \quad (1)$$

A well-known limitation called the latent collapse problem (Alemi et al., 2018; Dieng et al., 2019) affects the quality of latent variables Z . Consider the hypothetical case of arbitrary flexible encoders and decoders: then, posteriors with zero mutual information with respect to model inputs are valid maximizers of Equation (1). To prove this, it is sufficient to substitute the posteriors $q_{\psi^{A_i}}(z | x^{A_i}) = p_n(z)$ and $p_\theta(x|z) = p_D(x)$ into the Equation (1) to observe that the optimal value $\mathcal{L} = \int p_D(x) \log p_D(x) dx$ is achieved (Alemi et al., 2018; Dieng et al., 2019). The problem of information loss is exacerbated in the case of multi-modal VAEs (Daunhawer et al., 2022). Intuitively, even if the encoders $q_{\psi^{A_i}}(z | x^{A_i})$ carry relevant information about their inputs X^{A_i} , step ii) of the multi-modal encoding procedure described above induces a further information bottleneck. A fraction ω_i of the time, the latent variable Z will be a copy of Y_i , that only provides information about the subset X^{A_i} . No matter how good the encoding step is, the information about $X^{\{1, \dots, M\} \setminus A}$ that is not contained in X^{A_i} cannot be retrieved.

Furthermore, if the latent variable carries zero mutual information w.r.t. the multi-modal input, a coherent *conditional* generation of a set of modalities given others is impossible, since $\hat{X}^{A_1} \perp X^{A_2}$ for any generic sets A_1, A_2 . While the factorization $p_\theta(x|z) = \prod_{i=1}^M p_{\theta^i}(x^i|z)$, $\theta = \{\theta_1, \dots, \theta_M\}$ — where we use p_{θ^i} instead of $p_{\theta^i}^j$ to unclutter the notation — could enforce preservation of information and guarantee a better quality of the *jointly* generated data, in practice, the latent collapse phenomenon induces multi-modal VAEs to converge toward sub-optimal operating regime. When the posterior $q_\psi(z|x)$ collapses onto the uninformative prior $p_n(z)$, the ELBO in Equation (1) reduces to the sum of modality independent reconstruction terms $\sum_i \omega_i \sum_{j \in A_i} \int p_D^j(x^j) p_n(z) (\log p_{\theta^j}(x^j|z)) dz dx^j$.

In this case, flexible decoders can similarly ignore the latent variable and converge to the solution $p_{\theta^j}(x^j|z) = p_D^j(x^j)$ where, paradoxically, the quality of the approximation of the various marginal distributions is extremely high, while there is a complete lack of joint coherence.

General principles to avoid latent collapse consist in explicitly forcing the learning of informative encoders $q_\theta(z|x)$ via β -annealing of the Kullback-Leibler (KL) term in the ELBO and the reduction of the representational power of encoders and decoders. While β -annealing has been explored in the literature (Wu & Goodman, 2018) with limited improvements, reducing the flexibility of encoders/decoders clearly impacts the generation quality. Hence the presence of a trade-off: to improve coherence, the flexibility of encoders/decoders should be constrained, which in turns hurt generative quality. This trade-off has been recently addressed in the literature of multi-modal VAEs (Daunhawer et al., 2022; Palumbo et al., 2023), but our experimental results in Section 5 indicate that there is ample room for improvement, and that a new approach is truly needed.

3 OUR APPROACH: MULTI-MODAL LATENT DIFFUSION

We propose a new method for multi-modal generative modeling that, by design, does not suffer from the limitations discussed in Section 2. Our objective is to enable both high-quality and coherent joint/conditional data generation, using a simple design (see Appendix A for a schematic representation). As an overview, we use deterministic uni-modal autoencoders, whereby each modality X^i is encoded through its encoder e_{ψ^i} , which is a short form for $e_{\psi^i}^i$, into the modality specific latent variable Z^i and decoded into the corresponding $\hat{X}^i = d_{\theta^i}(Z^i)$. Our approach can be interpreted as a latent variable model where the different latent variables Z^i are concatenated as $Z = [Z^1, \dots, Z^M]$. This corresponds to the parametrization of the two conditional distributions as $q_\psi(z|x) = \prod_{i=1}^M \delta(z^i - e_{\psi^i}(x^i))$ and $p_\theta(\hat{x}|z) = \prod_{i=1}^M \delta(\hat{x}^i - d_{\theta^i}(z^i))$, respectively. Then, in place of an ELBO, we optimize the parameters of our autoencoders by minimizing the following sum of

modality specific losses:

$$\mathcal{L} = \sum_{i=1}^M \mathcal{L}_i, \quad \mathcal{L}_i = \int p_D^i(x^i) l^i(x^i - d_{\theta^i}(e_{\psi^i}(x^i))) dx^i, \quad (2)$$

where l^i can be any valid distance function, e.g, the square norm $\|\cdot\|^2$. Parameters ψ^i, θ^i are modality specific: then, minimization of Equation (2) corresponds to individual training of the different autoencoders. Since the mapping from input to latent is deterministic, there is no loss of information between X and Z .¹ Moreover, this choice avoids any form of interference in the back-propagated gradients corresponding to the uni-modal reconstruction losses. Consequently gradient conflicts issues (Javaloy et al., 2022), where stronger modalities pollute weaker ones, are avoided.

To enable such a simple design to become a generative model, it is sufficient to generate samples from the induced latent distribution $Z \sim q_\psi(z) = \int p_D(x) q_\psi(z|x) dx$ and decode them as $\hat{X} = d_\theta(Z) = [d_{\theta^1}(Z^1), \dots, d_{\theta^M}(Z^M)]$. To obtain such samples, we follow the two-stage procedure described in Loaiza-Ganem et al. (2022); Tran et al. (2021), where samples from the lower dimensional $q_\psi(z)$ are obtained through an appropriate generative model. We consider score-based diffusion models in latent space (Rombach et al., 2022; Vahdat et al., 2021) to solve this task, and call our approach Multi-modal Latent Diffusion (MLD). It may be helpful to clarify, at this point, that the two-stage training of MLD is carried out separately. Uni-modal deterministic autoencoders are pre-trained first, followed by the training of the score-based diffusion model, which is explained in more detail later.

To conclude the overview of our method, for joint data generation, one can sample from noise, perform backward diffusion, and then decode the generated multi-modal latent variable to obtain the corresponding data samples. For conditional data generation, given one modality, the reverse diffusion is guided by this modality, while the other modalities are generated by sampling from noise. The generated latent variable is then decoded to obtain data samples of the missing modality.

3.1 JOINT AND CONDITIONAL MULTI-MODAL LATENT DIFFUSION PROCESSES

In the first stage of our method, the deterministic encoders project the input modalities X^i into the corresponding latent spaces Z^i . This transformation induces a distribution $q_\psi(z)$ for the latent variable $Z = [Z^1, \dots, Z^M]$, resulting from the concatenation of uni-modal latent variables.

Joint generation. To generate a new sample for all modalities we use a simple score-based diffusion model in latent space (Sohl-Dickstein et al., 2015; Song et al., 2021b; Vahdat et al., 2021; Loaiza-Ganem et al., 2022; Tran et al., 2021). This requires reversing a stochastic noising process, starting from a simple, Gaussian distribution. Formally, the noising process is defined by a Stochastic Differential Equation (SDE) of the form:

$$dR_t = \alpha(t)R_t dt + g(t)dW_t, \quad R_0 \sim q(r, 0), \quad (3)$$

where $\alpha(t)R_t$ and $g(t)$ are the drift and diffusion terms, respectively, and W_t is a Wiener process. The time-varying probability density $q(r, t)$ of the stochastic process at time $t \in [0, T]$, where T is finite, satisfies the Fokker-Planck equation (Oksendal, 2013), with initial conditions $q(r, 0)$. We assume uniqueness and existence of a stationary distribution $\rho(r)$ for the process Equation (3).² The forward diffusion dynamics depend on the initial conditions $R_0 \sim q(r, 0)$. We consider $R_0 = Z$ to be the initial condition for the diffusion process, which is equivalent to $q(r, 0) = q_\psi(r)$. Under loose conditions (Anderson, 1982), a time-reversed stochastic process exists, with a new SDE of the form:

$$dR_t = (-\alpha(T-t)R_t + g^2(T-t)\nabla \log(q(R_t, T-t))) dt + g(T-t)dW_t, \quad R_0 \sim q(r, T), \quad (4)$$

indicating that, in principle, simulation of Equation (4) allows to generate samples from the desired distribution $q(r, 0)$. In practice, we use a **parametric score network** $s_\chi(r, t)$ to approximate the true score function, and we approximate $q(r, T)$ with the stationary distribution $\rho(r)$. Indeed, the generated data distribution $q(r, 0)$ is close (in KL sense) to the true density as described by Song et al. (2021a); Franzese et al. (2023):

$$\text{KL}[q_\psi(r) \parallel q(r, 0)] \leq \frac{1}{2} \int_0^T g^2(t) \mathbb{E}[\|s_\chi(R_t, t) - \nabla \log q(R_t, t)\|^2] dt + KL[q(r, T) \parallel \rho(r)], \quad (5)$$

¹Since the measures are not absolutely continuous w.r.t the Lebesgue measure, mutual information is $+\infty$.

²This is not necessary for the validity of the method Song et al. (2021a)

where the first term on the r.h.s is referred to as score-matching objective, and is the loss over which the score network is optimized, and the second is a vanishing term for $T \rightarrow \infty$.

To conclude, joint generation of all modalities is achieved through the simulation of the reverse-time SDE in Equation (4), followed by a simple decoding procedure. Indeed, optimally trained decoders (achieving zero in Equation (2)) can be used to transform $Z \sim q_\psi(z)$ into samples from $\int p_\theta(x|z)q_\psi(z)dz = p_D(x)$.

Conditional generation. Given a generic partition of all modalities into non overlapping sets $A_1 \cup A_2$, where $A_2 = (\{1, \dots, M\} \setminus A_1)$, conditional generation requires samples from the conditional distribution $q_\psi(z^{A_1} | z^{A_2})$, which are based on *masked* forward and backward diffusion processes.

Given conditioning latent modalities z^{A_2} , we consider a modified forward diffusion process with initial conditions $R_0 = \mathcal{C}(R_0^{A_1}, R_0^{A_2})$, with $R_0^{A_1} \sim q_\psi(r^{A_1} | z^{A_2})$, $R_0^{A_2} = z^{A_2}$. The composition operation $\mathcal{C}(\cdot)$ concatenates generated (R^{A_1}) and conditioning latents (z^{A_2}). As an illustration, consider $A_1 = \{1, 3, 5\}$, such that $X^{A_1} = \{X^1, X^3, X^5\}$, and $A_2 = \{2, 4, 6\}$ such that $X^{A_2} = \{X^2, X^4, X^6\}$. Then, $R_0 = \mathcal{C}(R_0^{A_1}, R_0^{A_2}) = \mathcal{C}(R_0^{A_1}, z^{A_2}) = [R_0^1, z^2, R_0^3, z^4, R_0^5, z^6]$.

More formally, we define the masked forward diffusion SDE:

$$dR_t = m(A_1) \odot [\alpha(t)R_t dt + g(t)dW_t], \quad q(r, 0) = q_\psi(r^{A_1} | z^{A_2})\delta(r^{A_2} - z^{A_2}). \quad (6)$$

The mask $m(A_1)$ contains M vectors u^i , one per modality, and with the corresponding cardinality. If modality $j \in A_1$, then $u^j = \mathbf{1}$, otherwise $u^j = \mathbf{0}$. Then, the effect of masking is to “freeze” throughout the diffusion process the part of the random variable R_t corresponding to the conditioning latent modalities z^{A_2} . We naturally associate to this modified forward process the conditional time varying density $q(r, t | z^{A_2}) = q(r^{A_1}, t | z^{A_2})\delta(r^{A_2} - z^{A_2})$.

To sample from $q_\psi(z^{A_1} | z^{A_2})$, we derive the reverse-time dynamics of Equation (6) as follows:

$$dR_t = m(A_1) \odot [(-\alpha(T-t)R_t + g^2(T-t)\nabla \log(q(R_t, T-t | z^{A_2}))) dt + g(T-t)dW_t], \quad (7)$$

with initial conditions $R_0 = \mathcal{C}(R_0^{A_1}, z^{A_2})$ and $R_0^{A_1} \sim q(r^{A_1}, T | z^{A_2})$. Then, we approximate $q(r^{A_1}, T | z^{A_2})$ by its corresponding steady state distribution $\rho(r^{A_1})$, and the true (conditional) score function $\nabla \log(q(r, t | z^{A_2}))$ by a conditional score network $s_\chi(r^{A_1}, t | z^{A_2})$.

4 GUIDANCE MECHANISMS TO LEARN THE CONDITIONAL SCORE NETWORK

A correctly optimized score network $s_\chi(r, t)$ allows, through simulation of Equation (4), to obtain samples from the joint distribution $q_\psi(z)$. Similarly, a *conditional* score network $s_\chi(r^{A_1}, t | z^{A_2})$ allows, through the simulation of Equation (7), to sample from $q_\psi(z^{A_1} | z^{A_2})$. In Section 4.1 we extend guidance mechanisms used in classical diffusion models to allow multi-modal conditional generation. A naïve alternative is to rely on the unconditional score network $s_\chi(r, t)$ for the conditional generation task, by casting it as an *in-painting* objective. Intuitively, any missing modality could be recovered in the same way as a uni-modal diffusion model can recover masked information. In Section 4.2 we discuss the implicit assumptions underlying in-painting from an information theoretic perspective, and argue that, in the context of multi-modal data, such assumptions are difficult to satisfy. Our intuition is corroborated by ample empirical evidence, where our method consistently outperform alternatives.

4.1 MULTI-TIME DIFFUSION

We propose a modification to the classifier-free guidance technique (Ho & Salimans, 2022) to learn a score network that can generate conditional and unconditional samples from any subset of modalities. Instead of training a separate score network for each possible combination of conditional modalities, which is computationally infeasible, we use a single architecture that accepts all modalities as inputs and a *multi-time vector* $\tau = [t_1, \dots, t_M]$. The multi-time vector serves two purposes: it is both a conditioning signal and the time at which we observe the diffusion process.

Training: learning the conditional score network relies on randomization. As discussed in Section 3.1, we consider an arbitrary partitioning of all modalities in two disjoint sets, A_1 and A_2 . The set A_2

contains randomly selected conditioning modalities, while the remaining modalities belong to set A_1 . Then, during training, the parametric score network estimates $\nabla \log(q(r, t | z^{A_2}))$, whereby the set A_2 is randomly chosen at every step. This is achieved by the *masked diffusion process* from Equation (6), which only diffuses modalities in A_1 . More formally, the score network input is $R_t = \mathcal{C}(R_t^{A_1}, Z^{A_2})$, along with a multi-time vector $\tau(A_1, t) = t[\mathbb{1}(1 \in A_1), \dots, \mathbb{1}(M \in A_1)]$. As a follow-up of the example in Section 3.1, given $A_1 = \{1, 3, 5\}$, such that $X^{A_1} = \{X^1, X^3, X^5\}$, and $A_2 = \{2, 4, 6\}$ such that $X^{A_2} = \{X^2, X^4, X^6\}$, then, $\tau(A_1, t) = [t, 0, t, 0, t, 0]$.

More precisely, the algorithm for the multi-time diffusion training (see A for the pseudo-code) proceeds as follows. At each step, a set of conditioning modalities A_2 is sampled from a predefined distribution ν , where $\nu(\emptyset) \stackrel{\text{def}}{=} \Pr(A_2 = \emptyset) = d$, and $\nu(U) \stackrel{\text{def}}{=} \Pr(A_2 = U) = (1-d)/(2^M - 1)$ with $U \in \mathcal{P}(\{1, \dots, M\}) \setminus \emptyset$, where $\mathcal{P}(\{1, \dots, M\})$ is the powerset of all modalities. The corresponding set A_1 and mask $m(A_1)$ are constructed, and a sample X is drawn from the training data-set. The corresponding latent variables $Z^{A_1} = \{e_{\psi}^i(X^i)\}_{i \in A_1}$ and $Z^{A_2} = \{e_{\psi}^i(X^i)\}_{i \in A_2}$ are computed using the pre-trained encoders, and a diffusion process starting from $R_0 = \mathcal{C}(Z^{A_1}, Z^{A_2})$ is simulated for a randomly chosen diffusion time t , using the conditional forward SDE with the mask $m(A_1)$. The score network is then fed the current state R_t and multi-time vector $\tau(A_1, t)$, and the difference between the score network’s prediction and the true score is computed, applying the mask $m(A_1)$. The score network parameters are updated using stochastic gradient descent, and this process is repeated for a total of L training steps. Clearly, when $A_2 = \emptyset$, training proceeds as for an un-masked diffusion process, since the mask $m(A_1)$ allows all latent variables to be diffused.

Conditional generation: any valid numerical integration scheme for Equation (7) can be used for conditional sampling (see A for an implementation using the Euler-Maruyama integrator). First, conditioning modalities in the set A_2 are encoded into the corresponding latent variables $z^{A_2} = \{e^j(x^j)\}_{j \in A_2}$. Then, numerical integration is performed with step-size $\Delta t = T/N$, starting from the initial conditions $R_0 = \mathcal{C}(R_0^{A_1}, z^{A_2})$, with $R_0^{A_1} \sim \rho(r^{A_1})$. At each integration step, the score network s_{χ} is fed the current state of the process and the multi-time vector $\tau(A_1, \cdot)$. Before updating the state, the masking is applied. Finally, the generated modalities are obtained thanks to the decoders as $\hat{X}^{A_1} = \{d_{\theta}^j(R_T^j)\}_{j \in A_1}$. Inference time conditional generation is not randomized: conditioning modalities are the ones that are available, whereas the remaining are the ones we wish to generate.

Any-to-any multi-modality has been recently studied through the composition of modality-specific diffusion models (Tang et al., 2023), by designing cross-attention and training procedures that allow arbitrary conditional generation. The work by Tang et al. (2023) relies on latent interpolation of input modalities, which is akin to mixture models, and uses it as conditioning signal for individual diffusion models. This is substantially different from the joint nature of the multi-modal latent diffusion we present in our work: instead of forcing entanglement through cross-attention between score networks, our model relies on joint diffusion process, whereby modalities naturally co-evolve according to the diffusion process. Another recent work (Wu et al., 2023) targets multi-modal conversational agents, whereby the strong, underlying assumption is to consider one modality, i.e., text, as a guide for the alignment and generation of other modalities. Even if conversational objectives are orthogonal to our work, techniques akin to instruction following for cross-generation, are an interesting illustration of the powerful capabilities of in-context learning of LLMs (Xie et al., 2022; Min et al., 2022).

4.2 IN-PAINTING AND ITS IMPLICIT ASSUMPTIONS

Under certain assumptions, given an unconditional score network $s_{\chi}(r, t)$ that approximates the true score $\nabla \log q(r, t)$, it is possible to obtain a conditional score network $s_{\chi}(r^{A_1}, t | z^{A_2})$, to approximate $\nabla \log q(r^{A_1}, t | z^{A_2})$. We start by observing the equality:

$$q(r^{A_1}, t | z^{A_2}) = \int q(\mathcal{C}(r^{A_1}, r^{A_2}), t | z^{A_2}) dr^{A_2} = \int \frac{q(z^{A_2} | \mathcal{C}(r^{A_1}, r^{A_2}), t)}{q_{\psi}(z^{A_2})} q(\mathcal{C}(r^{A_1}, r^{A_2}), t) dr^{A_2}, \quad (8)$$

where, with a slight abuse of notation, we indicate with $q(z^{A_2} | \mathcal{C}(r^{A_1}, r^{A_2}), t)$ the density associated to the event: the portion corresponding to A_2 of the latent variable Z is equal to z^{A_2} given that the whole diffused latent R_t at time t , is equal to $\mathcal{C}(r^{A_1}, r^{A_2})$. In the literature, the quantity $q(z^{A_2} | \mathcal{C}(r^{A_1}, r^{A_2}), t)$ is typically approximated by dropping its dependency on r^{A_1} . This approxima-

tion can be used to manipulate Equation (8) as $q(r^{A_1}, t | z^{A_2}) \simeq \int q(r^{A_2}, t | z^{A_2})q(r^{A_1}, t | r^{A_2}, t) dr$. Further Monte-Carlo approximations (Song et al., 2021b; Lugmayr et al., 2022) of the integral allow implementation of a practical scheme, where an approximate conditional score network is used to generate conditional samples. This approach, known in the literature as *in-painting*, provides high quality results in several *uni-modal* application domains (Song et al., 2021b; Lugmayr et al., 2022).

The KL divergence between $q(z^{A_2} | \mathcal{C}(r^{A_1}, r^{A_2}), t)$ and $q(z^{A_2} | r^{A_2}, t)$ quantifies, fixing r^{A_1}, r^{A_2} , the discrepancy between the true and approximated conditional probabilities. Similarly, the expected KL divergence $\Delta = \int q(r, t) \text{KL}[q(z^{A_2} | \mathcal{C}(r^{A_1}, r^{A_2}), t) | | q(z^{A_2} | r^{A_2}, t)] dr$, provides information about the average discrepancy. Simple manipulations allow to recast this as a discrepancy in terms of mutual information $\Delta = I(Z^{A_2}; R_t^{A_1}, R_t^{A_2}) - I(Z^{A_2}; R_t^{A_2})$. Information about Z^{A_2} is contained in $R_t^{A_2}$, as the latter is the result of a diffusion with the former as initial conditions, corresponding to the Markov chain $R_t^{A_2} \rightarrow Z^{A_2}$, and in $R_t^{A_1}$ through the Markov chain $Z^{A_2} \rightarrow Z^{A_1} \rightarrow R_t^{A_1}$. The positive quantity Δ is close to zero whenever the rate of loss of information w.r.t initial conditions is similar for the two subsets A_1, A_2 . In other terms, $\Delta \simeq 0$ whenever out of the whole R_t , the portion $R_t^{A_2}$ is a sufficient statistic for Z^{A_2} .

The assumptions underlying the approximation are in general not valid in the case of multi-modal learning, where the robustness to stochastic perturbations of latent variables corresponding to the various modalities can vary greatly. Our claim are supported empirically by an ample analysis on real data in B, where we show that multi-time diffusion approach consistently outperforms in-painting.

5 EXPERIMENTS

We compare our method MLD to MVAE Wu & Goodman (2018), MMVAE Shi et al. (2019), MOPOE Sutter et al. (2021), Hierarchical Generative Model (NEXUS) Vasco et al. (2022) and Multi-view Total Correlation Autoencoder (MVTCAE) Hwang et al. (2021), MMVAE+ Palumbo et al. (2023) re-implementing competitors in the same code base as our method, and selecting their best hyper-parameters (as indicated by the authors). For fair comparison, we use the same encoder/decoder architecture for all the models. For MLD, the score network is implemented using a simple stacked multilayer perceptron (MLP) with skip connections (see A for more details).

Evaluation metrics. *Coherence* is measured as in Shi et al. (2019); Sutter et al. (2021); Palumbo et al. (2023), using pre-trained classifiers on the generated data and checking the consistency of their outputs. *Generative quality* is computed using Fréchet Inception Distance (FID) Heusel et al. (2017) and Fréchet Audio Distance (FAD) Kilgour et al. (2019) scores for images and audio respectively. Full details on the metrics are included in C. All results are averaged over 5 seeds (We report standard deviation in E).

Results. Overall, MLD largely outperforms alternatives from the literature, **both** in terms of coherence and generative quality. VAE-based models suffer from a coherence–quality trad-off and modality collapse for highly heterogeneous data-sets. We proceed to show this on several standard benchmarks from the multi-modal VAE-based literature (see C for details on the data-sets).

The first data-set we consider is **MNIST-SVHN** ((Shi et al., 2019)), where the two modalities differ in complexity. High variability, noise and ambiguity makes attaining good coherence for the SVHN modality a challenging task. Overall, MLD outperforms all VAE-based alternatives in terms of coherency, especially in terms of joint generation and conditional generation of MNIST given SVHN, see Table 1. Mixture models (MMVAE, MOPOE) suffer from modality collapse (poor SVHN generation), whereas product of experts (MVAE, MVTCAE) generate better quality samples at the expense of SVHN to MNIST conditional coherence. Joint generation is poor for all VAE models. Interestingly, these models also fail at SVHN self-reconstruction which we discuss in E. MLD achieves the best performance also in terms of generation quality, as confirmed also by qualitative results (Figure 1) showing for example how MLD conditionally generates multiple SVHN digits within one sample, given the input MNIST image, whereas other methods fail to do so.

The Multi-modal Handwritten Digits data-set (**MHD**) (Vasco et al., 2022) contains gray-scale digit images, motion trajectory of the hand writing and sounds of the spoken digits. In our experiments, we do not use the label as a forth modality. While digit image and trajectory share a good amount of information, the sound modality contains a lot more of modality specific variation. Consequently,

Table 1: Generation coherence and quality for MNIST-SVHN (M :MNIST, S: SVHN). The generation quality is measured in terms of Fréchet Modality Distance (FMD) for MNIST and FID for SVHN.

Models	Coherence (% \uparrow)			Quality (\downarrow)			
	Joint	M \rightarrow S	S \rightarrow M	Joint(M)	Joint(S)	M \rightarrow S	S \rightarrow M
MVAE	38.19	48.21	28.57	13.34	68.9	68.0	13.66
MMVAE	37.82	11.72	67.55	25.89	146.82	393.33	53.37
MOPOE	39.93	12.27	68.82	20.11	129.2	373.73	43.34
NEXUS	40.0	16.68	70.67	13.84	98.13	281.28	53.41
MVTCAE	48.78	81.97	49.78	12.98	52.92	69.48	13.55
MMVAE+	17.64	13.23	29.69	26.60	121.77	240.90	35.11
MMVAE+(K=10)	41.59	55.3	56.41	19.05	67.13	75.9	18.16
MLD (ours)	85.22	83.79	79.13	3.93	56.36	57.2	3.67




Figure 1: Qualitative results for MNIST-SVHN. For each model we report: MNIST to SVHN conditional generation in the left, SVHN to MNIST conditional generation in the right.

conditional generation involving the sound modality, along with joint generation, are challenging tasks. Coherency-wise (Table 2) MLD outperforms all the competitors where the biggest difference is seen in joint and sound to other modalities generation (in the latter task MVTCAE performs better than other competitors but is still worse than MLD). MLD dominates alternatives also in terms of generation quality (Table 3). This is true both for image, sound modalities, for which some VAE-based models suffer in producing high quality results, demonstrating the limitation of these methods in handling highly heterogeneous modalities. MLD, in the other hand, achieves high generation quality for all modalities, possibly due to the independent training of the autoencoders avoiding interference.

Table 2: Generation Coherence (%) for MHD (Higher is better). Line above refer to the generated modality while the observed modalities subset are presented below.

Models	Joint	I (Image)			T (Trajectory)			S (Sound)		
		T	S	T,S	I	S	I,S	I	T	I,T
MVAE	37.77	11.68	26.46	28.4	95.55	26.66	96.58	58.87	10.76	58.16
MMVAE	34.78	99.7	69.69	84.74	<u>99.3</u>	85.46	92.39	49.95	50.14	50.17
MOPOE	48.84	<u>99.64</u>	68.67	<u>99.69</u>	99.28	<u>87.42</u>	99.35	50.73	51.5	56.97
NEXUS	26.56	94.58	<u>83.1</u>	95.27	88.51	76.82	93.27	70.06	75.84	89.48
MVTCAE	42.28	99.54	72.05	99.63	99.22	72.03	<u>99.39</u>	<u>92.58</u>	<u>93.07</u>	<u>94.78</u>
MMVAE+	41.67	98.05	84.16	91.88 \pm	97.47	81.16	89.31	64.34	65.42	64.88
MMVAE+(k=10)	42.60	99.44	89.75	94.7	99.44	89.58	95.01	87.15	87.99	87.57
MLD (ours)	98.34	99.45	<u>88.91</u>	99.88	99.58	<u>88.92</u>	99.91	97.63	97.7	98.01

The POLYMNIST data-set (Sutter et al., 2021) consists of 5 modalities synthetically generated by using MNIST digits and varying the background images. The homogeneous nature of the modalities is expected to mitigate gradient conflict issues in VAE-based models, and consequently reduce modality collapse. However, MLD still outperforms all alternatives, as shown Figure 2. Concerning generation coherence, MLD achieves the best performance in all cases with the single exception of a single observed modality. On the qualitative performance side, not only MLD is superior to alternatives, but its results are stable when more modalities are considered, a capability that not all competitors share.

Finally, we explore the Caltech Birds CUB (Shi et al., 2019) data-set, following the same experimentation protocol in Daunhawer et al. (2022) by using real bird images (instead of ResNet-features as in Shi et al. (2019)). Figure 3 presents qualitative results for caption to image conditional generation. MLD is the only model capable of generating bird images with convincing coherence. Clearly, none of the VAE-based methods is able to achieve sufficient caption to image conditional generation quality using the same simple autoencoder architecture. Note that an image autoencoder with larger capacity improves considerably MLD generative performance, suggesting that careful engineering applied to modality specific autoencoders is a promising avenue for future work. We report quantitative

Table 3: Generation quality for **MHD** in terms of FMD for image and trajectory modalities and FAD for the sound modality (Lower is better).

Models	I (Image)				T (Trajectory)				S (Sound)			
	Joint	T	S	T,S	Joint	I	S	LS	Joint	I	T	LT
MVAE	<u>94.9</u>	93.73	92.55	91.08	39.51	20.42	38.77	19.25	14.14	<u>14.13</u>	14.08	14.17
MMVAE	224.01	22.6	789.12	170.41	16.52	0.5	30.39	6.07	22.8	22.61	23.72	23.01
MOPOE	147.81	16.29	838.38	15.89	<u>13.92</u>	<u>0.52</u>	33.38	0.53	18.53	24.11	24.1	23.93
NEXUS	281.76	116.65	282.34	117.24	18.59	6.67	33.01	7.54	<u>13.99</u>	19.52	18.71	16.3
MVTCAE	121.85	<u>5.34</u>	<u>54.57</u>	<u>3.16</u>	19.49	0.62	<u>13.65</u>	0.75	15.88	14.22	<u>14.02</u>	<u>13.96</u>
MMVAE+	97.19	2.80	128.56	114.3	22.37	1.21	21.74	15.2	16.12	17.31	17.92	17.56
MMVAE+(K=10)	85.98	1.83	70.72	62.43	21.10	1.38	8.52	7.22	14.58	14.33	14.34	14.32
MLD	7.98	1.7	4.54	1.84	3.18	0.83	2.07	<u>0.6</u>	2.39	2.31	2.33	2.29

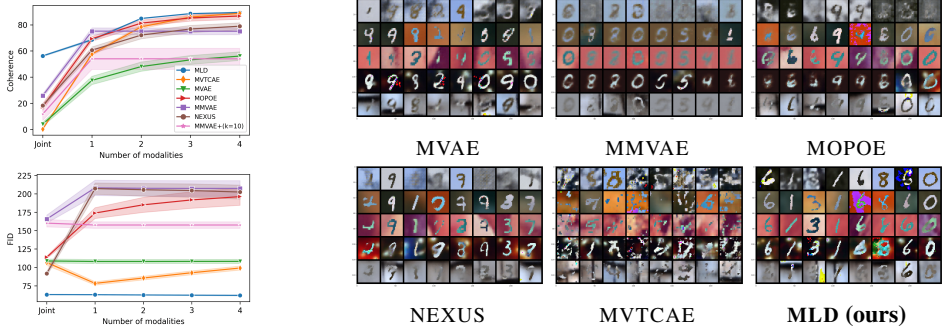


Figure 2: Results for **POLYMNIST** data-set. *Left*: a comparison of the generative coherence (% \uparrow) and quality in terms of FID (\downarrow) as a function of the number of inputs. We report the average performance following the leave-one-out strategy (see C). *Right*: are qualitative results for the joint generation of the 5 modalities.

results in E, where we show generation quality FID metric. Due to the unavailability of the labels in this data-set, coherence evaluation as with the previous data-sets is not possible. We then resort to CLIP-Score (CLIP-S) Hessel et al. (2021) an image-captioning metric, that, despite its limitations for the considered data-set Kim et al. (2022), shows that MLD outperforms competitors.

6 CONCLUSION AND LIMITATIONS

We have presented a new multi-modal generative model, Multimodal Latent Diffusion (MLD), to address the well-known coherence–quality tradeoff that is inherent in existing multi-modal VAE-based models. MLD uses a set of independently trained, uni-modal, deterministic autoencoders. Generative properties of our model stem from a masked diffusion process that operates on latent variables. We also developed a new multi-time training method to learn the conditional score network for multi-modal diffusion. An extensive experimental campaign on various real-life data-sets, provided compelling evidence on the effectiveness of MLD for multi-modal generative modeling. In all scenarios, including cases with loosely correlated modalities and high-resolution datasets, MLD consistently outperformed the alternatives from the state-of-the-art.




Figure 3: Qualitative results on **CUB** data-set. Caption used as condition to generate the bird images. **MLD*** denotes the version of our method using a powerful image autoencoder.

REFERENCES

- Alexander Alemi, Ben Poole, Ian Fischer, Joshua Dillon, Rif A Saurous, and Kevin Murphy. Fixing a broken elbow. In *International conference on machine learning*, pp. 159–168. PMLR, 2018.
- Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- Luigi Antelmi, Nicholas Ayache, Philippe Robert, and Marco Lorenzi. Sparse multi-channel variational autoencoder for the joint analysis of heterogeneous data. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 302–311. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/antelmi19a.html>.
- Shekoofeh Azizi, Simon Kornblith, Chitwan Saharia, Mohammad Norouzi, and David J. Fleet. Synthetic data from diffusion models improves imagenet classification, 2023.
- 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, 2023.
- Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models, 2023.
- Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T. Freeman, Michael Rubinstein, Yuanzhen Li, and Dilip Krishnan. Muse: Text-to-image generation via masked generative transformers, 2023.
- Matthieu Da Silva-Filarder, Andrea Ancora, Maurizio Filippone, and Pietro Michiardi. Multimodal variational autoencoders for sensor fusion and cross generation. In *2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pp. 1069–1076, 2021. doi: 10.1109/ICMLA52953.2021.00175.
- Imant Daunhawer, Thomas M. Sutter, Kieran Chin-Cheong, Emanuele Palumbo, and Julia E Vogt. On the limitations of multimodal VAEs. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=w-CPUXxRaj>.
- Adji B Dieng, Yoon Kim, Alexander M Rush, and David M Blei. Avoiding latent variable collapse with generative skip models. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2397–2405. PMLR, 2019.
- Emilien Dupont, Hyunjik Kim, S. M. Ali Eslami, Danilo Jimenez Rezende, and Dan Rosenbaum. From data to functa: Your data point is a function and you can treat it like one. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 5694–5725. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/dupont22a.html>.
- Giulio Franzese, Simone Rossi, Lixuan Yang, Alessandro Finamore, Dario Rossi, Maurizio Filippone, and Pietro Michiardi. How much is enough? a study on diffusion times in score-based generative models. *Entropy*, 25(4), 2023. ISSN 1099-4300. doi: 10.3390/e25040633. URL <https://www.mdpi.com/1099-4300/25/4/633>.
- Ruifei He, Shuyang Sun, Xin Yu, Chuhui Xue, Wenqing Zhang, Philip Torr, Song Bai, and XIAOJUAN QI. IS SYNTHETIC DATA FROM GENERATIVE MODELS READY FOR IMAGE RECOGNITION? In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=nUmCcZ5RKF>.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718*, 2021.

- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/8ald694707eb0fefe65871369074926d-Paper.pdf.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=rB6TpjAuSRy>.
- Minghui Hu, Chuanxia Zheng, Zuopeng Yang, Tat-Jen Cham, Heliang Zheng, Chaoyue Wang, Dacheng Tao, and Ponnuthurai N. Suganthan. Unified discrete diffusion for simultaneous vision-language generation. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=8JqINxA-2a>.
- Xun Huang, Arun Mallya, Ting-Chun Wang, and Ming-Yu Liu. Multimodal conditional image synthesis with product-of-experts gans. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVI*, pp. 91–109, Berlin, Heidelberg, 2022. Springer-Verlag. ISBN 978-3-031-19786-4. doi: 10.1007/978-3-031-19787-1_6. URL https://doi.org/10.1007/978-3-031-19787-1_6.
- 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.
- Adrian Javaloy, Maryam Meghdadi, and Isabel Valera. Mitigating modality collapse in multimodal VAEs via impartial optimization. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 9938–9964. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/javaloy22a.html>.
- Kevin Kilgour, Mauricio Zuluaga, Dominik Roblek, and Matthew Sharifi. Fréchet Audio Distance: A Reference-Free Metric for Evaluating Music Enhancement Algorithms. In *Proc. Interspeech 2019*, pp. 2350–2354, 2019. doi: 10.21437/Interspeech.2019-2219.
- Jin-Hwa Kim, Yunji Kim, Jiyoung Lee, Kang Min Yoo, and Sang-Woo Lee. Mutual information divergence: A unified metric for multimodal generative models. *arXiv preprint arXiv:2205.13445*, 2022.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6114>.
- Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5548–5557, 2019a. URL <https://api.semanticscholar.org/CorpusID:198967908>.
- Soochan Lee, Junsoo Ha, and Gunhee Kim. Harmonizing maximum likelihood with GANs for multimodal conditional generation. In *International Conference on Learning Representations*, 2019b. URL <https://openreview.net/forum?id=HJxyAjRcFX>.

- Gabriel Loaiza-Ganem, Brendan Leigh Ross, Jesse C Cresswell, and Anthony L. Caterini. Diagnosing and fixing manifold overfitting in deep generative models. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL <https://openreview.net/forum?id=0nEZCVshxS>.
- Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11461–11471, 2022.
- Sewon Min, Xinxin Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. Rethinking the role of demonstrations: What makes in-context learning work?, 2022.
- Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models, 2022.
- Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.
- Emanuele Palumbo, Imant Daunhawer, and Julia E Vogt. MMVAE+: Enhancing the generative quality of multimodal VAEs without compromises. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=sdQGxouELX>.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/radford21a.html>.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents, 2022.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, June 2022.
- Ludan Ruan, Yiyang Ma, Huan Yang, Huiguo He, Bei Liu, Jianlong Fu, Nicholas Jing Yuan, Qin Jin, and Baining Guo. Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation, 2023.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo-Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=08Yk-n5l2Al>.
- Mert Bulent Sariyildiz, Karteek Alahari, Diane Larlus, and Yannis Kalantidis. Fake it till you make it: Learning transferable representations from synthetic imagenet clones, 2023.
- Yuge Shi, Siddharth N, Brooks Paige, and Philip Torr. Variational mixture-of-experts autoencoders for multi-modal deep generative models. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/0ae775a8cb3b499ad1fca944e6f5c836-Paper.pdf.
- Yuge Shi, Brooks Paige, Philip Torr, and Siddharth N. Relating by contrasting: A data-efficient framework for multimodal generative models. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=vhKe9UFbrJo>.

- Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, Devi Parikh, Sonal Gupta, and Yaniv Taigman. Make-a-video: Text-to-video generation without text-video data, 2022.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/sohl-dickstein15.html>.
- Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 12438–12448. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/92c3b916311a5517d9290576e3ea37ad-Paper.pdf.
- Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in Neural Information Processing Systems*, 34:1415–1428, 2021a.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021b. URL <https://openreview.net/forum?id=PXTIG12RRHS>.
- Thomas M. Sutter, Imant Daunhawer, and Julia E. Vogt. Multimodal generative learning utilizing jensen-shannon-divergence. *CoRR*, abs/2006.08242, 2020. URL <https://arxiv.org/abs/2006.08242>.
- Thomas M. Sutter, Imant Daunhawer, and Julia E Vogt. Generalized multimodal ELBO. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=5Y21V0RDBV>.
- Zineng Tang, Ziyi Yang, Chenguang Zhu, Michael Zeng, and Mohit Bansal. Any-to-any generation via composable diffusion, 2023.
- Ming Tao, Hao Tang, Fei Wu, Xiao-Yuan Jing, Bing-Kun Bao, and Changsheng Xu. Df-gan: A simple and effective baseline for text-to-image synthesis, 2022.
- Ba-Hien Tran, Simone Rossi, Dimitrios Milios, Pietro Michiardi, Edwin V Bonilla, and Maurizio Filippone. Model selection for bayesian autoencoders. *Advances in Neural Information Processing Systems*, 34:19730–19742, 2021.
- Luan Tran, Xiaoming Liu, Jiayu Zhou, and Rong Jin. Missing modalities imputation via cascaded residual autoencoder. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=P9TYG0j-wtG>.
- Miguel Vasco, Hang Yin, Francisco S. Melo, and Ana Paiva. Leveraging hierarchy in multimodal generative models for effective cross-modality inference. *Neural Networks*, 146:238–255, 2 2022. ISSN 18792782. doi: 10.1016/j.neunet.2021.11.019.
- Ashvala Vinay and Alexander Lerch. Evaluating generative audio systems and their metrics. *arXiv preprint arXiv:2209.00130*, 2022.
- Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- Daniel Wesego and Amirmohammad Rooshenas. Score-based multimodal autoencoders, 2023.

Fuxiang Wu, Liu Liu, Fusheng Hao, Fengxiang He, and Jun Cheng. Text-to-image synthesis based on object-guided joint-decoding transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 18113–18122, June 2022.

Mike Wu and Noah Goodman. Multimodal generative models for scalable weakly-supervised learning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/1102a326d5f7c9e04fc3c89d0ede88c9-Paper.pdf.

Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. Next-gpt: Any-to-any multimodal llm, 2023.

Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. An explanation of in-context learning as implicit bayesian inference. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=RdJVFCHjUMI>.

Yue Zhang, Chengtao Peng, Qiuli Wang, Dan Song, Kaiyan Li, and S. Kevin Zhou. Unified multi-modal image synthesis for missing modality imputation, 2023.

A APPENDIX

MULTI-MODAL LATENT DIFFUSION — SUPPLEMENTARY MATERIAL

A DIFFUSION IN THE MULTIMODAL LATENT SPACE

In this section, we provide additional technical details of MLD. We first discuss a naive approach based on *In-painting* which uses only unconditional score network for both joint and conditional generation. We also discuss alternative training scheme based on a work from the caption-text translation literature Bao et al. (2023). Finally, we provide extra technical details for the score network architecture and sampling technique.

A.1 MODALITIES AUTO-ENCODERS

Each deterministic autoencoders used in the first stage of MLD uses a vector latent space with no size constraints. Instead, VAE-based models, generally require the latent space of each individual VAE to be exactly of the same size, to allow the definition of a joint latent space.

In our approach, before concatenation, the modality-specific latent spaces are *normalized* by element-wise mean and standard deviation. In practice, we use the statistics retrieved from the first training batch, which we found sufficient to gain sufficient statistical confidence. This operation allows the harmonization of different modality-specific latent spaces and, therefore, facilitate the learning of a joint score network.

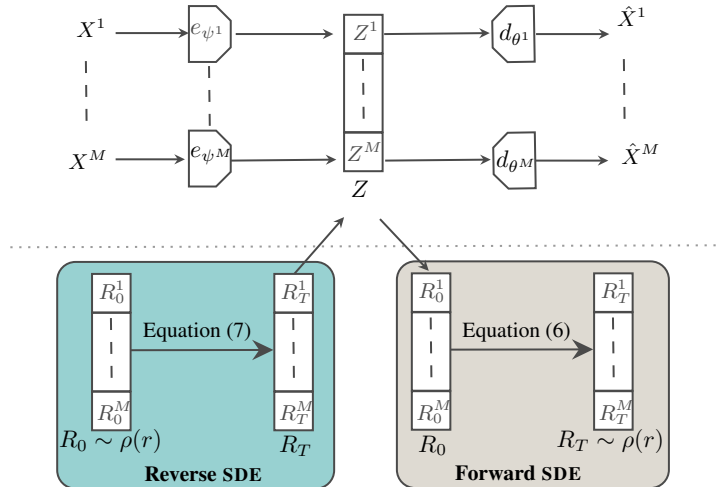


Figure 4: Multi-modal Latent Diffusion. Two-stage model involving: **Top**: deterministic, modality-specific encoder/decoders, **Bottom**: score-based diffusion model on the concatenated latent spaces.

A.2 MULTI-MODAL DIFFUSION SDE

In Section 3, we presented our multi-modal latent diffusion process allowing multi-modal joint and conditional generation. The role of the SDE is to gradually add noise to the data, perturbing its structure until attaining a noise distribution. In this work, we consider Variance preserving SDE (VPSDE) Song et al. (2021b). In this framework we have : $\rho(r) \sim \mathcal{N}(0; I)$, $\alpha(t) = -\frac{1}{2}\beta(t)$ and $g(t) = \sqrt{\beta(t)}$, where $\beta(t) = \beta_{min} + t(\beta_{max} - \beta_{min})$. Following (Ho et al., 2020; Song et al., 2021b), we set $\beta_{min} = 0.1$ and $\beta_{max} = 20$. With this configuration and by substitution of Equation (3), we obtain the following forward SDE:

$$dR_t = -\frac{1}{2}\beta(t)R_t dt + \sqrt{\beta(t)}dW_t, \quad t \in [0, T]. \quad (9)$$

The corresponding perturbation kernel is given by :

$$q(r|z, t) = \mathcal{N}(r; e^{-\frac{1}{4}t^2(\beta_{max}-\beta_{min})-\frac{1}{2}t\beta_{min}}z, (1 - e^{-\frac{1}{2}t^2(\beta_{max}-\beta_{min})-t\beta_{min}})\mathbf{I}). \quad (10)$$

The marginal score $\nabla \log q(R_t, t)$ is approximated by a score network $s_\chi(R_t, t)$ whose parameters χ can be optimized by minimizing the ELBO in Equation (5), where we found that using the same re-scaling as in Song et al. (2021b) is more stable.

The reverse process is described by a different SDE (Equation (4)). When using a variance-preserving SDE, Equation (4) specializes in:

$$dR_t = \left[\frac{1}{2}\beta(T-t)R_t + \beta(T-t)\nabla \log q(R_t, T-t) \right] dt + \sqrt{\beta(T-t)}dW_t, \quad (11)$$

With $R_0 \sim \rho(r)$ as initial condition and time t flows from $t = 0$ to $t = T$.

Once the parametric score network is optimized, through the simulation of Equation (11), sampling $R_T \sim q_\psi(r)$ is possible allowing **joint generation**. A numerical SDE solver can be used to sample R_T which can be fed to the modality specific decoders to jointly sample a set of $\hat{X} = \{d_\theta^i(R_T)\}_{i=0}^M$. As explained in Section 4.2, the use of the unconditional score network $s_\chi(R_t, t)$ allows **conditional generation** through the approximation described in Song et al. (2021b).

As described in Algorithm 1, one can generate a set of modalities A_1 conditioned on the available set of modalities A_2 . First, the available modalities are encoded into their respective latent space z^{A_2} , the initial missing part is sampled from the stationary distribution $R_0^{A_1} \sim \rho(r^{A_1})$, using an SDE solver (e.g. Euler-Maruyama), the reverse diffusion SDE (in Equation (11)) is discretized using a finite time steps $\Delta t = T/N$, starting from $t = 0$ and iterating until $t \approx T$. At each iteration, the available portion of the latent space is diffused and brought to the same noise level as $R_t^{A_1}$ allowing the use of the unconditional score network. Lastly, the reverse diffusion update is done. This process is repeated until arriving at $t \approx T$ and obtaining $R_T^{A_1} = \hat{Z}^{A_1}$ which can be decoded to recover \hat{x}^{A_1} . Note that the joint generation can be seen as a special case of Algorithm 1 with $A_2 = \emptyset$. We name this first approach Multi-modal Latent Diffusion with In-painting (MLD IN-PAINT) and provide extensive comparison with our method MLD in Appendix B.

Algorithm 1: MLD IN-PAINT conditional generation

```

Data:  $x^{A_2} = \{x^i\}_{i \in A_2}$ 
 $z^{A_2} \leftarrow \{e_{\phi_i}(x^i)\}_{i \in A_2}$  // Encode the available modalities  $X$  into their
latent space
 $A_1 \leftarrow \{1, \dots, M\} \setminus A_2$  // The set of modalities to generate
 $R_0 \leftarrow \mathcal{C}(R_0^{A_1}, z^{A_2}), \quad R_0^{A_1} \sim \rho(r^{A_1})$  // Compose the initial state
 $R \leftarrow R_0$ 
 $\Delta t \leftarrow T/N$ 
for  $n = 0$  to  $N - 1$  do
     $t' \leftarrow T - n \Delta t$ 
     $\bar{R} \sim q(r|R_0, t')$  // Diffuse the available portion of the latent
space (eq. (10))
     $R \leftarrow m(A_1) \odot R + (1 - m(A_1)) \odot \bar{R}$ 
     $\epsilon \sim \mathcal{N}(0; I)$  if  $n < (N - 1)$  else  $\epsilon = 0$ 
     $\Delta R \leftarrow \Delta t \left[ \frac{1}{2}\beta(t')R + \beta(t')s_\chi(R, t') \right] + \sqrt{\beta(t')}\Delta t \epsilon$ 
     $R \leftarrow R + \Delta R$  // The Euler-Maruyama update step
end
 $\hat{z}^{A_1} \leftarrow R^{A_1}$ 
Return  $\hat{X}^{A_1} = \{d_\theta^i(\hat{z}^i)\}_{i \in A_1}$ 

```

As discussed in Section 4.2, the approximation enabling the in-painting approach can be efficient in several domains but its generalization to the multi-modal latent space scenario is not trivial. We argue that this is due to the heterogeneity of modalities which induce different latent spaces characteristics. For different modality-specific latent spaces, the loss of information ratio can vary through the diffusion process. We verify this hypothesis through the following experiment.

Latent space robustness against diffusion perturbation: We analyse the effect of the forward diffusion perturbation on the latent space through time. We encode the modalities using their respective encoders to obtain their latent space $Z = [e_{\psi^1}(X^1) \dots e_{\psi^M}(X^M)]$. Given a time $t \in [0, T]$, we diffuse the different latent spaces by applying Equation (10) to get $R_t \sim q(r|z, t)$ with R_t being the perturbed version of the latent space at time t . We feed the modality specific decoders with the perturbed latent space $\hat{X}_t = \{d_{\theta}^i(R_t^i)\}_{i=1}^M$, \hat{X}_t being the output modalities generated using the perturbed latent space. To evaluate the information loss induced by the diffusion process on the different modalities, we assess the coherence preservation in the reconstructed modalities \hat{X}_t by computing the coherence (in %) as done in Section 5.

We expect to obtain high coherence results for $t \approx 0$, when compared to $t \approx T$, the information in the latent space being more preserved at the beginning of the diffusion process than at the last phase of the forward SDE where all dependencies on initial conditions vanish. Figure 5 shows the coherence as a function of the diffusion time $t \in [0, 1]$ for different modalities across multiple data-sets. We observe that within the same data-set, some modalities stand out with a specific level of robustness (using as a proxy the coherence level) against the diffusion perturbation in comparison with the remaining modalities from the same data-set. For instance, we remark that SVHN is less robust than MNIST which should manifest in an under-performance of SVHN to MNIST conditional generation. An intuition that we verify in Appendix B.




Figure 5: The coherence as a function of the diffusion process time for three datasets. The diffusion perturbation is applied on the modalities latent space after an element wise normalization.

A.3 MULTI-TIME MASKED MULTI-MODAL SDE

To learn the score network capable of both conditional and joint generation, we proposed in Section 4 a multi-time masked diffusion process.

Algorithm 2 presents a pseudo-code for the multi time masked training. The masked diffusion process is applied following a randomization with probability d . First, a subset of modalities A_2 is selected randomly to be the conditioning modalities and A_1 the remaining set of modalities to be the diffused modalities. The time t is sampled uniformly from $[0, T]$ and the portion of the latent space corresponding to the subset A_1 is diffused accordingly. Using the masking as shown in Algorithm 2, the portion of the latent space corresponding to the subset A_2 is not diffused and forced to be equal to $R_0^{A_2} = z^{A_2}$. The multi-time vector τ is constructed. Lastly, the score network is optimized by minimizing a masked loss corresponding to the diffused part of the latent space. With probability $(1 - d)$, all the modalities are diffused at the same time and $A_2 = \emptyset$. In order to calibrate the loss, given that the randomization of A_1 and A_2 can result in diffusing different sizes of the latent space, we re-weight the loss according to the cardinality of the diffused and freed portions of the latent space:

$$\Omega(A_1, A_2) = 1 + \frac{\dim(A_2)}{\dim(A_1)} \quad (12)$$

Where $\dim(\cdot)$ is the sum of each latent space cardinality of a given subset of modalities with $\dim(\emptyset) = 0$.

Algorithm 2: MLD Masked Multi-time diffusion training step

Data: $X = \{x^i\}_{i=1}^M$
Param: d
 $Z \leftarrow \{e_{\phi_i}(x^i)\}_{i=0}^M$ // Encode the modalities X into their latent space
 $A_2 \sim \nu$ // ν depends on the parameter d
 $A_1 \leftarrow \{1, \dots, M\} \setminus A_2$
 $t \sim \mathcal{U}[0, T]$
 $R \sim q(r|Z, t)$ // Diffuse the available portion of the latent space (Equation (10))
 $R \leftarrow m(A_1) \odot R + (1 - m(A_1)) \odot Z$ // Masked diffusion
 $\tau(A_1, t) \leftarrow [\mathbb{1}(1 \in A_1)t, \dots, \mathbb{1}(M \in A_1)t]$ // Construct the multi time vector
Return $\nabla_{\mathcal{X}} \left\{ \Omega(A_1, A_2) \quad \left\| m(A_1) \odot [s_{\mathcal{X}}(R, \tau(A_1, t)) - \nabla \log q(R, t|z^{A_2})] \right\|_2^2 \right\}$

The optimized score network can approximate both the conditional and unconditional true score:

$$s_{\mathcal{X}}(R_t, \tau(A_1, t)) \sim \nabla \log q(R_t, t | z^{A_2}). \quad (13)$$

The joint generation is a special case of the latter with $A_2 = \emptyset$:

$$s_{\mathcal{X}}(R_t, \tau(A_1, t)) \sim \nabla \log q(R_t, t) \quad , A_1 = \{1, \dots, M\} \quad (14)$$

Algorithm 3 describes the reverse conditional generation pseudo-code. It’s pertinent to compare this algorithm with Algorithm 1. The main difference resides in the use of the multi-time score network, enabling conditional generation with the multi-time vector playing the role of time information and conditioning signal. On the other hand, in Algorithm 1, we don’t have a conditional score network, therefore we resort to the approximation from Section 4.2, and use the unconditional score.

Algorithm 3: MLD conditional generation.

Data: $x^{A_2} \leftarrow \{x^i\}_{i \in A_2}$
 $z^{A_2} \leftarrow \{e_{\phi_i}(x^i)\}_{i \in A_2}$ // Encode the available modalities X into their latent space
 $A_1 \leftarrow \{1, \dots, M\} \setminus A_2$ // The set of modalities to be generated
 $R_0 \leftarrow \mathcal{C}(R_0^{A_1}, z^{A_2}), \quad R_0^{A_1} \sim \rho(r^{A_1})$ // Compose the initial latent space
 $R \leftarrow R_0$
 $\Delta t \leftarrow T/N$
for $n = 0$ **to** $N - 1$ **do**
 $t' \leftarrow T - n \Delta t$
 $\tau(A_1, t') \leftarrow [\mathbb{1}(1 \in A_1)t', \dots, \mathbb{1}(M \in A_1)t']$ // Construct the multi-time vector
 $\epsilon \sim \mathcal{N}(0; I)$ **if** $n < N$ **else** $\epsilon = 0$
 $\Delta R \leftarrow \Delta t \left[\frac{1}{2} \beta(t') R + \beta(t') s_{\mathcal{X}}(R, \tau(A_1, t')) \right] + \sqrt{\beta(t') \Delta t} \epsilon$
 $R \leftarrow R + \Delta R$ // The Euler-Maruyama update step
 $R \leftarrow m(A_1) \odot R + (1 - m(A_1)) \odot R_0$ // Update the portion corresponding to the unavailable modalities
end
 $\hat{z}^{A_1} = R^{A_1}$
Return $\hat{X}^{A_1} = \{d_{\theta}^i(\hat{z}^i)\}_{i \in A_1}$

A.4 UNI-DIFFUSER TRAINING

The work presented in Bao et al. (2023) is specialized for an image-caption application. The approach is based on a multi-modal diffusion model applied to a unified latent embedding, obtained via

pre-trained autoencoders, and incorporating pre-trained models (CLIP Radford et al. (2021) and GPT-2 Radford et al. (2019)). The unified latent space is composed of an image embedding, a CLIP image embedding and a CLIP text embedding. Note that the CLIP model is pre-trained on a pairs of multi-modal data (image-text), which is expected to enhance the generative performance. Since it is not trivial to have a jointly trained encoder similar to CLIP for any type of modality, the evaluation of this model on different modalities across different data-set (e.g. including audio) is not an easy task.

To compare to this work, we adapt the training scheme presented in Bao et al. (2023) to our MLD method. Instead of applying a masked multi-modal SDE for training the score network, every portion of the latent space is diffused according to a different time $t^i \sim \mathcal{U}(0, 1)$ and, therefore, the multi-time vector fed to the score network is $\tau(t) = [t^0 \sim \mathcal{U}(0, 1), \dots, t^M \sim \mathcal{U}(0, 1)]$. For fairness, we use the same score network and reverse process sampler as for our MLD version with multi-time training, and call this variant Multi-modal Latent Diffusion UniDiffuser (MLD UNI).

A.5 INTUITIVE SUMMARY: HOW DOES MLD CAPTURE MODALITY INTERACTIONS?

MLD treats the latent spaces of each modality as variables that evolve differently through the diffusion process according to a multi-time vector. The masked multi-time training enables the model to learn the score of all the combination of conditionally diffused modalities, using the frozen modalities as the conditioning signal, through a randomized scheme. By learning the score function of the diffused modalities at different time steps, the score model captures the correlation between the modalities. At test time, the diffusion time of each modality is chosen to modulate its influence on the generation, as follows.

For joint generation the model uses the unconditional score which corresponds to using the same diffusion time for all modalities. Thus, all the modalities influence each other equally. This ensures that modality interaction information is faithful to the one characterizing the observed data distribution.

The model can also generate modalities conditionally by using the conditional score, by freezing the conditioning modalities during the reverse process. The frozen state is similar to the final state of the reverse process where information is not perturbed, thus the influence of the conditioning modalities is maximal. Subsequently, the generated modalities reflect the necessary information from the conditioning modalities and achieve the desired correlation.

A.6 TECHNICAL DETAILS

Sampling schedule: We use the sampling schedule proposed in Lugmayr et al. (2022), which has shown to improve the coherence of the conditional and joint generation. We use the best parameters suggested by the authors: $N = 250$ time-steps, applied $r = 10$ re-sampling times with jump size $j = 10$. For readability in algorithm 1 and algorithm 3, we present pseudo code with a linear sampling schedule which can be easily adapted to any other schedule.

Training the score network: Inspired by the architecture from (Dupont et al., 2022), we use simple Residual MLP blocks with skip connections as our score network (see Figure 6). We fix the **width** and **number of blocks** proportionally to the number of the modalities and the latent space size. As in Song & Ermon (2020), we use Exponential moving average (EMA) of model parameters with a momentum parameter $m = 0.999$.

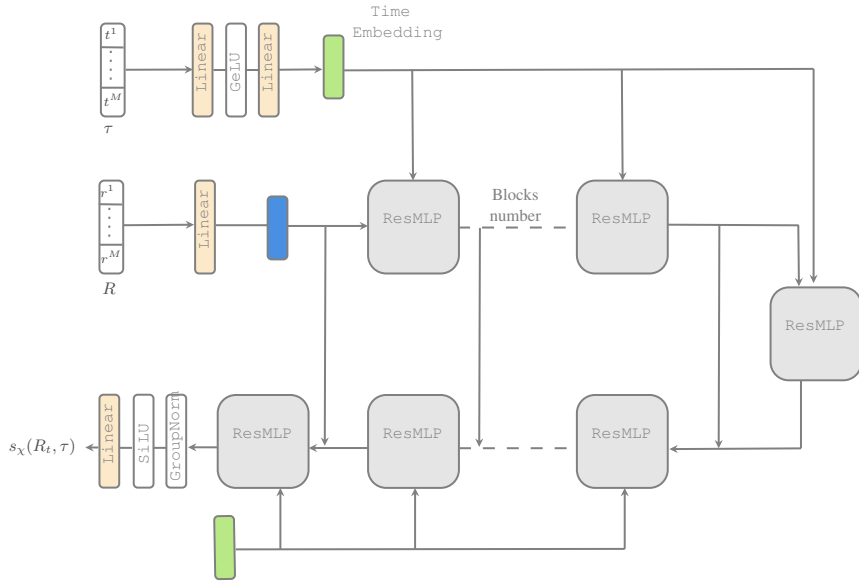


Figure 6: Score network s_x architecture used in our MLD implementation. Residual MLP block architecture is shown in Figure 7.




Figure 7: Architecture of ResMLP block.

B MLD ABLATIONS STUDY

In this section, we compare MLD with two variants presented in Appendix A : MLD IN-PAINT, a naive approach without our proposed *multi-time masked* SDE, MLD UNI a variant of our method using the same training scheme of Bao et al. (2023). We also analyse the effect of the d randomization parameter on MLD performance through ablations study.

B.1 MLD AND ITS VARIANTS

Table 4 summarizes the different approaches adopted in each variant. All the considered models share the same deterministic autoencoders trained during the first stage.

For fairness, our evaluation was done using the same configuration and code basis of MLD. This includes: the autoencoder architectures and latent space size (similar to Section 5), the same score network (Figure 6) is used across experiments, with MLD IN-PAINT using the same architecture with one time dimension instead of the multi-time vector. In all the variants, the joint and conditional generation are conducted using the same reverse sampling schedule described in Appendix A.6.

Table 4: MLD and its variants ablation study

Model	Multi-time diffusion	Training	Conditional and joint generation
MLD IN-PAINT	x	Equation (5)	Algorithm 1
MLD UNI	✓	Bao et al. (2023)	Algorithm 3
MLD	✓	Algorithm 2	Algorithm 3

Results In some cases, the MLD variants can match the joint generation performance of MLD but, overall, they are less efficient and have noticeable weaknesses: MLD IN-PAINT under-performs in conditional generation when considering relatively complex modalities, MLD UNI is not able to leverage the presence of multiple modalities to improve cross generation, especially for data-sets with a large number of modalities. On the other hand, MLD is able to overcome all these limitations.

MNIST-SVHN. In Table 5, MLD achieves the best results and dominates cross generation performance. We observe that MLD IN-PAINT lacks coherence for SVHN to MNIST conditional generation, a results we expected by analysing the experiment in Figure 5. MLD UNI, despite the use of a multi-time diffusion process, under-performs our method, which indicates the effectiveness of our masked diffusion process in learning the conditional score network. Since all the models use the same deterministic autoencoders, the observed generative quality performance are relatively similar (See Figure 8 for qualitative results).

Table 5: Generation Coherence and Quality for MNIST-SVHN (M is for MNIST and S for SVHN). The generation quality is measured in terms of FMD for MNIST and FID for SVHN.

Models	Coherence (% \uparrow)			Quality (\downarrow)			
	Joint	M \rightarrow S	S \rightarrow M	Joint(M)	Joint(S)	M \rightarrow S	S \rightarrow M
MLD-Inpaint	85.53 \pm 0.22	81.76 \pm 0.23	63.28 \pm 1.16	3.85 \pm 0.02	60.86 \pm 1.27	59.86 \pm 1.18	3.55 \pm 0.11
MLD-Uni	82.19 \pm 0.97	79.31 \pm 1.21	72.78 \pm 1.81	4.1 \pm 0.17	57.41 \pm 1.43	57.84 \pm 1.57	4.84 \pm 0.28
MLD	85.22 \pm 0.5	83.79 \pm 0.62	79.13 \pm 0.38	3.93 \pm 0.12	56.36 \pm 1.63	57.2 \pm 1.47	3.67 \pm 0.14



Figure 8: Qualitative results for **MNIST-SVHN**. For each model we report: MNIST to SVHN conditional generation in the left, SVHN to MNIST conditional generation in the right.

MHD. Table 6 shows the performance results for the MHD data-set in terms of generative coherence. MLD achieves the best joint generation coherence and, along with MLD UNI, they dominate the cross generation coherence. MLD IN-PAINT shows a lack of coherence when conditioning on the sound modality alone, a predictable result since this is a more difficult configuration, as the sound modality is loosely correlated to other modalities. We also observe that MLD IN-PAINT performs worse than the two other alternatives when conditioned on the trajectory modality, which is the smallest modality in terms of latent size. This indicates another limitation of the naive approach regarding coherent generation when handling different latent spaces sizes, a weakness our method MLD overcomes. Table 7 presents the qualitative generative performance which are homogeneous across the variants with MLD, achieving either the best or second best performance.

Table 6: Generation Coherence ($\% \uparrow$) for MHD (Higher is better). Line above refers to the generated modality while the observed modalities subset are presented below.

Models	Joint	I (Image)			T (Trajectory)			S (Sound)		
		T	S	T,S	I	S	I,S	I	T	I,T
MLD-Inpaint	96.88 \pm 0.35	63.9 \pm 1.7	56.52 \pm 1.89	95.83 \pm 0.48	99.58 \pm 0.1	56.51 \pm 1.89	99.89 \pm 0.04	95.81 \pm 0.25	56.51 \pm 1.89	96.38 \pm 0.35
MLD-Uni	97.69 \pm 0.26	99.91 \pm 0.04	89.87 \pm 0.38	99.92 \pm 0.04	99.68 \pm 0.1	89.78 \pm 0.45	99.38 \pm 0.31	97.54 \pm 0.2	97.69 \pm 0.41	97.79 \pm 0.41
MLD	98.34 \pm 0.22	99.45 \pm 0.09	88.91 \pm 0.54	99.88 \pm 0.04	99.58 \pm 0.03	88.92 \pm 0.53	99.91 \pm 0.02	97.63 \pm 0.14	97.7 \pm 0.34	98.01 \pm 0.21

Table 7: Generation quality for MHD. The metrics reported are FMD for Image and Trajectory modalities and FAD for the sound modalities (Lower is better).

Models	I (Image)				T (Trajectory)			S (Sound)				
	Joint	T	S	T,S	Joint	I	S	I,S	Joint	I	T	I,T
MLD-Inpaint	5.35 \pm 1.35	6.23 \pm 1.13	4.76 \pm 0.68	3.53 \pm 0.36	1.59 \pm 0.12	0.6 \pm 0.05	1.81 \pm 0.13	0.54 \pm 0.06	2.41 \pm 0.07	2.5 \pm 0.04	2.52 \pm 0.02	2.49 \pm 0.05
MLD-Uni	7.91 \pm 2.2	1.65 \pm 0.33	6.29 \pm 1.38	3.06 \pm 0.54	2.53 \pm 0.5	1.18 \pm 0.26	3.18 \pm 0.77	2.84 \pm 1.14	2.11 \pm 0.08	2.25 \pm 0.05	2.1 \pm 0.0	2.15 \pm 0.01
MLD	7.98 \pm 1.41	1.7 \pm 0.14	4.54 \pm 0.45	1.84 \pm 0.27	3.18 \pm 0.18	0.83 \pm 0.03	2.07 \pm 0.26	0.6 \pm 0.05	2.39 \pm 0.1	2.31 \pm 0.07	2.33 \pm 0.11	2.29 \pm 0.06

POLYMNIST. In Figure 9, we remark the superiority of MLD in both generative coherence and quality. MLD-Uni is not able to leverage the presence of a large number of modalities in conditional generation coherence. Interestingly, an increase in the number of input modalities impacts negatively the performance of MLD UNI.

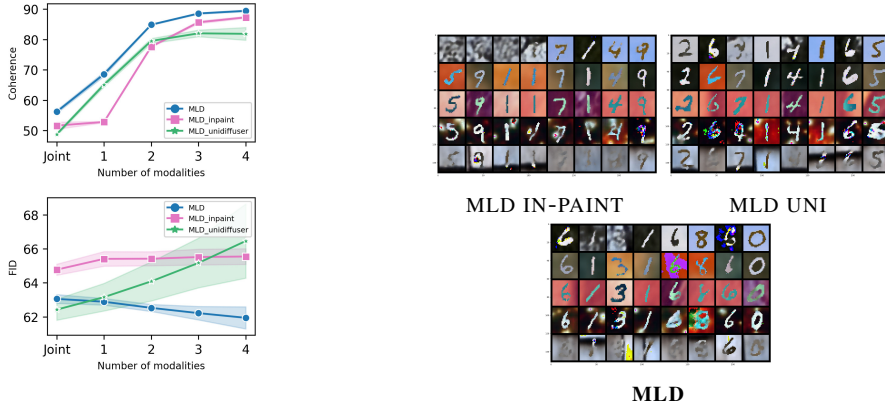


Figure 9: Results for **POLYMNIST** data-set. *Left*: a comparison of the generative coherence ($\% \uparrow$) and quality in terms of FID (\downarrow) as a function of the number of modality input. We report the average performance following the leave-one-out strategy (see Appendix C). *Right*: are qualitative results for the joint generation of the 5 modalities.

CUB. Figure 10 shows qualitative results for caption to image conditional generation. All the variants are based on the same first stage autoencoders, and the generative performance in terms of quality are comparable.

B.2 RANDOMIZATION d -ABLATIONS STUDY

The d parameter controls the randomization of the *multi-time masked diffusion process* during training in Algorithm 2. With probability d , the concatenated latent space corresponding to all the modalities is diffused at the same time. With probability $(1 - d)$, a portion of the latent space corresponding to a random subset of the modalities is not diffused and frozen during the training step. To study the parameter d and its effect on the performance of our MLD model, we use $d \in \{0.1, \dots, 0.9\}$. Figure 11 shows the d -ablentions study results on the **MNIST-SVHN** dataset. We report the performance results averaged over 5 independent seeds as a function of the probability $(1 - d)$: **Left**: the conditional

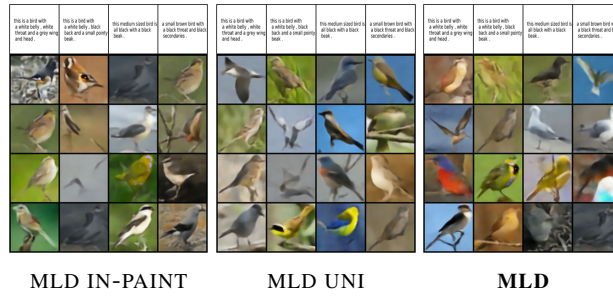
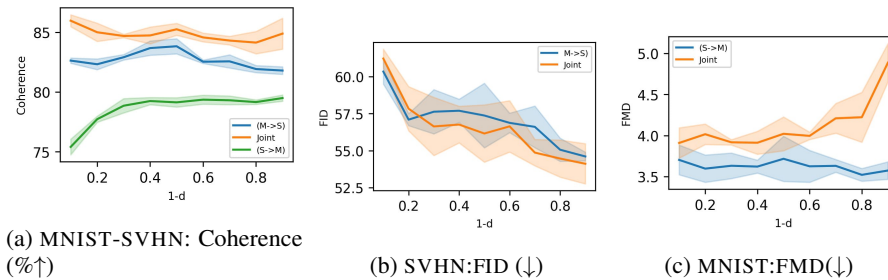


Figure 10: Qualitative results on CUB data-set. Caption used as condition to generate the bird images.

and joint coherence for **MNIST-SVHN** dataset. **Middle**: the quality performance in terms of FID for SVHN generation. **Right**: the quality performance in terms of FMD for MNIST generation.

We observe that higher value for $1 - d$ thus greater probability of applying the *multi-time masked diffusion*, improves the SVHN to MNIST conditional generation coherence. This confirms that the masked multi-time training enables better conditional generation. Overall, on the **MNIST-SVHN** dataset, MLD shows weak sensibility to the d parameter whenever the value of $d \in [0.2, 0.7]$.

Figure 11: The randomization parameter d ablations study on **MNIST-SVHN**.

C DATASETS AND EVALUATION PROTOCOL

C.1 DATASETS DESCRIPTION

MNIST-SVHN Shi et al. (2019) is constructed using pairs of MNIST and SVHN, sharing the same digit class (See Figure 12a). Each instance of a digit class (in either dataset) is randomly paired with 20 instances of the same digit class from the other data-set. SVHN modality samples are obtained from house numbers in Google Street View images, characterized by a variety of colors, shapes and angles. A high number of SVHN samples are noisy and can contain different digits within the same sample due to the imperfect cropping of the original full house number image. One challenge of this data-set for multi-modal generative models is to learn to extract digit number and reconstruct a coherent MNIST modality.

MHD Vasco et al. (2022) is composed of 3 modalities: synthetically generated images and motion trajectories of handwritten digits associated with their speech sounds. The images are gray scale $1 \times 28 \times 28$ and the handwriting trajectory are represented by a 1×200 vector. The spoken digits sound is 1s long audio processed as Mel-Spectrograms constructed with a hopping window of 512 ms with 128 Mel Bins resulting in a $1 \times 128 \times 32$ representation. This benchmark is the closest to a real world multi-modal sensors scenario because of the presence of three completely different modalities, the audio modality representing a complex data type. Therefore, similar to SVHN, the conditional generation of sound to coherent images or trajectories represents a challenging use case.

POLYMNIST Sutter et al. (2021) is an extended version of the MNIST data-set to 5 modalities. Each modality is constructed using a randomly set of MNIST digits with an overlay over a random crop from a modality specific, 3 channel image background. This synthetic generated data-set allows the evaluation of the scalability of multi-modal generative models to large number of modalities. Although this data-set is composed of only images, the different modality-specific background having different textures, results in different levels of difficulty. In Figure 12c, the digits numbers are more difficult to distinguish in modality 1 and 5 than in the remaining modalities.

CUB Shi et al. (2019) is comprised of bird images and their associated text captions. The work in Shi et al. (2019) used a simplified version based on pre-computed ResNet-features. We follow Daunhawer et al. (2022) and conduct all our experiments on the real image data instead. Each image from the 11,788 photos of birds from Caltech-Birds Wah et al. (2011) are resized to $3 \times 64 \times 64$ image size and coupled with 10 textual descriptions of the respective bird (See Figure 12d).

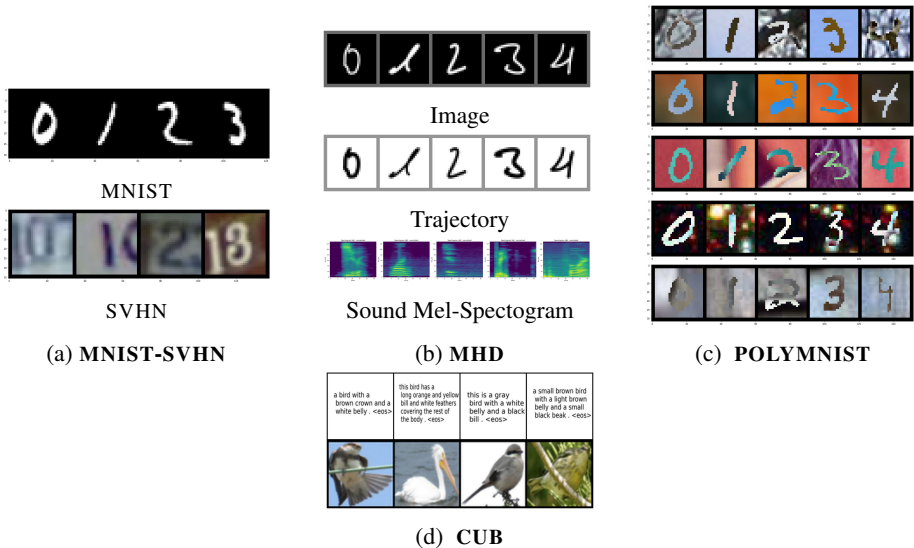


Figure 12: Illustrative example of the Datasets used for the evaluation

C.2 EVALUATION METRICS

Multimodal generative models are evaluated in terms of generative coherence and quality.

C.2.1 GENERATION COHERENCE

We measure *coherence* by verifying that generated data (both for joint and conditional generations) share the same information across modalities. Following Shi et al. (2019); Sutter et al. (2021); Hwang et al. (2021); Vasco et al. (2022); Daunhawer et al. (2022), we consider the class label of the modalities as the shared information and use pre-trained classifiers to extract the label information from the generated samples and compare it across modalities.

For **MNIST-SVHN**, **MHD** and **POLYMNIST**, the shared semantic information is the digit class number. Single modality classifiers are trained to classify the digit number of a given modality sample. To compute the conditional generation of modality m with a subset of modalities A , we feed the modality specific pre-trained classifier C_m with the conditional generated sample \hat{X}^m . The predicted label class is compared to the ground truth label y_{X^A} which is the label of modalities of the subset X^A . For N samples, the matching rate average establishes the coherence. For all the experiments, N is equal to the length of the test-set.

$$Coherence(\hat{X}^m|X^A) = \frac{1}{N} \sum_1^N \mathbb{1}_{\{C_m(\hat{X}^m)=y_{X^A}\}} \quad (15)$$

The **joint generation coherence** is measured by feeding the generated samples of each modality to their specific trained classifier. The rate with which all classifiers output the same predicted digit label for N generations is considered as the joint generation coherence.

The leave one out coherence: is the conditional generation coherence using all the possible subsets excluding the generated modality: $Coherence(\hat{X}^m|X^A)$ with $A = \{1, \dots, M\} \setminus m$. Due to the large number of modalities in **POLYMNIST**, similar to Sutter et al. (2021); Hwang et al. (2021); Daunhawer et al. (2022) we compute the average **leave one out coherence** conditional coherence as a function of the input modalities subset size.

Due to the unavailability of labels in the **CUB** data-set, we use CLIP-S Hessel et al. (2021) a state of the art metric for image captioning evaluation.

C.2.2 GENERATION QUALITY

For each modality, we consider the following metrics:

- **RGB Images:** FID Heusel et al. (2017) is the state-of-the-art standard metric to evaluate image generation quality of generative models.
- **Audio:** FAD Kilgour et al. (2019), is state-of-the-art standard metric in the evaluation of audio generation. FAD performs well in terms of robustness against noise and is consistent with human judgments Vinay & Lerch (2022). Similar to FID, a Fréchet distance is computed but VGGish (audio classifier model) embeddings are used instead.
- **Other modalities** For other modality types, we derive FMD (Fréchet Modality Distance), a similar metric to FID and FAD. We compute the **Fréchet distance** between the statistics retrieved from the activations of the modality specific pre-trained classifiers used for coherence evaluation. FMD is used to evaluate the generative quality of MNIST modality in **MNIST-SVHN** and image and trajectory modalities in **MHD** data-set.

For conditional generation, we compute the quality metric (FID,FAD or FMD) using the conditionally generated modality and the real data. For joint generation, we use the randomly generated modality and randomly selected same number of samples from the real data.

For **CUB**, we use 10000 samples to evaluate the generation quality in terms of FID. In the remaining experiments, we use 5000 samples to evaluate the performance in terms of FID, FAD or FMD.

D IMPLEMENTATION DETAILS

We report in this section the implementation details for each benchmark. We used the same unified code-base for all the baselines, using the *PyTorch* framework. The VAE implementation is adapted from the official code whenever it’s available (MVAE, MMVAE and MOPOE as in ³, MVTCAE ⁴ and NEXUS⁵). For fairness, MLD and all the VAE-based models use the same autoencoder architecture. We use the best hyper-parameters suggested by the authors. Across all the data-sets, we use the *Adam optimizer* Kingma & Ba (2014) for training.

D.1 MLD

MLD uses the same autoencoders architecture used for VAE-based models, except that these are deterministic autoencoders. The autoencoders are trained using the same reconstruction loss term as for the VAE-based models. Table 8 and Table 9 summarize the hyper-parameters used during the two phases of MLD training. Note that for the image modality in the CUB dataset, to overcome over-fitting in training the deterministic autoencoder, data augmentation was necessary (we used *TrivialAugmentWide* from the Torchvision library).

Table 8: MLD: The deterministic autoencoders hyper-parameters

Dataset	Modality	Latent space	Batch size	Lr	Epochs	Weight decay
MNIST-SVHN	MNIST	16	128	1e-3	150	
	SVHN	64				
MHD	Image	64	64	1e-3	500	
	Trajectory	16				
	Sound	128				
POLYMNIST	All modalities	160	128	1e-3	300	
CUB	Caption	32	128	1e-3	500	1e-6
	Image	64		1e-4	300	
CelebAMask-HQ	Image	256	64	1e-3	200	
	Mask	128				
	Attributes	32				

Table 9: MLD: The score network hyper-parameters

Dataset	d	Blocks	Width	Time embed	Batch size	Lr	Epochs
MNIST-SVHN	0.5	2	512	256	128	1e-4	150
MHD	0.3	2	1024	512	128		3000
POLYMNIST	0.5	2	1536	512	256		3000
CUB	0.7	2	1024	512	64		3000
CelebAMask-HQ	0.5	2	1536	512	64		3000

D.2 VAE-BASED MODELS

For **MNIST-SVHN**, we follow Sutter et al. (2021); Shi et al. (2019) and use the same autoencoder architecture and pre-trained classifier. The latent space size is set to 20, $\beta = 5.0$. For MVTCAE $\alpha = \frac{5}{6}$. For both modalities, the likelihood is estimated using Laplace distribution. For NEXUS, we use the same modalities latent space sizes as in MLD, the joint NEXUS latent space is set to 20, $\beta_i = 1.0$ and $\beta_c = 5.0$. We train all the VAE-models for 150 epochs with 256 batch size and learning rate of $1e - 3$.

³<https://github.com/thomassutter/MoPoE>

⁴<https://github.com/gr8joo/MVTCAE>

⁵https://github.com/miguelsvasco/nexus_pytorch

For **MHD**, we reuse the autoencoders architecture and pre-trained classifier of Vasco et al. (2022). We adopt the hyper-parameters of Vasco et al. (2022) to train NEXUS model with the same settings, besides discarding the label modality. For the remaining VAE-based models, the latent space size is set to 128, $\beta = 1.0$ and $\alpha = \frac{5}{6}$ for MVTCAE. For all the modalities, Mean square error (MSE) is used to compute the reconstruction loss, similar to Vasco et al. (2022). These models are trained for 600 epochs with 128 batch size and learning rate of $1e - 3$.

For **POLYMNIST**, we use the same autoencoders architecture and pretrained classifier used by Sutter et al. (2021); Hwang et al. (2021). We set the latent space size to 512, $\beta = 2.5$ and $\alpha = \frac{5}{6}$ for MVTCAE. For all the modalities, the likelihood is estimated using Laplace distribution. For NEXUS, we use the same modality latent space size as in MLD, the joint NEXUS latent space to 64, $\beta_i = 1.0$ and $\beta_c = 2.5$. We train all the models for 300 epochs with 256 batch size and learning rate of $1e - 3$.

For **CUB**, we use the same autoencoders architecture and implementation settings as in Daunhawer et al. (2022). Laplace and one-hot categorical distributions are used to estimate likelihoods of the image and caption modalities respectively. The latent space size is set to 64, $\beta = 9.0$ for MVAE, MVTCAE and MOPOE and $\beta = 1$ for MMVAE. We set $\alpha = \frac{5}{6}$ for MVTCAE. For NEXUS, we use the same modalities latent space sizes as in MLD, the joint NEXUS latent space is set to 64, $\beta_i = 1.0$ and $\beta_c = 1$. We train all the models for 150 epochs with 64 batch size, with learning rate of $5e - 4$ for MVAE, MVTCAE and MOPOE and $1e - 3$ for the remaining models.

Finally, note that in the official implementation of Sutter et al. (2021) and Hwang et al. (2021), for the **POLYMNIST** and **MNIST-SVHN** data-sets, the classifiers were used for evaluation using dropout. In our implementation, we make sure to deactivate dropout during evaluation step.

D.3 MLD WITH POWERFULL AUTOENCODER

Here we provide more detail about the CUB experiment using more powerful autoencoder denoted MLD* in Figure 3. We use an architecture similar to Rombach et al. (2022) adapted to (64X64) resolution images. We modified the autoencoder architecture to be deterministic and train the model with a simple Mean square error loss. We kept the same configuration of the CUB experiment described in the previous experiment on the same dataset including the text autoencoder, score network and hyper-parameters. We also perform experiments with the same settings on (128X128) resolution images. We included the qualitative results in fig. 25.

D.4 COMPUTATION RESOURCES

In our experiments, we used 4 A100 GPUs, for a total of roughly 4 months of experiments.

E ADDITIONAL RESULTS

In this section, we report detailed results for all of our experiments, including standard deviation and additional qualitative samples, for all the data-sets and all the methods we compared in our work.

E.1 MNIST-SVHN

E.1.1 SELF RECONSTRUCTION

In Table 10 we report results about *self-coherence*, which we use to support the arguments from Section 2. This metric is used to measure the loss of information due to latent collapse, by showing the ability of all competing models to reconstruct an arbitrary modality given the same modality or a set thereof as an input. For our MLD model, the self-reconstruction is done without using the diffusion model component: the modality is encoded using its deterministic encoder and the decoder is fed with the latent space to get the reconstruction.

We observe that VAE based models fail at reconstructing SVHN given SVHN. This is especially more visible for product of experts based models (MVAE and MVTCAE). In MLD, the deterministic autoencoders do not suffer from such weakness and achieve overall the best performance.

Figure 13 shows qualitative results for the self-generation. We remark that some samples generated using VAE-based models, the digits differs from the ones in the input sample, indicating information loss due to the latent collapse. For example, in the case of MVAE, generation of the MNIST digit 3, in MVTCAE generation of the SVHN digit 2.

Table 10: Self-generation coherence and quality for **MNIST-SVHN** (M :MNIST, S: SVHN). The generation quality is measured in terms of FMD for MNIST and FID for SVHN.

Models	Coherence (% \uparrow)				Quality (\downarrow)			
	M \rightarrow M	M,S \rightarrow M	S \rightarrow S	M,S \rightarrow M	M \rightarrow M	M,S \rightarrow M	S \rightarrow S	M,S \rightarrow M
MVAE	86.92 \pm 0.8	88.03 \pm 0.78	40.62 \pm 0.99	68.01 \pm 1.29	10.75 \pm 1.04	10.79 \pm 1.02	60.22 \pm 1.01	59.0 \pm 0.6
MMVAE	87.22 \pm 1.87	77.35 \pm 4.19	67.31 \pm 6.93	39.44 \pm 3.43	12.15 \pm 1.25	20.24 \pm 1.04	58.1 \pm 3.14	171.42 \pm 4.55
MOPOE	89.95 \pm 0.84	91.71 \pm 0.77	67.26 \pm 0.8	83.58 \pm 0.44	9.39 \pm 0.76	10.1 \pm 0.73	53.19 \pm 1.06	57.34 \pm 1.35
NEXUS	92.63 \pm 0.45	93.59 \pm 0.4	68.31 \pm 0.46	83.13 \pm 0.58	4.92 \pm 0.61	5.16 \pm 0.59	85.67 \pm 2.74	97.86 \pm 2.86
MVTCAE	94.33 \pm 0.18	95.18 \pm 0.19	47.47 \pm 0.76	86.6 \pm 0.23	4.67 \pm 0.35	4.94 \pm 0.37	52.29 \pm 1.17	53.55 \pm 1.19
MLD	96.73 \pm 0.0	96.73 \pm 0.0	82.19 \pm 0.0	82.19 \pm 0.0	2.25 \pm 0.03	2.25 \pm 0.03	48.47 \pm 0.63	48.47 \pm 0.63

E.1.2 DETAILED RESULTS

Table 11: Generative Coherence for **MNIST-SVHN**. We report the detailed version of Table 1 with standard deviation for 5 independent runs with different seeds.

Models	Coherence (% \uparrow)			Quality (\downarrow)			
	Joint	M \rightarrow S	S \rightarrow M	Joint(M)	Joint(S)	M \rightarrow S	S \rightarrow M
MVAE	38.19 \pm 2.27	48.21 \pm 2.56	28.57 \pm 1.46	13.34 \pm 0.93	68.0 \pm 0.99	68.9 \pm 1.84	13.66 \pm 0.95
MMVAE	37.82 \pm 1.19	11.72 \pm 0.33	67.55 \pm 9.22	25.89 \pm 0.46	146.82 \pm 4.76	393.33 \pm 4.86	53.37 \pm 1.87
MOPOE	39.93 \pm 1.54	12.27 \pm 0.68	68.82 \pm 0.39	20.11 \pm 0.96	129.2 \pm 6.33	373.73 \pm 26.42	43.34 \pm 1.72
NEXUS	40.0 \pm 2.74	16.68 \pm 5.93	70.67 \pm 0.77	13.84 \pm 1.41	98.13 \pm 5.9	281.28 \pm 16.07	53.41 \pm 1.54
MVTCAE	48.78 \pm 1	81.97 \pm 0.32	49.78 \pm 0.88	12.98 \pm 0.68	52.92 \pm 1.39	69.48 \pm 1.64	13.55 \pm 0.8
MMVAE+	17.64 \pm 4.12	13.23 \pm 4.96	29.69 \pm 5.08	26.60 \pm 2.58	121.77 \pm 37.77	240.90 \pm 85.74	35.11 \pm 4.25
MMVAE+(K=10)	41.59 \pm 4.89	55.3 \pm 9.89	56.41 \pm 5.37	19.05 \pm 1.10	67.13 \pm 4.58	75.9 \pm 12.91	18.16 \pm 2.20
MLD	85.22 \pm 0.5	83.79 \pm 0.62	79.13 \pm 0.38	3.93 \pm 0.12	56.36 \pm 1.63	57.2 \pm 1.47	3.67 \pm 0.14




Figure 13: Self-generation qualitative results for MNIST-SVHN. For each model we report: MNIST to MNIST conditional generation in the left, SVHN to SVHN conditional generation in the right.

E.2 MHD

Table 12: Generative Coherence for MHD. We report the detailed version of Table 2 with standard deviation for 5 independent runs with different seeds.

Models	Joint	I (Image)			T (Trajectory)			S (Sound)		
		T	S	T,S	I	S	I,S	I	T	I,T
MVAE	37.77 \pm 3.32	11.68 \pm 0.35	26.46 \pm 1.84	28.4 \pm 1.47	95.55 \pm 1.39	26.66 \pm 1.72	96.58 \pm 1.06	58.87 \pm 4.89	10.39 \pm 0.42	58.16 \pm 5.24
MMVAE	34.78 \pm 0.83	99.7 \pm 0.03	69.69 \pm 1.66	84.74 \pm 0.95	<u>99.3</u> \pm 0.07	85.46 \pm 1.57	92.39 \pm 0.95	49.95 \pm 0.79	50.14 \pm 0.89	50.17 \pm 0.99
MOPOE	48.84 \pm 0.36	<u>99.64</u> \pm 0.08	68.67 \pm 2.07	<u>99.69</u> \pm 0.04	99.28 \pm 0.08	<u>87.42</u> \pm 0.41	99.35 \pm 0.04	50.73 \pm 3.72	51.5 \pm 3.52	56.97 \pm 6.34
NEXUS	26.56 \pm 1.71	94.58 \pm 0.34	<u>83.1</u> \pm 0.74	95.27 \pm 0.52	88.51 \pm 0.64	76.82 \pm 3.63	93.27 \pm 0.91	70.06 \pm 2.83	75.84 \pm 2.53	89.48 \pm 3.24
MVTCAE	42.28 \pm 1.12	99.54 \pm 0.07	72.05 \pm 0.95	99.63 \pm 0.05	99.22 \pm 0.08	72.03 \pm 0.48	<u>99.39</u> \pm 0.02	<u>92.58</u> \pm 0.47	<u>93.07</u> \pm 0.36	<u>94.78</u> \pm 0.25
MMVAE+	41.67 \pm 2.3	98.05 \pm 0.19	84.16 \pm 0.57	91.88 \pm	97.47 \pm 0.89	81.16 \pm 2.24	89.31 \pm 1.54	64.34 \pm 4.46	65.42 \pm 5.42	64.88 \pm 4.93
MMVAE+(k=10)	42.60 \pm 2.5	99.44 \pm 0.07	89.75 \pm 0.75	94.7 \pm 0.72	99.44 \pm 0.18	89.58 \pm 0.4	95.01 \pm 0.30	87.15 \pm 2.81	87.99 \pm 2.55	87.57 \pm 2.09
MLD	98.34 \pm 0.22	99.45 \pm 0.09	<u>88.91</u> \pm 0.54	99.88 \pm 0.04	99.58 \pm 0.03	<u>88.92</u> \pm 0.53	99.91 \pm 0.02	97.63 \pm 0.14	97.7 \pm 0.34	98.01 \pm 0.21

Table 13: Generative quality for MHD. We report the detailed version of Table 3 with standard deviation for 5 independent runs with different seeds.

Models	I (Image)				T (Trajectory)				S (Sound)			
	Joint	T	S	T,S	Joint	I	S	I,S	Joint	I	T	I,T
MVAE	94.9 \pm 7.37	93.73 \pm 5.44	92.55 \pm 7.37	91.08 \pm 10.24	39.51 \pm 6.04	20.42 \pm 4.42	38.77 \pm 6.29	19.25 \pm 4.26	14.14 \pm 0.25	14.13 \pm 0.19	14.08 \pm 0.24	14.17 \pm 4.26
MMVAE	224.01 \pm 12.58	22.6 \pm 1.3	789.12 \pm 12.58	170.41 \pm 8.06	16.52 \pm 1.17	0.5 \pm 0.05	30.39 \pm 1.38	6.07 \pm 0.37	22.8 \pm 0.39	22.61 \pm 0.75	23.72 \pm 0.86	23.01 \pm 0.67
MOPOE	147.81 \pm 10.37	16.29 \pm 0.85	838.38 \pm 10.84	15.89 \pm 1.96	<u>13.92</u> \pm 0.96	0.52 \pm 0.12	33.38 \pm 1.14	0.53 \pm 0.1	18.53 \pm 0.27	24.11 \pm 0.4	24.1 \pm 0.41	23.93 \pm 0.87
NEXUS	281.76 \pm 12.69	116.65 \pm 9.99	282.34 \pm 12.69	117.24 \pm 8.53	18.59 \pm 2.16	6.67 \pm 0.23	33.01 \pm 3.41	7.54 \pm 0.29	<u>13.99</u> \pm 0.9	19.52 \pm 0.14	18.71 \pm 0.24	16.3 \pm 0.59
MVTCAE	121.85 \pm 3.44	5.34 \pm 0.33	54.57 \pm 7.79	3.16 \pm 0.26	19.49 \pm 0.67	0.62 \pm 0.1	13.65 \pm 1.24	0.75 \pm 0.13	15.88 \pm 0.19	14.22 \pm 0.27	14.02 \pm 0.14	13.96 \pm 0.28
MMVAE+	97.19 \pm 12.37	2.80 \pm 0.42	128.56 \pm 4.47	114.3 \pm 11.4	22.37 \pm 1.87	1.21 \pm 0.22	21.74 \pm 3.49	15.2 \pm 1.15	16.12 \pm 0.40	17.31 \pm 0.62	17.92 \pm 0.19	17.56 \pm 0.48
MMVAE+(K=10)	85.98 \pm 1.25	1.83 \pm 0.26	70.72 \pm 1.76	62.43 \pm 3.4	21.10 \pm 1.25	1.38 \pm 0.34	8.52 \pm 0.79	7.22 \pm 1.6	14.58 \pm 0.47	14.33 \pm 0.51	14.34 \pm 0.42	14.32 \pm 0.6
MLD (ours)	7.98 \pm 1.41	1.7 \pm 0.14	4.54 \pm 0.45	1.84 \pm 0.27	3.18 \pm 0.18	0.83 \pm 0.03	2.07 \pm 0.26	<u>0.6</u> \pm 0.05	2.39 \pm 0.1	2.31 \pm 0.07	2.33 \pm 0.11	2.29 \pm 0.06

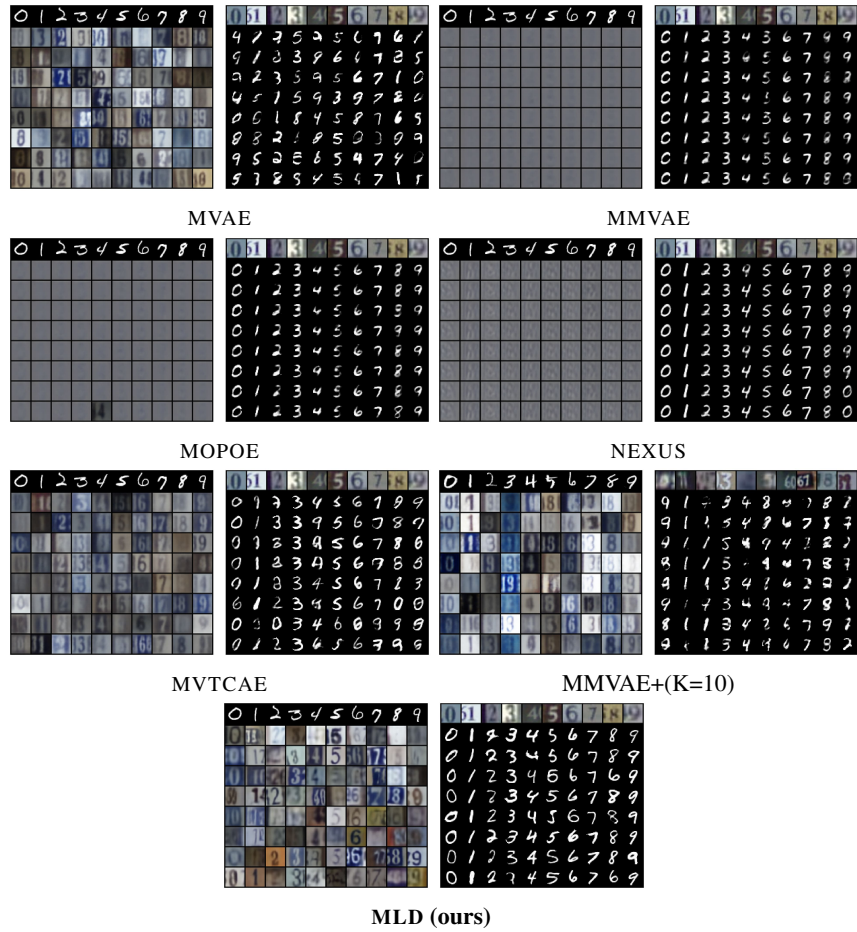
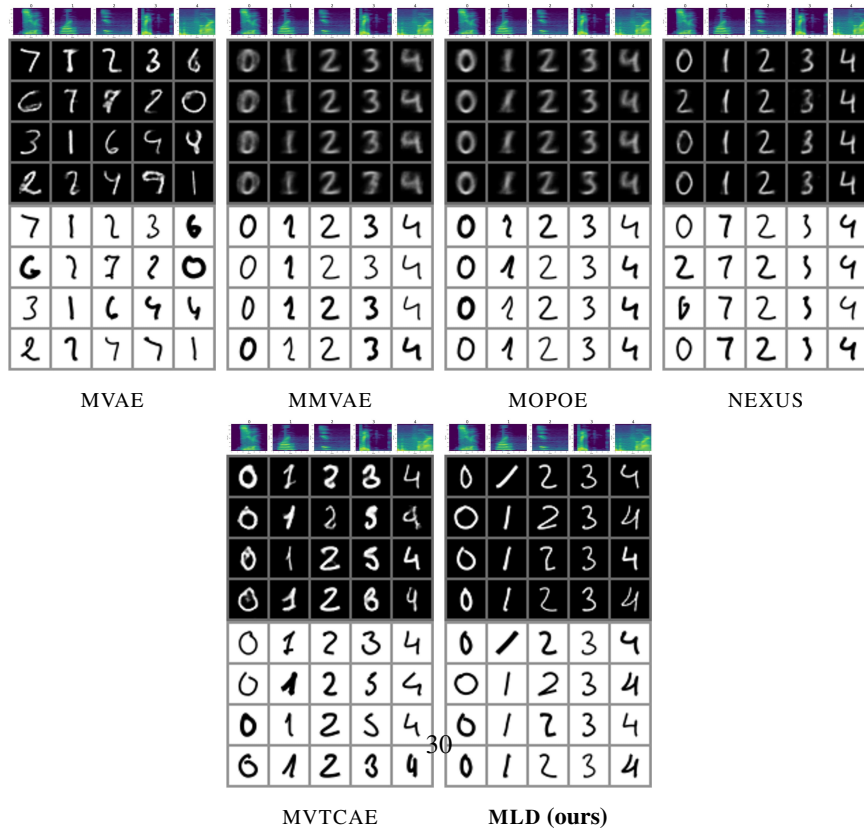


Figure 14: Additional qualitative results for MNIST-SVHN. For each model we report: MNIST to SVHN conditional generation in the left, SVHN to MNIST conditional generation in the right.



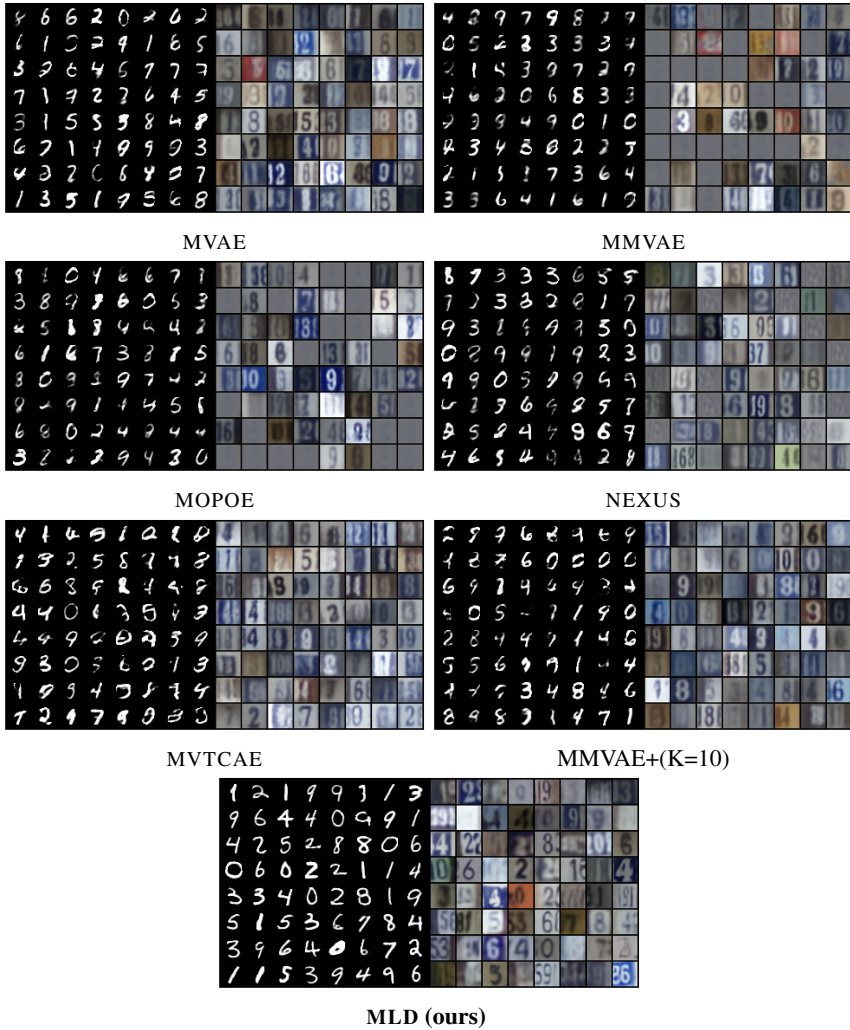


Figure 15: Qualitative results for MNIST-SVHN joint generation.

E.3 POLYMNIST

Table 14: Generation Coherence (%) for POLYMNIST (Higher is better) used for the plots in Figure 2 and Figure 9. We report the average *leave one out coherence* as a function of the number of observed modalities. *Joint* refers to random generation of the 5 modalities simultaneously.

Models	Coherence (%↑)				
	Joint	1	2	3	4
MVAE	4.0±1.49	37.51±3.16	48.06±3.55	53.19±3.37	56.09±3.31
MMVAE	25.8±1.43	75.15 ±2.54	75.14±2.47	75.09±2.6	75.09±2.58
MOPOE	17.32±2.47	69.37±1.85	81.29±2.34	85.26±2.36	86.7±2.39
NEXUS	18.24±0.89	60.61±2.51	72.14±2.79	76.81±2.75	78.92±2.64
MVTCAE	0.21±0.05	57.66±1.06	78.44±1.31	85.97 ±1.43	88.81±1.49
MMVAE+	26.28±2.19	54.74±0.5	54.06±0.33	55.2±1.32	53.17±0.75
MMVAE+(K=10)	14.53±4.94	58.93±6.3	59.42±8.8	60.77±8.03	58.24±7.42
MLD IN-PAINT	51.65±1.16	52.85±0.23	77.65±0.24	85.66±0.43	87.29±0.29
MLD UNI	48.79±0.43	65.12±0.7	79.52±0.8	82.03±1.19	81.86±2.09
MLD	56.23 ±0.52	68.58±0.72	84.87 ±0.19	88.56 ±0.12	89.43 ±0.27




Figure 16: Joint generation qualitative results for **MHD**. The three modalities are randomly generated simultaneously (**Top row**: image, **Middle row**: trajectory vector converted into image, **Bottom row**: sound Mel-Spectrogram).

Table 15: Generation quality (FID ↓) for **POLYMNIST** (lower is better) used for the plots in Figure 2 and Figure 9. Similar to Table 14, we report the average *leave one out* FID as a function of the number of observed modalities. *Joint* refers to random generation quality of the 5 modalities simultaneously.

Models	Quality (↓)				
	Joint	1	2	3	4
MVAE	108.74 \pm 2.73	108.06 \pm 2.79	108.05 \pm 2.73	108.14 \pm 2.71	108.18 \pm 2.85
MMVAE	165.74 \pm 5.4	208.16 \pm 10.41	207.5 \pm 10.57	207.35 \pm 10.59	207.38 \pm 10.58
MOPOE	113.77 \pm 1.62	173.87 \pm 7.34	185.06 \pm 10.21	191.72 \pm 11.26	196.17 \pm 11.66
NEXUS	91.66 \pm 2.93	207.14 \pm 7.71	205.54 \pm 8.6	204.46 \pm 9.08	202.43 \pm 9.49
MVTCAE	106.55 \pm 3.83	78.3 \pm 2.35	85.55 \pm 2.51	92.73 \pm 2.65	99.13 \pm 2.72
MMVAE+	168.88 \pm 0.12	165.67 \pm 0.14	166.5 \pm 0.18	165.53 \pm 0.55	165.3 \pm 0.33
MMVAE+ (K=10)	156.55 \pm 3.58	154.42 \pm 2.73	153.1 \pm 3.01	153.06 \pm 2.88	154.9 \pm 2.9
MLD IN-PAINT	64.78 \pm 0.33	65.41 \pm 0.43	65.42 \pm 0.41	65.52 \pm 0.46	65.55 \pm 0.46
MLD UNI	62.42 \pm 0.62	63.16 \pm 0.81	64.09 \pm 1.15	65.17 \pm 1.46	66.46 \pm 2.18
MLD	63.05 \pm 0.26	62.89 \pm 0.2	62.53 \pm 0.21	62.22 \pm 0.39	61.94 \pm 0.65

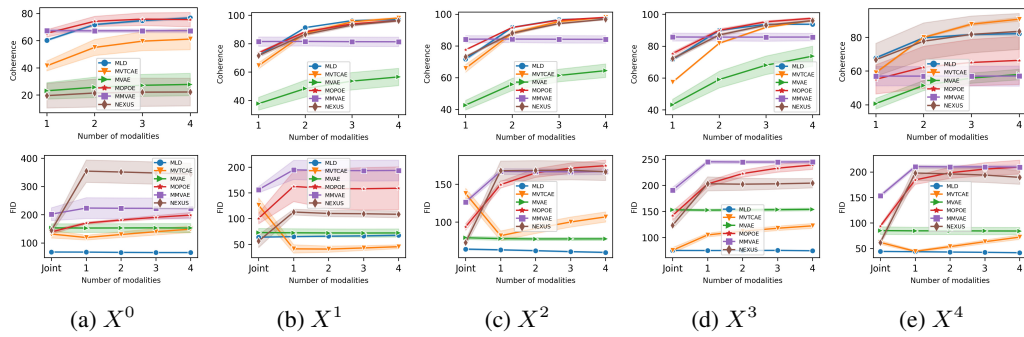


Figure 18: **Top:** Generation Coherence (%) for **POLYMNIST** (Higher is better). **Bottom:** Generation quality (FID) (Lower is better). We report the average *leave one out* performance as a function of the number of observed modalities for each modality X^i . *Joint* refers to random generation of the 5 modalities simultaneously.

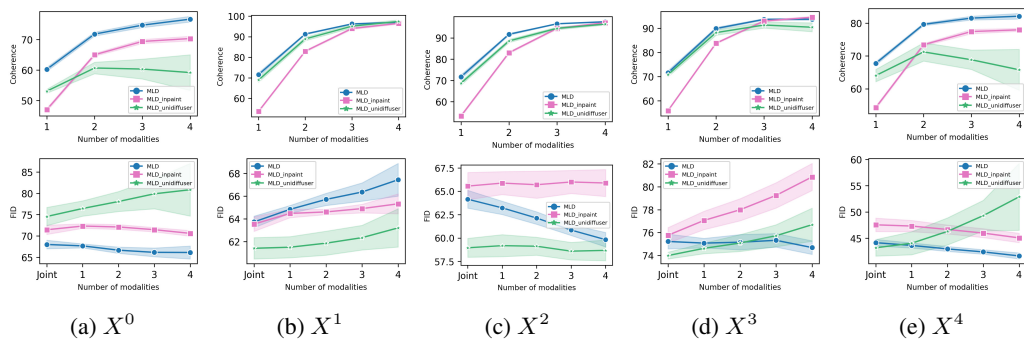


Figure 19: **Top:** Generation Coherence (%) for **POLYMNIST** (Higher is better). **Bottom:** Generation quality (FID) (Lower is better). We report the average *leave one out* performance as a function of the number of observed modalities for each modality X^i . *Joint* refers to random generation of the 5 modalities simultaneously.

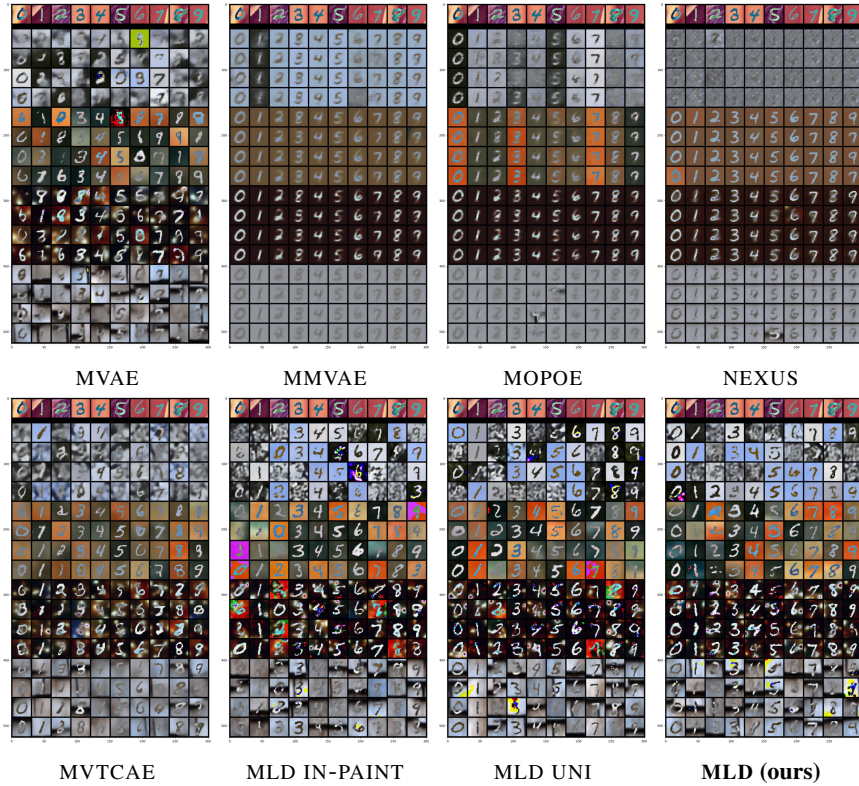


Figure 20: Conditional generation qualitative results for **POLYMNIST**. The modality X^2 (dirst row) is used as the condition to generate the 4 remaining modalities(The rows below).

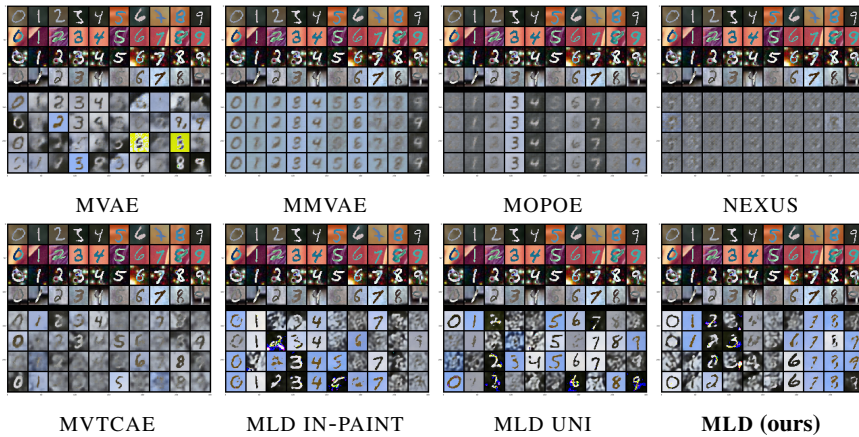


Figure 21: Conditional generation qualitative results for **POLYMNIST**. The subset of modalities X^1, X^2, X^3, X^4 (First 4 rows) are used as the condition to generate the modality X^0 (The rows below).

E.3.1 ADDITIONAL EXPERIMENTS WITH PALUMBO ET AL. (2023) ARCHITECTURE

In our experiments on POLYMNIST, we used the same architecture as in Sutter et al. (2021) Hwang et al. (2021) to ensure a fair settings for all the baselines. In Palumbo et al. (2023), the experiments on POLYMNIST are conducted using a different autoencoder architecture which is based on Resnets instead of a a sequence of convolutions layers based autoencoder. We investigate in this section, the performance of MMVAE+ and our MLD using this architecture. For MMVAE+, we keep the same settings as in Palumbo et al. (2023) including the autoencoder architecture, latent size, and importance sampling $K=10$ with doubly reparameterized gradient estimator (DReG). For MLD, we use the same autoencoder architecture with latent size equal to 160. In Appendix E.3.1, we observe that while the new architecture autoencoder enhance the MMVAE+ performance, our MLD performance is improved as well. Similarly to previous results, MLD achieves simultaneously the best generative coherence and quality.




Figure 22: Results for POLYMNIST data-set. Left: a comparison of the generative coherence (↑) and quality in terms of FID (↓) as a function of the number of inputs.

E.4 CUB

Models	Coherence (↑)			Quality (↓)	
	Joint	Image → Caption	Caption → Image	Joint → Image	Caption → Image
MVAE	0.66	0.70	0.64	158.91	158.88
MMVAE	0.66	0.69	0.62	277.8	212.57
MOPOE	0.64	0.68	0.55	279.78	179.04
NEXUS	0.65	0.69	0.59	147.96	262.9
MVTCAE	0.65	0.70	0.65	155.75	168.17
MMVAE+	0.61	0.68	0.65	188.63	247.44
MMVAE+(K=10)	0.63	0.68	0.62	172.21	178.88
MLD IN-PAINT	0.69	0.69	0.68	69.16	68.33
MLD UNI	0.69	0.69	0.69	64.09	61.92
MLD	0.69	0.69	0.69	63.47	62.62
MLD*	0.70	0.69	0.69	22.19	22.50

Table 16: Generation Coherence (CLIP-S : Higher is better) and Quality (FID ↓ Lower is better) for CUB dataset. **MLD*** denotes the version of our method using a more powerful image autoencoder.



Figure 23: Qualitative results for joint generation on CUB.

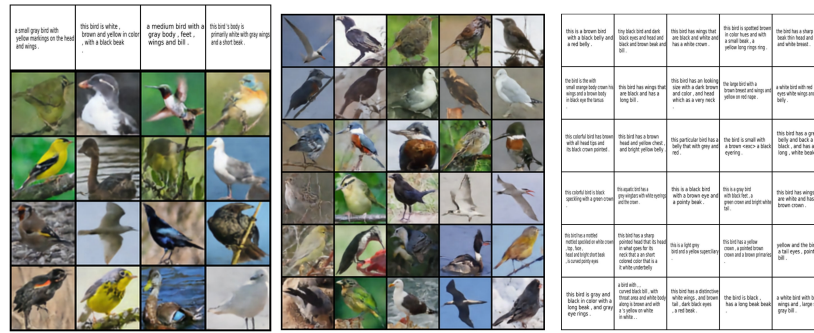


Figure 24: Qualitative results of MLD* on CUB data-set with powerful image autoencoder.

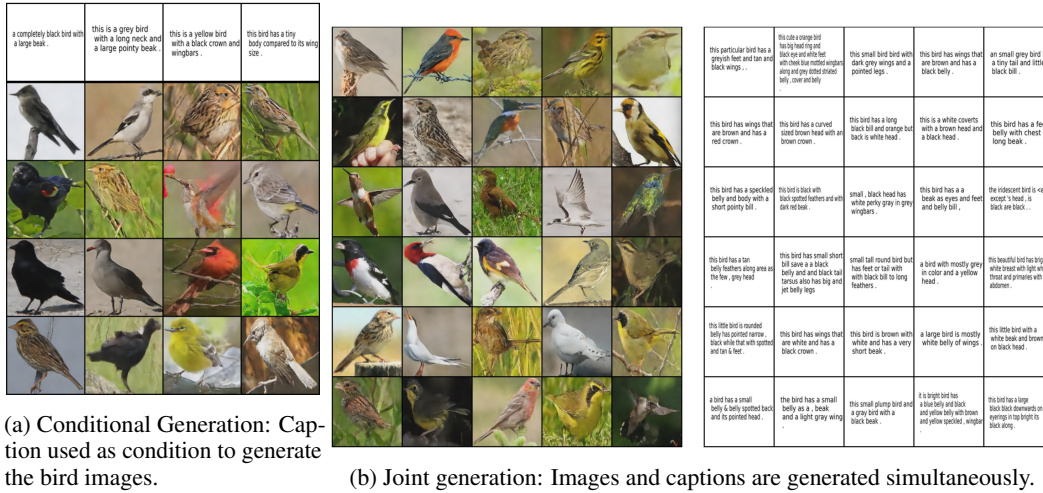


Figure 25: Qualitative results of MLD* on CUB data-set with 128x128 resolution with powerful image autoencoder.

E.5 CELEBAMASK-HQ

In this section, we present additional experiments on the CelebAMask-HQ dataset (Lee et al., 2019a), which consists of face images, each having a segmentation mask, and text attributes, so 3 modalities. We follow the same experimentation protocol as in (Wesego & Rooshenas, 2023) including the autoencoder base architecture. Note that for MLD we use deterministic autoencoders instead of variational autoencoders (Lee et al., 2019a). Similarly, the CelebAMask-HQ dataset is restricted to take into account 18 out of 40 attributes from the original dataset and the images are resized to 128×128 resolution, as done in (Wu & Goodman, 2018; Wesego & Rooshenas, 2023). Please refer to Appendix D, for additional implementation details of MLD.

The Image generation quality is evaluated in terms of FID score. The attributes and the mask having binary values, are evaluated in terms of $F1$ Score against the ground truth. The competitors performance results are reported from Wesego & Rooshenas (2023).

The quantitative results in Table 17 show that MLD outperforms the competitors on the generation quality. It achieves the best $F1$ score in the attributes generation given Image and Mask modalities. The mask generation best performance is achieved by MOPOE. MLD achieves the second best performance on the mask generation conditioning on both image and attributes modalities. Overall, MLD stands out with the best image quality generation while being on-par with competitors in the mask and attribute generation coherence.



(a) Image

(b) Mask

Heavy_Makeup, Mouth_Slightly_Open, Smiling, Wavy_Hair	Bushy_Eyebrows, Male, Wavy_Hair	Heavy_Makeup, Mouth_Slightly_Open	Brown_Hair, Heavy_Makeup, Wavy_Hair	Brown_Hair, Wavy_Hair
Brown_Hair, Heavy_Makeup, Mouth_Slightly_Open, Smiling	Black_Hair, Male	Heavy_Makeup, Smiling	Brown_Hair, Male, Mouth_Slightly_Open, Smiling, Straight_Hair	Heavy_Makeup, Mouth_Slightly_Open, Smiling, Wavy_Hair
Male, Straight_Hair	Bangs, Brown_Hair, Heavy_Makeup, Mouth_Slightly_Open, Smiling, Wavy_Hair	Bangs, Blond_Hair, Heavy_Makeup, Mouth_Slightly_Open, Smiling, Wavy_Hair	Heavy_Makeup, Mouth_Slightly_Open, Smiling, Wavy_Hair	Blond_Hair, Heavy_Makeup, Mouth_Slightly_Open, Wavy_Hair
Heavy_Makeup, Mouth_Slightly_Open, Smiling	Brown_Hair, Heavy_Makeup, Mouth_Slightly_Open, Smiling, Wavy_Hair	Heavy_Makeup, Mouth_Slightly_Open, Smiling	Brown_Hair, Heavy_Makeup, Straight_Hair	Brown_Hair, Wavy_Hair
Bangs, Brown_Hair, Heavy_Makeup, Smiling	Black_Hair, Male, Mouth_Slightly_Open, Smiling	Blond_Hair, Heavy_Makeup, Mouth_Slightly_Open, Smiling, Wavy_Hair	Heavy_Makeup, Wavy_Hair	Brown_Hair, Male, Straight_Hair

(c) Attributes

Figure 26: Joint (Unconditional) generation qualitative results of MLD on CelebAMask-HQ.

Models	Attributes			Image			Mask	
	Img + Mask F1	Img F1	Att + Mask FID	Mask FID	Att FID	Joint FID	Img+Att F1	Img F1
Weseago & Rooshenas (2023)								
SBM-RAE	0.62	0.6	84.9	86.4	85.6	84.2	0.83	0.82
SBM-RAE-C	0.66	0.64	83.6	82.8	83.1	84.2	0.83	0.82
SBM-VAE	0.62	0.58	81.6	81.9	78.7	79.1	0.83	0.83
SBM-VAE-C	0.69	0.66	82.4	81.7	76.3	79.1	0.84	0.84
MOPOE	0.68	0.71	114.9	101.1	186.8	164.8	0.85	0.92
MVTCAE	0.71	0.69	94	84.2	87.2	162.2	0.89	0.89
MMVAE+	0.64	0.61	133	97.3	153	103.7	0.82	0.89
Supervised classifier		0.79					0.94	
MLD (ours)	0.72	0.69	52.75	51.73	53.09	54.27	0.87	0.87

Table 17: Quantitative results on CelebAMask-HQ dataset. The performance is measured in terms of FID (\downarrow) and F1 score (\uparrow). The first row is the generated modality while the second row is the modalities used as condition. Supervised classifier designates a classifier performance to predict the attributes or the mask from an image.



(a) Generated Images

Figure 27: (Attributes \rightarrow Image) Conditional generation of MLD on CelebAMask-HQ.



(a) Generated Images

Figure 28: (Attributes,Mask \rightarrow Image) Conditional generation of MLD on CelebAMask-HQ.



(a) Generated Images

Figure 29: (Mask \rightarrow Image) Conditional generation of **MLD** on CelebAMask-HQ.

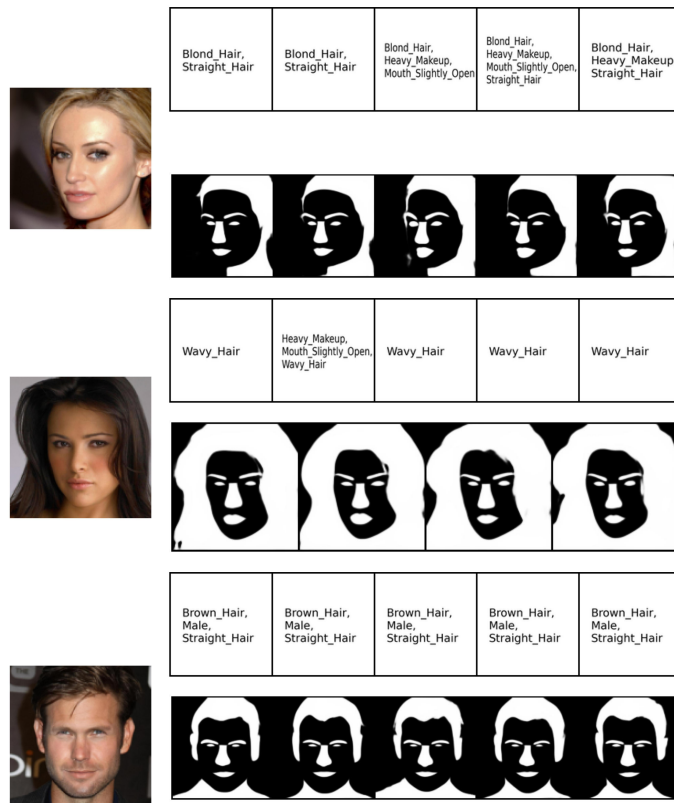


Figure 30: (Image \rightarrow Attribute,Mask) Conditional generation of **MLD** on CelebAMask-HQ.